

Copyright is owned by the Author of the thesis. Permission is given for a copy to be downloaded by an individual for the purpose of research and private study only. The thesis may not be reproduced elsewhere without the permission of the Author.

COORDINATION IN VEHICLE ROUTING

A THESIS PRESENTED IN PARTIAL FULFILMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN
OPERATIONS RESEARCH
AT
MASSEY UNIVERSITY, PALMERSTON NORTH
NEW ZEALAND

CATHERINE MARGARET RIVERS

2002

ABSTRACT

Coordination involves the re-deployment of payload between depot and customer, and includes split deliveries, load transfers and load swapping, facilitated by the establishment of coordination sites at strategic locations. Real-world coordination includes mid-air refuelling, the use of temporary replenishment sites, trailers left for later uplift by their towing vehicles, bulk re-suppliers travelling to field operatives, fleet re-supply, and couriers swapping loads on the side of the road.

This thesis models the coordination process and investigates the basic types of coordination in single depot, pure delivery systems in both the Euclidean plane and the rectilinear grid network. Strategies are developed for dealing with dynamic situations in the rectilinear grid, which are based on the pre-processing of scenarios in order that dispatchers may select a suitable response from an existing selection at the time that dynamic values are revealed.

In addition, a procedure is suggested that reduces the number of edges and vertices of a rectilinear grid to those that may be useful within a coordination hull.

ACKNOWLEDGEMENTS

I am grateful to Lieutenant Colonel Janet L. Castell, MNZM, and Captain Greg Davies for early discussions on the application of logistics within the New Zealand Army.

Professor Don Barnes provided valuable advice about mobile coordination at sea.

Roading contractors, Higgins Contractors Ltd, gave me practical on-site experience of coordination in an industrial setting.

Mrs Marlene Gordon, and her team at Marton Coachlines Ltd, gave me the opportunity to investigate coordination in the passenger transport area, and I am grateful for the cooperation of their numerous passengers who participated with enthusiasm.

I thank the myriad courier drivers who were patient enough to listen to my theories and kind enough to make suggestions and to allow me to analyse their work practices.

Major Rio M. Thalieb, Indonesian Air Force, Wing Commander David Huang, Taiwanese Air Force, and Flight Lieutenant Meta Suntornsaratoon, Royal Thai Air Force, all attached to the Australian Defence Force Academy, gave continued support and critical appraisal, especially in comparing my procedures with evolutionary approaches.

I acknowledge with thanks the supervision of this research by Dr John W. Giffin, Massey University, and Dr Mahyar A. Amouzegar, Rand Corporation, USA.

TABLE OF CONTENTS

List of Figures	xv
List of Tables	xxiii
List of Procedures	xxxi
Glossary	xxxiii
1 Introduction	1
2 Literature Review	5
2.1 The vehicle routing problem	5
2.2 Coordination in vehicle routing problems	14
2.2.1 Split deliveries	15
2.2.2 Coordinated public transport	18
2.2.3 Mid-air refuelling	19
2.2.4 Pre-emptive routing	22
2.2.5 Inventory routing with satellite facilities	22
3 Aspects of Coordination	23
3.1 Predicted operational benefits	24
3.2 Payload requirements	26
3.2.1 Compatibility	27
3.2.2 Volume	28
3.2.3 Mass	28
3.2.4 Divisibility	28
3.2.5 Need to return to the depot	28
3.2.6 Need for specific vehicle type	29

3.2.7	Social acceptability	29
3.2.8	Legal acceptability	29
3.2.9	Support vehicles	30
3.2.10	En route processing	30
3.2.11	Urgency	30
3.2.12	Commodity classes in this study	30
3.3	Coordination sites	31
3.3.1	Satellite facility	33
3.3.2	Node	34
3.3.2.1	Sequential multiple vehicle	34
3.3.2.2	Simultaneous multiple vehicle	36
3.3.2.3	Single vehicle	37
3.3.3	Existing route	38
3.3.3.1	Chance meeting on an existing route (serendipity) .	38
3.3.3.2	Pre-determined meeting on existing routes	38
3.3.3.3	Post-determined meeting on existing routes	39
3.3.3.4	Repair	39
3.3.4	Not on existing route	39
3.3.4.1	Pre-determined meeting involving detour	39
3.3.4.2	Post-determined meeting involving detour	40
3.3.5	Mobile	41
3.3.5.1	Spatially mobile	41
3.3.5.2	Temporally mobile	41
3.3.6	Driver-defined	41
3.3.6.1	Single vehicle driver-defined	41
3.3.6.2	Multiple vehicle driver-defined	42
3.3.7	Single vehicle	42
3.3.8	Direction of transfer	43
3.3.8.1	One-way transfer	43
3.3.8.2	Two-way transfer	43
3.3.8.3	Many-way transfer	44
3.4	Classification	45
3.5	The coordination model	46
3.6	Performance indicators	47
3.7	Spatial and temporal inter-dependence	47
3.8	Route order and orientation	48
3.9	Existence and identity of a suitable node	53
3.9.1	Node-based approach	53

3.9.2	Route-based approach	54
3.9.3	Intersection-based approach	59
3.9.4	Limitations	59
4	Coordination in the Euclidean Plane	61
4.1	Node location	61
4.2	Euclidean distances	62
4.3	Depot location	62
4.4	Customer demand	63
4.5	Delivery times	63
4.6	Coordination opportunities in the Euclidean plane	64
4.7	Model construction	66
4.7.1	Initial data	67
4.7.2	Route construction	67
4.7.2.1	Route improvement by intra-route 3-opt exchange .	68
4.7.2.2	Route improvement by inter-route node transfer	69
4.7.2.3	Route improvement by inter-route node exchange ..	70
4.7.2.4	Route improvement by inter-route path exchange ..	72
4.7.2.5	Route improvement by split delivery	77
4.7.3	Route scheduling	89
4.7.4	Incorporation of coordination	89
4.8	Size of near radius	91
4.9	Orientation and route order	94
4.10	Number of nodes	97
4.11	Problem design	98
4.12	Load swapping	103
4.12.1	Reduction in total distance	103
4.12.2	Reduction in fleet size	110
4.12.3	Transfer	115
4.13	Obviating return to depot	119
4.13.1	Single-noded routes	119
4.14	Coordination site as interim step	125
4.14.1	Cascading interim steps	134
4.14.2	Stability of intersections	140
4.14.3	Limitations and advantages of coordination as an interim step	142
4.15	Appraisal of Euclidean coordination	143

5	Coordination in the Rectilinear Grid Network	149
5.1	Introduction	149
5.2	Grid construction and definition	150
5.3	Node location on the grid	151
5.4	Grid network distances	152
5.5	Depot location	161
5.6	Connectedness of the grid	163
5.7	Grid reduction	164
5.7.1	Vertex classification	165
5.7.2	Reduction procedures	168
5.7.3	Effect on path proximity	180
5.7.4	Uniqueness of reduction	181
5.7.5	Effect of node location	183
5.8	Coordination opportunities	187
5.9	Solution construction	188
5.9.1	Initial data	188
5.9.2	Route construction	188
5.9.2.1	Inter-route node transfer	191
5.9.2.2	Inter-route node exchange	192
5.9.2.3	Intra-route edge exchange	193
5.9.2.4	Inter-route edge exchange	195
5.9.3	Route scheduling	195
5.9.4	Incorporation of coordination	198
5.10	Maximum distance constraint	202
5.11	Maximum vehicle capacity constraint	204
5.11.1	Coordination by methods other than split delivery	204
5.11.2	Coordination by split delivery	207
5.11.3	Mixed systems with common demand	209
5.12	Problem design	211
5.13	Load swapping	212
5.14	Deferral of return to depot	219
5.15	Split delivery	221
5.15.1	Reducing path duplication	221
5.15.2	Fleet reduction	223
5.16	Alternative coordination sites	227
5.17	Grid density	232
5.17.1	Medium grid	234
5.17.1	Sparse grid	240

5.17.3	Dense grid	241
5.17.4	Comparison of splitting with other forms of coordination .	243
5.18	Order of application of coordination	244
5.19	Quality of initial solution	248
5.20	Appraisal of rectilinear coordination	252

6 An Extension – Dynamic Coordination in the Rectilinear Grid Network 255

6.1	Introduction	255
6.1.1	Dynamic and stochastic vehicle routing	255
6.1.2	Dynamic coordination	259
6.2	Problem description	264
6.3	Nomenclature	266
6.4	DC1 Only time is known	267
6.4.1	Time constraint	269
6.4.2	Multiple placement	269
6.4.3	Problem generation	270
6.4.4	Incorporation of coordination	271
6.4.4.1	Routing scheme A	272
6.4.4.2	Routing scheme B	273
6.4.4.3	Routing scheme C	273
6.4.4.4	Routing scheme D	274
6.4.4.5	Routing scheme E	275
6.4.4.6	Routing scheme F	275
6.4.4.7	Routing strategy for DC1	276
6.4.5	Conclusions	277
6.5	DC2 Only demand is unknown	278
6.5.1	Problem generation	279
6.5.2	Problem solution	279
6.5.3	Demand estimation	280
6.5.4	Initial route construction	280
6.5.5	Coordination site location	282
6.5.6	Incorporation of coordination	284
6.5.6.1	Routing scheme A	285
6.5.6.2	Routing scheme B	285
6.5.6.3	Routing scheme C	286
6.5.6.4	Routing scheme D	288
6.5.6.5	Routing strategy for DC2	290

6.5.7	Limitations and accommodations	291
6.5.8	Conclusions	292
6.6	DC3 Only location is unknown	293
6.6.1	Problem generation	294
6.6.2	Problem solution	295
6.6.3	Initial route construction	298
6.6.4	Behaviour at $t = t_{rev}$	300
6.6.5	Incorporation of coordination	301
6.6.5.1	Routing scheme A	302
6.6.5.2	Routing scheme B	303
6.6.5.3	Routing scheme C	305
6.6.5.4	Routing strategy for DC3	305
6.6.6	Conclusions	306
6.7	DC4 Only location is known	308
6.7.1	Problem generation	309
6.7.2	Problem solution	309
6.7.3	Demand estimation	310
6.7.4	Choosing time values	311
6.7.5	Initial route construction	311
6.7.6	Time $t=15$ scheme	311
6.7.7	Time $t=25$ scheme	311
6.7.7.1	Routing scheme E	312
6.7.7.2	Routing scheme F	313
6.7.7.3	Routing scheme G	313
6.7.7.4	Routing scheme H	314
6.7.7.5	Routing strategy $t=25$	314
6.7.8	Time $t=35$ scheme	315
6.7.8.1	Routing schemes	316
6.7.8.2	Routing strategy $t=35$	316
6.7.9	Routing strategy for DC4	316
6.7.10	Limitations and accommodations	318
6.7.11	Conclusions	318
6.8	DC5 Only demand is known	319
6.8.1	Problem generation	319
6.8.2	Problem solution	319
6.8.3	Initial route construction	320
6.8.4	$t_{rev}=10$	321
6.8.4.1	Routing scheme A	322

6.8.4.2	Routing scheme B	323
6.8.4.3	Routing strategy	324
6.8.5	$t_{rev} \approx 20$	324
6.8.5.1	Routing scheme C	325
6.8.5.2	Routing scheme D	326
6.8.5.3	Routing scheme E	327
6.8.5.4	Routing strategy	327
6.8.6	$t_{rev} \approx 30$	328
6.8.6.1	Routing strategy	328
6.8.7	$t_{rev} \approx 40$	329
6.8.7.1	Routing scheme F	330
6.8.7.2	Routing scheme G	331
6.8.7.3	Routing scheme H	331
6.8.7.4	Routing strategy	332
6.8.8	Routing strategy for DC5	334
6.8.9	Limitations and accommodations	335
6.8.10	Conclusions	336
6.9	DC6 Only time is known	337
6.9.1	Problem generation	337
6.9.2	Problem solution	338
6.9.3	Initial route construction	338
6.9.4	Determining demand	339
6.9.5	Routing strategy	339
6.9.6	Conclusions	340
6.10	DC7 Nothing is known	341
6.10.1	Problem generation	342
6.10.2	Problem solution	343
6.10.3	Determining time	343
6.10.4	Conclusions	344
6.11	Multiple nodes with dynamic attributes	345
6.11.1	Same category of dynamism	345
6.11.2	Different categories of dynamism	346
6.12	Appraisal of dynamic coordination in the rectilinear grid network .	348
7	Conclusions and Recommendations	351
7.1	Conclusions	351
7.2	Recommendations for future research	358

A	Reduction of rectilinear grids	361
B	Changing depot location	373
C	Data for examples in Chapter 4	383
D	Data and worked example for standard problem in Chapter 5	413
E	Grid density comparison	427
F	Data for examples in Chapter 6	491
	Bibliography	493

LIST OF FIGURES

2.1	The optimal refuelling point is restricted to the spherically convex hull defined by the origin, destination and tanker base. ..	21
3.1	These bitumen carriers are transferring part of the load from one vehicle to a second.	24
3.2	Mid-air refuelling is a form of mobile coordination.	32
3.3	Coordination can be categorised in many ways.	33
3.4	Sequential multiple vehicle at node.	35
3.5	Meeting on an existing route.	38
3.6	Detour.	40
3.7	Single vehicle coordination.	42
3.8	Single vehicle coordination involving detour.	43
3.9	Three-way transfer between courier vans.	44
3.10	Many-way transfer.	45
3.11	A single-depot, ten-node system in the Euclidean plane.	50
3.12	One snapshot of the system from Figure 3.11.	51
3.13	Another snapshot.	52
3.14	Node-based approach.	53
3.15	Route-based approach.	55
3.16	The situation prior to exchange.	56
3.17	Options C1 and C2.	57
3.18	Options C3 and C4.	57
3.19	Option C5.	58
3.20	Option C6.	58
3.21	Intersection-based approach.	59
3.22	Paths in the rectilinear grid.	60
4.1	Euclidean convex hull.	64
4.2	Zones of influence.	65
4.3	3-opt exchange.	68

4.4	Inter-route node transfer.	69
4.5	Inter-route node exchange.	70
4.6	Inter-route path exchange.....	75
4.7	Split delivery.	81
4.8	Effect of split delivery during construction and improvement using eight different construction methods.	85
4.9	Sparse node density (10 nodes) without (upper) and with (lower) split deliveries.....	86
4.10	Medium node density (30 nodes) without (upper) and with (lower) split deliveries.....	87
4.11	Dense node density (75 nodes) without (upper) and with (lower) split deliveries.....	88
4.12	The near radius.....	91
4.13	Changing the value of the near radius.....	94
4.14	Overall view of feasible routes.	95
4.15	Snapshot of system in Figure 4.14 and Table 4.5.	97
4.16	Effect of order of application of load splitting and creation of interim step.....	102
4.17	Before application of load swapping.	104
4.18	Load swapping at node 9.....	106
4.19	The preferred exchange plan.	108
4.20	Single-routed fleet before reduction in total distance by load swapping.....	109
4.21	Reduction in distance by load swapping at node 5.	109
4.22	Before coordination.	111
4.23	Load swapping reducing the fleet size.	114
4.24	Node 9 is distant from the other nodes serviced by the same vehicle.	116
4.25	After servicing node 28, the red vehicle is able to continue to a coordination site at node 8, collect demand(9) and deliver it before returning to the depot.....	118
4.26	Situation before coordination.	120
4.27	Creation of coordination sites at nodes c and i has removed the requirement for an inter-route return to depot.	121
4.28	The multi-routed red vehicle services three single-noded routes.	122
4.29	One inter-route return to depot is removed for a vehicle servicing multiple single-noded routes.	123
4.30	Removal of inter-route return to depot between two multi-	

noded routes.....	124
4.31 Investigation of the formation of an interim step.	125
4.32 Coordination as an interim step.	126
4.33 Structures resulting from dissolution of the interim step.	127
4.34 The two type (iii) dissolution structures.	128
4.35 Before coordination.....	129
4.36 The metastable state for intersection 1 in Table 4.21.	130
4.37 After the application of interim step.	133
4.38 During the investigation of coordination as an interim step, eight intersections will be investigated.	134
4.39 After a first application of coordination as an interim step in combination (a), three intersections have been eliminated and two have been changed.....	137
4.40 A second interim step formation.	138
4.41 The result of cascading investigations of using coordination sites as interim nodes.....	139
4.42 Stability of intersections.	140
4.43 The number of intersections is unchanged.	141
5.1 Paths of equal length between nodes.	150
5.2 Different routes sharing common edges.....	150
5.3 Rectilinear and staircase paths.	153
5.4 Grid section.	155
5.5 Finding the shortest distance between two nodes.	160
5.6 Slight changes in depot location can produce different outcomes.	162
5.7 Sparse grid network showing a modified coordination hull.....	164
5.8 Vertex classification.	166
5.9 Classification changes with edge removal.	167
5.10 A no-choice-string is removed if it is no shorter than another path between the same endpoints.	170
5.11 Of two no-choice-strings with the same endpoints, no more than one is retained in the reduced grid.....	170
5.12 No-choice-strings that include essential points.	172
5.13 Node or depot at end of no-choice-string.....	173
5.14 Several nodes on a no-choice-string.....	174
5.15 Two nodes on a no-choice-string.	175
5.16 Edge between nodes is retained.	176
5.17 In this reduced grid, the red edges are not part of any	

shortest paths between nodes.....	176
5.18 Grid before and after reduction.	177
5.19 Effect of number of nodes on grid reduction.....	178
5.20 Grid reduction is affected by node location and edge density.....	179
5.21 Equivalent paths.	180
5.22 Piecewise reduction of a grid section.	182
5.23 Effect of node location on a complete grid.	183
5.24 Geographical features may be recognised in the structure of a rectilinear grid.	184
5.25 The effect of nodes on opposite sides of a stream is to retain some semblance of the stream in the reduced grid.	185
5.26 Before and after grid reduction. A compact result.	186
5.27 Before and after grid reduction. A spread result.	187
5.28 Inter-route node transfer.	191
5.29 Inter-route node exchange.	192
5.30 Inter-route node exchange requiring a different placement of a transferred node.	193
5.31 Intra-route edge exchange.	194
5.32 Intra-route edge exchange in which the exchanged edges share a common node.....	194
5.33 Inter route edge exchange. Edges <i>b</i> and <i>e</i> are swapped between routes.	195
5.34 Scheduling vehicles using two different objectives.	197
5.35 A site can have more than one classification depending on how it is used.	201
5.36 Effect of maxVehDist on fleet size.....	203
5.37 Different values of maxVehCap may produce different schemes.	209
5.38 Effect of CD on R and minimum number of routes for a system with maxVehCap=235.....	211
5.39 Three-vehicle swapping.	212
5.40 Intersection at points and regions.....	213
5.41 Routes prior to swapping.	214
5.42 Routes with load swapping.	215
5.43 Routes prior to detour and swapping.....	216
5.44 Routes with detour and load swapping.	217
5.45 Multi-routed vehicle prior to coordination.....	218
5.46 No inter-route return to depot for red vehicle.....	219
5.47 Positioning a type 1 site.	220

5.48	Before coordination. Several stretches of overlap.....	221
5.49	Reduced overlap.....	223
5.50	Four vehicles service the routes before coordination.....	224
5.51	Routes after coordination, showing fleet reduction.	226
5.52	Vertices acceptable to the red route as coordination sites.	228
5.53	Vertices acceptable to the green route as coordination sites.	228
5.54	Potential coordination sites.	229
5.55	Distance reduction values at potential coordination sites.....	229
5.56	Green vehicle detours to upper left coordination site.	230
5.57	Red vehicle detours to lower left coordination site.....	230
5.58	Red vehicle detours to lower right coordination site.....	231
5.59	Three equivalent paths for the green vehicle.	232
5.60	Medium density grid.....	234
5.61	Arrows indicate the direction of shortest paths to the depot D1. .	235
5.62	Arrows indicate the direction of shortest paths to the depot D2. .	237
5.63	Limited access to a depot can increase the route density near it.	238
5.64	Variation of distance reduction with CD and depot location.....	238
5.65	Sequential and simultaneous coordination.	239
5.66	Sparse grid.	241
5.67	Dense grid.	242
5.68	Grid density and relative effectiveness of coordination methods..	244
5.69	Good (upper, distance 222) and poor (middle, distance 236) solutions can be improved by coordination (lower, distance 198).	249
5.70	Grid structure prevents coordination.	250
5.71	Routing scheme prevents coordination.	250
5.72	Coordination hulls of grid networks and Euclidean systems do not necessarily coincide.	252
6.1	Diversion from route ABC at location A' to route AA'DBC.....	259
6.2	New request occurs at time t but is not actioned until time $t+\delta t$.	260
6.3	Example structure.	266
6.4	Multiple placement in a route of a node with unknown time attribute.	270
6.5	DC1.	271
6.6	Initial routes generated with an estimated demand for node 15..	281
6.7	Progress at $t=15$	282
6.8	Initial routes incorporating coordination points which may be used if the demand of node 15 exceeds 82.	284

6.9	Routing scheme C.	287
6.10	Routing scheme D.	290
6.11	Upper: Red discs indicate the extent of 10 units from the depot. Lower: Blue lines indicate the coverage of a coordination site established at S at $t=10$; a site established at T at $t=10$ covers all but the two indicated vertices.	298
6.12	Initial routes generated with geographic coordination sites, S and T.	300
6.13	Situation at $t=t_{rev}$	301
6.14	Routing scheme A – the teal vehicle services node 15.	302
6.15	Routing scheme B – the violet vehicle services node 15.	303
6.16	Routing scheme C1 – the blue vehicle services node 15.	304
6.17	The scheme adopted depends on the location of node 15.	306
6.18	Deterministic routing scheme for a node at X with demand 50. ...	308
6.19	Situation at $t=25$	312
6.20	Situation at $t=35$	315
6.21	Initial routes.	321
6.22	Vertices for which scheme A is effective at $t=10$	322
6.23	Vertices for which scheme B is effective.	323
6.24	Vertices for which scheme C is effective at $t=20$	326
6.25	Vehicle locations and serviced nodes at $t=30$	328
6.26	Vehicle locations and serviced nodes at $t=40$	330
6.27	Effective schemes for vertices at $t=40$	332
6.28	The vehicle to service node 15, with non-zero demand less than 8, depends on the vertex at which the node is located.	340
6.29	Three types of dynamism exist in this structure.	347
A.1	Sparse density grid, 438 edges, 10 nodes, 1 depot.	361
A.2	Cul de sacs removed.	362
A.3	NCS removed without considering NCS involving essential points	362
A.4	Cul de sacs removed.	363
A.5	NCS removed without considering NCS involving essential points	363
A.6	Removed NCS with essential points.	364
A.7	Removed NCS, thus creating other NCS.	364
A.8	Final grid has 96 edges.	365
A.9	Medium density grid, 1 depot, 30 nodes, 643 edges.	366
A.10	Cul de sacs removed.	366
A.11	First reduction by removal of NCS.	367

A.12	Second reduction by removal of NCS.....	367
A.13	Removal of remaining NCS not involving essential points.....	368
A.14	Removal of NCS involving nodes.....	368
A.15	Removal of NCS created in previous step.	369
A.16	Removal of NCS created in previous step. Final grid.	369
A.17	High density grid. 790 edges. 1 depot, 30 nodes.	370
A.18	Removal of NCS not involving essential points.....	370
A.19	Removal of NCS.	371
A.20	Removal of NCS.	371
A.21	Removal of NCS.	372
A.22	Removal of NCS, including loop.	372
A.23	Final reduced high density grid. 451 edges.	373
A.24	Initial full density grid. 840 edges.	374
A.25	Removal of NCS not involving essential points.....	374
A.26	Removal of NCS involving essential points.	375
A.27	Removal of NCS created in previous step.....	375
A.28	Removal of NCS created in previous step.....	376
A.29	Final grid. 475 edges. Each outside corner houses a node.....	376
B.1	Locations for type 1 sites	381
B.2	Final routes involving coordination, depot at [10,8]	381
B.3	Routes after coordination, depot at [12,12]	384
C.1	Without coordination	388
C.2	With coordination	389
C.3	Initial routes.	391
C.4	Split deliveries at node.	392
C.5	Initial routes.	393
C.6	Routes including coordination.....	394
C.7	Initial routes.	395
C.8	Near radius=2.....	396
C.9	Near radius=5.	397
C.10	Near radius=10.	399
C.11	Near radius=15.	399
C.12	Near radius=20.	400
C.13	Near radius=25.	401
C.14	Near radius=30.	402
C.15	Initial routes.	407

C.16	Interim step	407
C.17	Load splitting the interim step	409
C.18	Load splitting the non-coordinated routes	411
C.19	Metastable state configurations	414
C.20	Routes yielding greatest reductions	415
D.1	The standard layout for medium grid problems with depot located at $D, (10,10)$	417
D.2	Coordination sites are indicated by arrows.	428
E.1	Sparse grid layout.....	432
E.2	Medium density grid.....	441
E.3	Coordination site locations used in improvement by coordination for systems with depot at $D1$	442
E.4	Coordination site locations used in improvement by coordination for systems with depot at $D2$	463
E.5	Dense grid layout.....	484
E.6	Coordination site locations on the dense grid.....	485
F.1	Example structure.....	495

LIST OF TABLES

1.1	Scope of this study	2
3.1	Route order for a three-routed vehicle.....	48
3.2	Routing scheme for the system in Figure 3.12.....	51
3.3	Changes to the routes caused by using a different combination of route order and orientation.	52
3.4	Combination options.	56
4.1	Effect of depot location on routes for a single depot system of 25 nodes located randomly across the plane.....	62
4.2	Effect of route construction technique on average values for fleet size and total distance.....	84
4.3	Effect of varying the size of the near radius.	92
4.4	Close List.....	96
4.5	Relationship between environment and coordination type in the Euclidean plane.	99
4.6	Effect of order of application of coordination type.	100
4.7	Without coordination.....	104
4.8	Progressive lading at node 9.	106
4.9	Results of coordination to reduce distance travelled.	107
4.10	Single-routed fleet before reduction in total distance by load swapping.	109
4.11	Distance reduction by load swapping at node 5.	110
4.12	Before coordination.....	111
4.13	Possible combinations for investigation of load swapping at node b.	112
4.14	After coordination at node b.....	113
4.15	Routing scheme prior to transfer.	116
4.16	A coordination site is established at node 8.....	118
4.17	Without coordination.	121
4.18	After coordination.	122

4.19	Multi-routed vehicle servicing only single-noded routes.	122
4.20	Route data for the scheme shown in Figure 4.35.	129
4.21	Interim step investigation.....	130
4.22	Fixed links from the dissolution of the metastable state shown in Figure 4.36.....	130
4.23	Routes resulting from the dissolution of intersection 1 in Figure 4.35.	132
4.24	Using coordination point as interim step.....	132
4.25	Route structure before coordination for the routes in Figure 4.38.	135
4.26	Options.	136
4.27	Interim step investigation of the results of combination (a).	137
5.1	Slight differences in depot location can have large differences in outcomes.	163
5.2	Classification of vertices with respect to the number of edges attached to them.....	165
5.3	Route generation using different clustering methods.....	190
5.4	Surplus table for a system with p vehicles.	201
5.5	Effect of maximum distance constraint on fleet size.....	203
5.6	Typical effect of maxVehCap on total distance and fleet size.	205
5.7	Effect of changing maxVehCap from 240 to 245.	206
5.8	Effect of individual route distances on total distance and fleet size (maxVehDist=50).....	207
5.9	Effect of changing maxVehCap from 235 to 250.	208
5.10	Effect of maxVehCap on total distance and fleet size for the standard problem in Appendix D.....	209
5.11	Activity at the coordination site in three-way swapping.	213
5.12	Prior to swapping.	214
5.13	Load swapping at coordination site, x.....	215
5.14	Prior to detour and swapping.	215
5.15	Detour and load swapping.	217
5.16	Prior to coordination.	218
5.17	No inter-route return to depot for red vehicle.....	219
5.18	Positioning a type 1 site.	220
5.19	Before coordination.	222
5.20	Reduced overlap.....	222
5.21	Without coordination.....	225
5.22	Fleet reduction after coordination.	225

5.23	Comparison of solutions.	231
5.24	Reductions in distance and fleet size using splitting and other forms of coordination in grids of different density when all nodes have the same demand (CD=150).	233
5.25	Reductions in distance (%) and fleet size using splitting and other forms of coordination in the medium grid.	236
5.26	Routing schemes for the system with depot at D1 and high demand nodes 13 and 23, differing only in the value of CD.	239
5.27	Reductions in distance (%) and fleet size using splitting and other forms of coordination in the sparse and dense grids.	242
5.28	Generic examples of possible outcomes after reducing the number of multi-routed vehicles.	247
6.1	Dynamic coordination classification.	266
6.2	Routing scheme A for DC1 example.	272
6.3	Routing scheme B for DC1 example.	273
6.4	Routing scheme C for DC1 example.	274
6.5	Routing scheme D for DC1 example.	274
6.6	Routing scheme E for DC1 example.	275
6.7	Routing scheme F for DC1 example.	275
6.8	Routing scheme for static DC1 example without coordination.	276
6.9	Initial routes using an estimated value for Demand(15) without coordination.	281
6.10	Progress at t=15.	282
6.11	Routing scheme A.	285
6.12	Routing scheme B.	286
6.13	Routing scheme C.	287
6.14	Routing scheme D.	289
6.15	Routing strategy for DC2 example.	290
6.16	Routing strategy based on limitations of node 6.	291
6.17	Assignments for nodes with known location.	296
6.18	Initial routes using coordination sites at S and T but otherwise without coordination.	299
6.19	Teal vehicle services node 15.	302
6.20	Violet vehicle services node 15.	303
6.21	Blue vehicle services node 15. Scheme C1.	304
6.22	Black vehicle services node 15. Scheme C2.	305
6.23	Deterministic routing for system with node 15 at vertex X	

	(in Figure 6.18) and demand 50.....	308
6.24	Routing scheme E.....	313
6.25	Routing scheme F.....	313
6.26	Routing scheme G.	314
6.27	Routing scheme H.	314
6.28	Routing strategy for $t_{rev}=25$	315
6.29	Routing strategy for $t_{rev}=35$	316
6.30	Rendering schemes consistent.	316
6.31	Consistent schemes.	317
6.32	Routing strategy for DC4 example.	318
6.33	Initial routes for DC5 example.....	320
6.34	Routing scheme A for DC5 example.	322
6.35	Routing scheme B for DC5 example.	323
6.36	Routing scheme C for DC5 example.	325
6.37	Routing scheme D for DC5 example.	326
6.38	Routing scheme E for DC5 example.	327
6.39	Routing scheme F for DC5 example.	330
6.40	Routing scheme G for DC5 example.	331
6.41	Routing scheme H for DC5 example.	331
6.42	Surplus table for DC6 example.	339
6.43	Assignments for nodes with known location and demand.....	346
B.1	Demands and node locations	377
B.2	Route structure without coordination, depot at [10,8].	378
B.3	Surplus table for depot at [10,8]	379
B.4	Route structure with coordination, depot at [10,8].....	380
B.5	Route structure before coordination, depot at [12,12].	382
B.6	Surplus table, depot at [12,12].	383
B.7	Route structure after coordination, depot at [12,12].	384
B.8	Effect of depot location.	385
C.1	Without coordination.....	387
C.2	Split deliveries to nodes 1, 2 and 6.	388
C.3	Initial routes.	390
C.4	Split deliveries.	391
C.5	Without coordination.....	393
C.6	Coordination.	394
C.7	Initial routes.	395

C.8	Near radius=2.	396
C.9	Near radius=5.	397
C.10	Near radius=10.	398
C.11	Near radius=15.	399
C.12	Near radius=20.	400
C.13	Near radius=25.	401
C.14	Near radius=30.	402
C.15	Initial routes.	403
C.16	Orientation and route order.	404
C.17	Locations and demands.	405
C.18	Initial routes.	406
C.19	Interim step.	408
C.20	Edges broken and created.	409
C.21	Load splitting after interim step.	410
C.22	Load splitting the non-coordinated routes.	412
C.23	Locations and demands.	415
C.24	Options.	416
D.1	Node locations and demands.	418
D.2	Initial load plan.	419
D.3	Surplus table for initial plan improved with split deliveries.	419
D.4	Surplus table.	420
D.5	Surplus table with total distance 306.	421
D.6	Surplus table with total distance 298.	422
D.7	Surplus table with total distance 316.	423
D.8	Surplus table for vehicles.	423
D.9	Surplus table.	424
D.10	Surplus table.	424
D.11	Surplus table with total distance 288.	425
D.12	Comparison with original.	425
D.13	Surplus table for re-constructed routes.	426
D.14	Surplus table.	428
D.15	Final routes and their surplus table, with total distance 288.	429
E.1	Sparse grid. All demands=150. No coordination	433
E.2	Improvement by coordination	434
E.3	All other demands=100. No coordination	435
E.4	All other demands=100. Coordination	435

E.5	All other demands=90. No coordination	436
E.6	All other demands=90. Coordination	436
E.7	All other demands=80. No coordination	437
E.8	All other demands=80. Coordination	437
E.9	All other demands=70. No coordination	438
E.10	All other demands=70. Coordination	438
E.11	All other demands=60. No coordination	439
E.12	All other demands=60. Coordination	439
E.13	All other demands=50. No coordination	440
E.14	All other demands=50. Coordination	440
E.15	Medium grid. D1. All demands=150. No coordination	443
E.16	Improvement by coordination	444
E.17	All other demands=100. No coordination	445
E.18	All other demands=100. Coordination	445
E.19	All other demands=90. No coordination	446
E.20	All other demands=90. Coordination	446
E.21	All other demands=80. No coordination	447
E.22	All other demands=80. Coordination	447
E.23	All other demands=70. No coordination	448
E.24	All other demands=70. Coordination	448
E.25	All other demands=60. No coordination	449
E.26	All other demands=60. Coordination	449
E.27	All other demands=50. No coordination	450
E.28	All other demands=50. Coordination	450
E.29	All other demands=100. No coordination	451
E.30	All other demands=100. Coordination	451
E.31	All other demands=90. No coordination	452
E.32	All other demands=90. Coordination	452
E.33	All other demands=80. No coordination	453
E.34	All other demands=80. Coordination	453
E.35	All other demands=70. No coordination	454
E.36	All other demands=70. Coordination	454
E.37	All other demands=60. No coordination	455
E.38	All other demands=60. Coordination	455
E.39	All other demands=50. No coordination	456
E.40	All other demands=50. Coordination	456
E.41	All other demands=100. No coordination	457
E.42	All other demands=100. Coordination	457

E.43	All other demands=90. No coordination	458
E.44	All other demands=90. Coordination	458
E.45	All other demands=80. No coordination	459
E.46	All other demands=80. Coordination	459
E.47	All other demands=70. No coordination	460
E.48	All other demands=70. Coordination	460
E.49	All other demands=60. No coordination	461
E.50	All other demands=60. Coordination	461
E.51	All other demands=50. No coordination	462
E.52	All other demands=50. Coordination	462
E.53	Medium grid. D2. All demands=150. No coordination	464
E.54	Improvement by coordination	465
E.55	All other demands=100. No coordination	466
E.56	All other demands=100. Coordination	466
E.57	All other demands=90. No coordination	467
E.58	All other demands=90. Coordination	467
E.59	All other demands=80. No coordination	468
E.60	All other demands=80. Coordination	468
E.61	All other demands=70. No coordination	469
E.62	All other demands=70. Coordination	469
E.63	All other demands=60. No coordination	470
E.64	All other demands=60. Coordination	470
E.65	All other demands=50. No coordination	471
E.66	All other demands=50. Coordination	471
E.67	All other demands=100. No coordination	472
E.68	All other demands=100. Coordination	472
E.69	All other demands=90. No coordination	473
E.70	All other demands=90. Coordination	473
E.71	All other demands=80. No coordination	474
E.72	All other demands=80. Coordination	474
E.73	All other demands=70. No coordination	475
E.74	All other demands=70. Coordination	475
E.75	All other demands=60. No coordination	476
E.76	All other demands=60. Coordination	476
E.77	All other demands=50. No coordination	477
E.78	All other demands=50. Coordination	477
E.79	All other demands=100. No coordination	478
E.80	All other demands=100. Coordination	478

E.81	All other demands=90. No coordination	479
E.82	All other demands=90. Coordination	479
E.83	All other demands=80. No coordination	480
E.84	All other demands=80. Coordination	480
E.85	All other demands=70. No coordination	481
E.86	All other demands=70. Coordination	481
E.87	All other demands=60. No coordination	482
E.88	All other demands=60. Coordination	482
E.89	All other demands=50. No coordination	483
E.90	All other demands=50. Coordination	483
E.91	Dense grid. All demands=150. No coordination	486
E.92	Improvement by coordination	487
E.93	All other demands=100. No coordination	488
E.94	All other demands=100. Coordination	488
E.94	All other demands=90. No coordination	489
E.96	All other demands=90. Coordination	489
E.97	All other demands=80. No coordination	490
E.98	All other demands=80. Coordination	490
E.99	All other demands=70. No coordination	491
E.100	All other demands=70. Coordination	481
E.101	All other demands=60. No coordination	492
E.102	All other demands=60. Coordination	492
E.103	All other demands=50. No coordination	493
E.104	All other demands=50. Coordination	493

LIST OF PROCEDURES

3.1	Node-based approach for determining a coordination site at a node	54
3.2	Route-based approach for determining a coordination site at a node	55
4.1	Inter-route node transfer procedure.....	70
4.2	Inter-route node exchange procedure	72
4.3	Inter-route path exchange	77
4.4	Split delivery during route construction	79
4.5	Split delivery during route improvement.....	83
4.6	Route scheduling.	90
5.1	Finding the length of shortest paths between nodes.	156
5.2	Creating a path between two points.	159
6.1	Parallel tabu search approach.....	261
6.2	The inventory routing problem with satellite facilities.	263
6.3	Conservative DC1.	269
6.4	DC2.	280
6.5	Co-locating a geographic coordination site and a node.....	297
6.6	Combination approach for DC3.	297
6.7	Routing strategy DC3 example.	306
6.8	DC4. Only the location is known.	310
6.9	DC5. Only the demand is known.	320
6.10	Routing strategy based on the location of node 15 at $t=10$	324
6.11	Routing strategy based on the location of node 15 at $t=20$	327
6.12	Routing strategy based on the location of node 15 at $t=30$	329
6.13	Routing strategy based on the location of node 15 at $t=40$	333
6.14	Routing strategy for DC5 example.	334
6.15	DC6 solution.	338
6.16	DC7 solution	343

GLOSSARY

Centroid	Geographical centre, weighted by population.
Close list	List of vehicles that approach within the near radius of each other, and the time step at which this happens.
Commodity class	Type of payload.
Coordinated system	A routing and scheduling system to which coordination has been applied.
Coordination	A relaxation of the standard VRP in which the load for a customer is allowed to vary its integrity or carriage before arrival at that customer.
Coordination hull	Area of a graph within which lie all nodes and depots and the region in which coordination can occur.
Coordination point	Location at which coordination may occur.
Coordination site	Location at which coordination does occur.
Cul de sac	A section of path that has an intersection at only one end.
Delivery time	The time required to deliver a customer's demand - a linear function of the demand and the distance travelled to reach the customer's node.
Demand	The requirement by a node of goods.

Edge	A link between two vertices in a graph.
Essential point	Node or depot. Such a point must not be eliminated from a graph.
Existing route	A route that has been identified without recourse to a change mechanism, e.g. load swapping, that may be applied.
Goods	Items that are delivered to customers. May be divided into different types.
Graph	A combination of nodes, vertices and edges representing a system of customers, depots and vehicle routes.
Grid reduction	Removal of non-essential edges and vertices from a rectilinear grid.
Horizon	Temporal limitation.
Load plan	Result of routing and scheduling. The final assignment of nodes to routes and to vehicles.
Made edge	An edge that has been created, i.e. incorporated into a tour, usually after other edges.
maxVehCap	Variable representing maximum vehicle capacity.
maxVehDist	Variable representing maximum vehicle distance.
Maximum vehicle capacity	The maximum amount of payload that can be carried at any one time by a vehicle.
Maximum vehicle distance	The maximum distance permissible by law (or by other means) that any one vehicle may travel during one work cycle.

Metastable state	A transition state between two stable states.
Multi-routed vehicle	A vehicle that services more than one route.
Node	Customer, represented as a location on a two-dimensional graph.
Orientation	An indication of the direction of travel of a vehicle around a route.
Path	The actual locus of a vehicle, including nodes visited and all intervening points.
Phase space	The phase space of a system is the space of all possible states or behaviours.
Pre-determined	Decided in advance.
Near radius	Distance from node or path that is investigated for coordination possibilities. [§4.8]
No Choice String	A section of a path the has no intersections except at the ends.
Nodal zone of influence	The area defined by a minimum distance from a node within which coordination has no significant effect.
Non-coordinated	Something to which coordination has not been applied.
Payload	A generic term for anything carried by a vehicle that can be collected and/or delivered.
Rectilinear grid	A mesh that consists only of straight lines intersecting at right angles.
Routing scheme	Result of routing and scheduling for a system of nodes, depots, demands and vehicles.

Route structure	The actual nodes attended (even if not serviced) by each vehicle, in the order in which they are visited.
Service	To satisfy the demand (of a node).
Spatial	Concerned with distance and physical location.
Split delivery	The demand of a node is delivered in batches, possibly by more than one vehicle (at the same or different times) or by the same vehicle at different times. Where split deliveries are permitted, the demand for a particular node could even be split en route, carried by different vehicles, re-combined and delivered intact to the node.
Surplus table	Table indicating the differences between actual payloads and/or distances and potential (capacity) values.
Temporal	Concerned with time.
Threshold	Spatial limitation.
Transfer	The movement of goods from one location to another.
Uplift	Collection by a vehicle, e.g. of goods.
Vehicle	Any object capable of carrying payload, e.g. truck, plane, bus, van.
Work cycle	The length of time in which vehicles may be away from the depot. At the beginning of the work cycle, vehicles leave the depot to service nodes; at the end of the work cycle, all vehicles must have been returned to the depot. Normally has the value of one working day.

CHAPTER 1

INTRODUCTION

The delivery of goods to customers occurs generally in accordance with routine: Customer places the order, Sales accept the order and pass instructions to Dispatch, who check the inventory; if the stock is available, it is put on a delivery truck at the depot; Driver takes the vehicle from the depot to the Customer; Customer accepts the goods; Driver proceeds to next task. From an inventory control perspective, goods are considered to be delivered when they leave the depot but, since distribution costs often play a significant role in operating costs, much effort in recent years has been directed towards reducing these by making distribution more efficient. Logistical capabilities are thus assuming a high profile in corporate strategy, and efficient use of vehicle routing in supply chain management can improve corporate profit and customer satisfaction.

In practice, it is not uncommon for transport operators to create temporary, or even extemporaneous, "mini-depots" at convenient locations between dispatch depot and customers. This technique enables loads or part-loads to be swapped or deposited for later collection, and relieves operating costs by reducing fleet size and total distance travelled. In some situations, such as mid-air refuelling, it can increase the available area of operations. Such re-deployment of materiel between depot and customer is to be termed 'coordination' and includes the use of split delivery (as in [DT89]).

It is the intent of this thesis to introduce the concept of coordination, to show that it is an effective technique, and to investigate the types of coordination found in single depot, pure delivery, fully deterministic static systems in both Euclidean and rectilinear grid network space. In addition, the rectilinear grid network is investigated in a dynamic environment in which some or all of the main nodal attributes (location, demand, time at which all attributes become known) of some nodes are unknown at the time the vehicles leave the depot. Procedures are developed that enable both total distance travelled and fleet size to be reduced in coordinated systems relative to those without coordination.

Furthermore, a procedure is presented that reduces a rectilinear grid to a smaller, yet equally effective, grid, thereby reducing the number of edges that need to be considered and improving the efficiency of calculations involving paths within that grid. Significant computational experiences will demonstrate that the use of coordination as an improvement procedure requires only initial routes and locations of depot and customer nodes.

The following Table indicates factors both considered and not dealt with herein.

Considered	Not dealt with
▪ heuristic solutions; ▪ nodes requiring delivery only; ▪ single uncapacitated depot; ▪ fleets of identical vehicles; ▪ fleet sizes less than 15; ▪ single period planning; ▪ single commodity; ▪ coordination at nodes in static, deterministic Euclidean systems; ▪ coordination at nodes and vertices in static, deterministic and dynamic rectilinear systems; ▪ coordination sites require emptying before the end of the work period; ▪ simultaneous departure of vehicles from depot; ▪ instant communication between customers, vehicles and dispatch depot; ▪ paths (roads) of infinite capacity; ▪ paths are two-way and intersections are unrestricted; ▪ objective function involves reduction in total distance and/or fleet size; ▪ unlimited splitting of demand with the proviso that an already-split demand is not further split;	▪ optimisation; ▪ meta-heuristics; ▪ delayed departure times; ▪ pickup, or mixed delivery/pickup; ▪ capacitated depots; ▪ multiple depots; ▪ prioritised delivery; ▪ mixed commodities; ▪ time windows; ▪ spherical coordinates; ▪ fleets of non-identical vehicles; ▪ congestion; ▪ restrictions on the size or number of splits in individual demands; ▪ inter-vertex coordination sites in rectilinear systems.

Table 1.1: Scope of this study.

The following real-life problems may be used to demonstrate the application (or otherwise) of coordination.

- Problem 1:

A roading company supplies its road crews with bitumen sourced from a coastal site and wishes to extend its area of operations inland. The distances involved are large and essentially Euclidean. The payload is identical for each road gang, and the vehicles involved are custom-built with heating devices that prevent the bitumen cooling and solidifying before delivery. All locations are known in advance of the tendering process and do not change (other than by elimination at the conclusion of each job).

- Problem 2:

A delivery company wants to reduce its fleet size without adverse effect on its client base. It operates within a reasonably flat city with two-way streets that can be represented by a rectilinear grid. Individual packages are much smaller than the vehicles' capacities, and splitting of demands is not permitted. The demands of customers are known at least one day in advance of delivery.

- Problem 3:

A stock transport company wants to reduce its operating costs and provide a faster service to customers. On sales days, many of its customers require service at the same time in order to have their animals delivered to the auction yards before the sales start. Post-auction transport times vary throughout the day. Individual stock trucks contain internal baffles that enable them to carry animals from several farms simultaneously. Each truck may include a trailer of similar design. All collections from farms prior to the sales are known well in advance, and most of the post-sales deliveries to the farms are also booked tentatively in advance.

- Problem 4:

A bus company servicing rural schools wants to expand its service to an extra school but does not want to extend its total operating time by more than 20 minutes each day because of the effect on drivers and pupils. Although the schools are generally close together, the pupils live in a widely distributed area. Each of the schools starts at about the same time in the morning, but the secondary school finishes 30 minutes later than the primary schools. Once

pupils are collected from their homes or schools, they must be supervised until they are delivered to their destinations. The number, location and destination of passengers is known before the start of the school year and does not vary significantly thereafter.

- Problem 5:

A company producing small runs of high quality clothing for the tourist market wishes to improve its service by reducing the time for delivery of its goods to tourist centres at opposite ends of the country. Because the goods are fashion items, they must be on retail shelves as soon as possible and en route storage is inappropriate. The distances involved are large and essentially Euclidean, and demands are known six days in advance of delivery dates.

CHAPTER 2

LITERATURE REVIEW

This chapter discusses previous attempts at finding workable solutions to the vehicle routing problem and identifies earlier work on coordination.

2.1 The vehicle routing problem

In 1959, Dantzig and Ramser [DR59] formulated the truck dispatching problem, later known as the vehicle routing problem (VRP), and presented a heuristic solution to make efficient use of payload space on vehicles. In the many static, deterministic VRP solutions which followed, the emphasis has been on minimising the total distance travelled, although much effort has also been expended on minimising waiting times at customer nodes, fleet sizes, and combinations of these. Comprehensive surveys are found in [B83] and [GA88] and classification schemes can be found in [B81] and [DLS90].

Because the VRP is an NP-hard problem (see [GJ79]), there is unlikely to be a procedure that will solve it in polynomial time, and much effort has been applied to using inexact measures to create good approximations to optimal solutions. Nodes are allocated to routes in a routing phase, and routes are allocated to vehicles in a scheduling phase, where the routing phase accommodates spatial characteristics and the scheduling phase accommodates temporal characteristics. Heuristic procedures can be classified (see [TE98]) into the three major classes of construction procedures, improvement procedures and composite procedures.

The Capacitated Vehicle Routing Problem (CVRP) generally consists of designing suitable vehicle routes for a system that satisfy a set of constraints in order to minimise total cost. Typical constraints include

- the quantity of goods carried must not exceed the carrying capacity of a vehicle,

- the total distance travelled by a vehicle must not exceed some pre-set (user-defined) limit,
- each customer must be visited exactly once by one vehicle,
- vehicles begin and end their routes at a central depot,

although each of these can be relaxed under some circumstances. Usually, the VRP is solved by deciding which vehicles will deliver to particular customers and the order in which they should do this, and routes are constructed and improved until a satisfactory feasible solution is found. This study investigates the improvement method of coordination, and the discussion below indicates some of the route construction and improvement methods used in the investigation, as well as some of the methods that, while not used here, have been influential therein.

Construction procedures (see [GB80] and [C76]) generate feasible solutions starting from a distance matrix but, although the methods are fast, the resulting solutions are often poor and require improvement.

Simple iterative improvement procedures, such as 3-opt (see [LK73]), start with an initial solution and improve it by edge exchange methods. Where a better solution is found, it replaces the previous solution and the process is repeated until no further improvements are found. These improvement procedures have the disadvantage of becoming trapped at local optima, as they are based upon a pure descent strategy.

Composite procedures (see [GB80]) combine construction and improvement procedures by producing good initial solutions quickly and then applying a computationally expensive improvement method to give very good results in a relatively fast time.

In applying coordination in this study, initial routes are considered to have been constructed and improved by a variety of methods, as discussed in the relevant chapters.

In the early route construction method known as the parallel savings heuristic of Clarke and Wright (see [CW64]), distance, capacity and fleet size constraints are ignored and tours are formed simultaneously, with initially each node being served from the depot (node 0) by a distinct vehicle. In a system with distance matrix $[c_{ij}]$, for each combination of two customers, i and j , the saving resulting from servicing both by the same vehicle is calculated by

$$S_{ij} = C_{0i} + C_{0j} - C_{ij}.$$

All savings are sorted in decreasing order and customers i and j corresponding to the highest unused saving are merged if constraints are not violated. Clusters grow simultaneously but only two are merged at each iteration.

Altinkemer and Gavish (see [AG91]) combine the savings approach with matching based procedures, showing that the average performance of tour partitioning based procedures is significantly inferior to parallel savings based heuristics. Their parallel savings heuristics merge multiple clusters of nodes at each iteration, and solution of a matching problem determines the number of clusters to be merged.

The Sweep algorithm (see [GM74]), used in this study and shown in Procedure 2.1, constructs routes for VRPs with one or several depots, and nodes located in the Euclidean plane. Nodes are represented by their polar coordinates, (θ_i, p_i) , where θ_i is the angle and p_i is the ray length from the origin.

Sweep	
Step One	The value $\theta^* = 0$ is assigned to an arbitrary node i^* and all remaining angles are computed. Rank the nodes in order of increasing value of θ_i .
Step Two	Choose an unused vehicle k .
Step Three	Starting from the unrouted node having the smallest angle, assign nodes to the vehicle as long as its capacity is not exceeded. If unrouted nodes remain, go to Step One.
Step Four	Optimise each vehicle route separately by solving the corresponding TSP. Perform node exchanges between adjacent routes if this saves distance. Re-optimize and stop.

Procedure 2.1: Sweep algorithm for route construction.

While some instances of the CVRP can be solved exactly, the time required is often excessive as the problem is NP-hard. Many of the best available exact algorithms are based on dynamic programming and the most promising meta-heuristic methods for the CVRP are based on Tabu Search (see Procedure 2.2 for a simple version). Although Tabu Search is not used in this study, there is nothing inherent in the coordination process that prevents its incorporation as an improvement step into a Tabu Search procedure. The Tabu Search procedure constructs a sequence of solutions before implementing an improvement step. Because successive vehicle routes may not be feasible, their degree of departure from feasibility is measured by means of a penalty function, and a tabu list is

generated that prevents the same feasible solution from being re-visited. Basic difficulties with this heuristic are the design of a suitable tabu list, the efficient management of this list, and the selection of an appropriate move in Step Two.

Rochat and Taillard (see [RT95]) first generate promising solutions to the VRP using Tabu Search procedures (see [G89]), then extract single vehicle routes from this population of solutions and combine some of these to define a partial starting solution for another application of Tabu Search. Repeating the process a number of times provides some of the routes generated as candidates for the final VRP solution.

Tabu Search	
Step One	Create back and forth routes between the depot and customers. Let x be this solution with cost $F(x)$. Set the tabu list $T := \emptyset$.
Step Two	Define $N(x)$, the neighbourhood of x , as the set of all solutions that can be generated by a move from x . If $N(x) \setminus T = \emptyset$, go to Step Three. Otherwise, identify the least cost solution y in $N(x) \setminus T$ and set $x := y$. Update the best known solution, if necessary.
Step Three	If the maximum number of iterations since the beginning of the process or since the last update has been reached, go to Step Four. Otherwise update T and go to Step Two.
Step Four	Attempt to improve each of the vehicle routes.

Procedure 2.2: Elementary Tabu Search.

Taillard et al (see [TLG96]) provide an algorithm based on this to solve the VRP variant in which the same vehicle may be assigned to several routes during a given planning period, where the total duration of a route, including service and travel times, does not exceed a pre-set limit. Initially, a large set of good routes are generated, from which a subset is selected and combined into feasible working days using a bin packing heuristic.

Raft (see [R82]) provides a modular algorithm by decomposing the solution of the VRP into five main decision areas: route assignment, depot assignment, vehicle assignment, delivery period, and detailed route design. Each part is solved separately and then connected by an iterative procedure, allowing multiple use of vehicles and postponement of delivery. This accommodates fixed depot and fixed vehicle assignments for each customer, stochastic demand and stochastic service.

Reinelt (see[R94]) gives a summary of heuristic solutions to the TSP, including Simulated Annealing, evolutionary strategies and Genetic Algorithms, Tabu Search, and Neural Networks. Tsubakitani and Evans ([TE98]) use jump search to apply local search methods to different starting points. Algorithms that solve the TSP to optimality are not stable in the sense that similar sized problems can take vastly different times to solve, and problems with a hundred nodes can take hours of CPU time, unacceptable to practitioners in the field. Also, problems encountered in practice may be far larger than this.

Unlike the TSP, the VRP needs to account for variants (see [HP97]) that include:

- multiple use of vehicles,
- route capacities,
- time windows,
- dynamic routing, and
- randomness,

all of which are suitable for improvement by coordination, as demonstrated in this study although the use of time windows will be left for future research and randomness is mentioned in passing in the discussion on dynamic coordination in Chapter 6. For more information on stochastic vehicle routing see [B95].

Solutions for the VRP are often derived from those for the TSP ([L92]), and heuristic approaches to the VRP are characterised by the broad range of approaches, but although the VRP is a generalisation of the TSP, the concepts that led to the progress in solving large scale travelling salesman problems cannot be applied to vehicle routing problems because all the improvement heuristics used for the TSP can be used for no more than local improvement within one of the tours. Recent efforts to develop solutions to the VRP have incorporated artificial intelligence concepts, and include Simulated Annealing, evolutionary strategies and Neural Network heuristics.

Simulated annealing (see [KGV83], [W98], [F99]) derives its origins from theoretical physics where Monte-Carlo methods are employed to simulate phenomena in statistical mechanics. Evolutionary strategies attempt to imitate evolution in nature, and are based on analogues of mutation and selection. Genetic Algorithms (see[G89a]) construct tours from a suitable combination of two or more tours. Neural Network heuristics (see [P93], [S99]) are modelled on the functioning of the human brain; based on the inputs that a neuron receives, a certain output is computed which is propagated to other neurons. A variety of models addresses the

activation status of neurons, determination of outputs and propagation of signals in the net with the basic goal to realise some form of learning mechanism, but Osman (see [O96]) suggests that the quality of the solutions might not compare to the solutions produced by the best heuristics.

Because of the inter-dependence of the temporal and spatial characteristics of coordination problems, their solution is described in this study using simple systems with fleet sizes limited by the underlying software to 15 vehicles, and their application in Tabu Search, evolutionary strategies and neural networking is an area for future research.

Nygard et al (see [NJK90]) provide empirical evidence that, for systems with unchanged demands, the performance of various heuristics is sensitive to the statistical distributions of the geographical locations of nodes in the Euclidean plane and that no heuristic consistently out-performs all others.

Lund et al (see [LMR96]) describe the difference between dynamic and stochastic VRPs by the responsive nature of the VRP solutions: the stochastic VRP determines the a priori two stage route using stochastic programming techniques and maintains the solution already obtained even if the data changes. Since all information, deterministic as well as stochastic, is incorporated in the model before the determination of the routes, stochastic models are static. In contrast, dynamic models have inputs (and consequently, solutions) that change over time. The degree of dynamism determines the style of heuristic: weakly dynamic systems are solvable using adaptations of static procedures, rerun each time a revision of the input occurs. They term this "myopic" or "rolling horizon" optimisation, and determine the degree of dynamism of a system by creating a simulator that generates new (real-time) requests and noting the effect of changing the distribution of advance and real-time requests. Their work is based on the presumption (as in [Ps95]) that a static updating procedure is only suitable for problems with very weak (quasi-static) degrees of dynamism. They show that when the number of dynamic requests make up more than about 30% of the total number of requests, the solution of dynamic routing problems based on adaptation of static procedures is no longer effective.

Psaraftis ([Ps95]) highlights the recent advances in technology that enable decision-makers to receive dynamically the information needed to create suitable routes and schedules, and cites electronic data interchange (EDI), geographic information systems (GIS), global positioning systems (GPS) and intelligent vehicle-highway systems (IVHS) as technological improvements that have enhanced the possibilities for dynamic routing. Such advances have shifted the emphasis of more recent researchers from static deterministic VRP solutions to

dynamic models. Although the TSP can be solved optimally for a few thousand nodes, the VRP resists optimal solutions for problems above 50 customers. Advances in column generation methods and the creative use of cutting planes and bounding techniques (as in [D99]) enable some VRPs to be solved to optimality.

Dror and Powell [DP93] compare the deterministic environment with dynamic and stochastic environments and note that longer range planning problems can be represented more realistically by a dynamic model than by a static one. They further note that the major difference between static problems and dynamic and stochastic problems is that, unlike the standard VRP, stochastic and dynamic vehicle routing problems are often characterised by the lack of a well defined objective function. In addition, several modelling frameworks may be applied to the same problem.

Minkoff (see [M93]) models the dynamic delivery dispatching problem (DDP) as a Markov decision process in which vehicles are periodically sent out to replenish stock at a set of customers experiencing stochastic demand. When one or more vehicles are ready to be dispatched, the dispatcher learns the current inventory level at each customer, then selects a routing scheme from a pre-assembled list of feasible schemes. The vehicles then execute their routes, which are not changed once assigned. When more vehicles become available, the process repeats itself. The creation of schemes and the routing of vehicles is performed in advance (rather than as information becomes available), with the transportation costs representing an input into the decision. Eligible routes are thus an input to the DDP, as in set partitioning-based approaches to vehicle routing (see [CJR81] and [AMS89]).

Unlike the DDP, vehicle routes are created in the course of solving the Inventory Routing Problem (IRP) (see [M93]). The IRP tends to cover a finite number of time periods, whereas the DDP is an infinite-horizon problem. The DDP commits only to immediate dispatches, basing future dispatches on later inventory levels, whereas the IRP decides several time periods worth of dispatches at once and executes them over the dispatch horizon. The IRP with satellite facilities (IRPSF) permits the creation of 'sub-depots' for the replenishment of stock but not for the storage of vehicles, and (as a form of coordination) is discussed in §2.2.5.

Dynamic VRP has also been studied under the titles of allocation, routing, and dispatching. Powell (see [P88]) explores dynamic vehicle allocation, which concerns itself with the deployment of a fleet of vehicles over time to maximise profit, and decisions include where to direct each vehicle in each time period, and whether or not to carry loads which may be there.

Bertsimas and van Ryzin ([BV91]) develop the dynamic travelling repairman problem to model waiting time considerations in the provision of service to

customers distributed over a Euclidean plane, where demand for service is stochastic and a single, uncapacitated vehicle, travelling at constant velocity must service demands whose time of arrival, location and on-site service are stochastic. The objective is to find a policy to service demands over an infinite horizon that minimises the expected system time (wait plus service) of the demands. They extend ([BV93]) this to the case where a fleet of vehicles is used and each vehicle can visit at most a certain number of customers before it must return to a specified depot location, and find that solutions depend strongly on the geometry of the problem (depot location, size and shape of the service region).

Many real-world problems involve considerable uncertainty in the problems data. For example, locations may be known only probabilistically in advance and the demands they place on the vehicle capacity may be random. In addition, requests for service often arrive sequentially in time, and these events may be stochastic. Where static VRP problems are often solved from the dispatcher's perspective in which total distance is minimised, dynamic problems are often more concerned with the customer's perspective, and minimising the total delivery time may be more appropriate.

Psaraftis and Tsitsiklis (see [PT93]) investigate a minimum cost path problem designed on an acyclic network whose costs are known functions of environment variables, and allow any vehicle to wait at a node until the environment becomes favourable for travel to the next node. This applies where information about the environment is available at each node that is visited, and has applications in maritime travel.

Berge and Hopperstad (see [BH93]) report on demand driven dispatch (D^3) in the airline industry, a concept that attempts to assign aeroplane capacity to flight schedules to meet fluctuating market needs using a family of aircraft rather than individual planes. This enables aircraft assignments to be swapped at short notice when aircraft share the same airport. The variability of demand is high but the demand forecast improves as flight departure approaches. Aircraft are assigned dynamically to flights to match the predicted final demands. Using a family of aircraft with enough operating flexibility to permit dynamic reassignment of the fleet means that the updated forecasts can be matched more easily. Before this approach, individual aircraft were scheduled, resulting in system loads of about 65% capacity. This is not a pure dispatch scheme since the schedule must be flown no matter what the demand. Aircraft positions are constrained at the beginning of the index week but unrestricted at the end of the week.

Savelsbergh and Sol (see [SS98]) report on an industry-specific pickup and delivery planning module, DRIVE: Dynamic Routing for Independent Vehicles, that

works iteratively with a base plan founded on current routes and demands and uses estimates of future workload. The base plan has a time horizon of about one hour. As soon as the base plan is accepted, the short term decisions are made permanent, dividing the routes into a head, the part that will be executed as planned, and a tail, which may change in the future.

In the stochastic VRP (SVRP) (see [DLT89]), customers require delivery amounts which are uncertain until a vehicle arrives. This opens up the possibility of route failure, where a vehicle runs out of the stock it is delivering before all customers on its route have been replenished. In the case of route failure, the vehicle may have to return to the depot, reload and go back to the customer where failure occurred, or a new routing sequence may have to be planned for the remaining customers. The incidence of modification caused by route failure can be reduced by allowing breaks to occur, in which case the vehicle returns to the depot for replenishment before it resumes its deliveries whether or not failure has occurred.

Two solution frameworks for SVRPs are stochastic programming (see [B97] for a recent survey) and Markov decision processes. Stochastic programming can be further divided into chance constrained programming ([SG83]) and stochastic programming with recourse.

Although deterministic approximations are widely used to solve stochastic and dynamic problems, Bertsimas et al [BCP95] show that a priori approaches are beneficial in circumstances where solving repeated deterministic instances is impractical or otherwise undesirable. In a priori solutions, the route is constructed prior to the realisation of the demands. They developed such solutions to the stochastic VRP, known as PVRP (Probabilistic Vehicle Routing Problem), in a Euclidean framework based on variants of Bertsimas's probabilistic cyclic heuristic (see [B92]). The cyclic heuristic, which is based on deterministic optimal tour partitioning (see [HR85]), begins with a TSP tour of n demand points, and the vehicle visits points in the order of this tour, returning to the depot whenever capacity is reached. Since each of the n points is a potential starting point, the cyclic heuristic considers the n resulting visitation orders and selects the one with the lowest cost. [B92] and [HR85] imply that if the initial tour for the cyclic heuristic is based on the optimal solution to the TSP, then the heuristic is itself asymptotically optimal under given conditions of n and vehicle capacity. Using a heuristic to find an initial TSP tour degrades the quality of the solution slightly, but this can be offset by local improvement and establishing threshold return strategies via dynamic programming. Each time the vehicle departs a demand node carrying a partial load, the driver faces the decision of whether to proceed to the next point or

to return to the depot only partially full. A threshold is constructed for each node such that if the load after visiting the node is at least equal to the threshold for that node, it is optimal to return to the depot before filling further demands. By comparing the results of a priori PVRP heuristics with two a posteriori heuristics, clustering ([C85]) and savings ([CW64]), [BCP95] showed that for the data sets investigated, the cyclic heuristic with dynamic programming but no local improvement gave the best combination of speed and quality. Adding local improvement gave better results but at a significant cost in running times. The cyclic heuristic was consistently superior to clustering, consistent with the idea that cluster first, route second heuristics can be asymptotically sub-optimal, and the savings heuristic generally gave the best results.

Gendreau et al (see [GLS95]) consider the VRP with stochastic demands and customers, a problem that arises in LTL collection routes. Each customer has a known probability of presence and a random demand. Using a stochastic programming framework, the problem is solved in two stages: in a first stage, planned collection routes are designed, and in a second stage, when the set of present customers is known, these routes are followed as planned by skipping the absent customers. Whenever the vehicle capacity is attained or exceeded, the vehicle returns to the depot and resumes its collections along the planned route, incurring a penalty. The problem is to design a first stage solution in order to minimise the expected total cost of the second stage solution.

The concept of creating an initial base plan that can be modified, when circumstances change, by selecting a suitable scheme from a set of a priori schemes is adopted in this study in the investigation of coordination in dynamic VRPs. The routing strategy includes the construction of thresholds to determine appropriate schemes. This is discussed in more detail in Chapter 6.

2.2 Coordination in vehicle routing problems

The coordination process involves routing and scheduling problems where the load for a customer is not transported directly from depot to customer by one vehicle, and forms one of many relaxations (see [DLS90]) of the CVRP:

- customers may be visited more than once,
- the total demand carried by a vehicle may exceed the capacity of that vehicle but the payload being carried at any moment must never exceed the vehicle's capacity.

Although each vehicle starts and ends the work cycle at the depot, it will be discussed in later chapters that because vehicles may be used many times, each route need not begin and end there. Approaches to this issue are discussed in more detail below and include:

(a) Split deliveries

During split delivery, the load for a particular customer is carried by more than one vehicle, or by one vehicle making more than one visit to that customer.

(b) Coordinated public transport

Resources are shared by having two vehicles, each with its own depot, each collect loads for both routes.

(c) Mid-air refuelling

The fuel to enable a vehicle to travel must be transferred from another aircraft.

(d) Pre-emptive routing

Part of the load is left for later uplift.

(e) Inventory routing with satellite facilities

Temporary storage depots are created that exist to replenish stock but are not used for the storage of vehicles.

2.2.1 Split deliveries

The split delivery vehicle routing problem (SDVRP) has been shown (see [DT89]) to produce reductions in both total distance travelled and fleet size, and forms an important part of the discussion on coordination. As a relaxation of the CVRP, the SDVRP allows deliveries to a node to be carried by more than one vehicle, and will be extended in later chapters of this study to permit the splitting and/or recombination to take place at sites other than the depot and intended node.

Dror and Trudeau ([DT89]) apply k -split interchange and route addition in an improvement procedure to a feasible initial solution. The k -split interchange is the technique by which demand at a particular node is split between k vehicles, each on a unique route, thus removing that node from the route that originally serviced it. No two routes can have more than one split demand point in common. The original $k+1$ route costs are compared with the k split-route costs and with the diminished route that that originally serviced the node:

$$SAV_k(p) = \sum_{t=1}^k (C_{i_t j_t} - C_{i_t p} - C_{p i_t}) + C_{b_p p} + C_{p a_p} - C_{b_p a_p}$$

where b_p and a_p are the points immediately before and after (respectively) the point p under consideration, $[c_{ij}]$ is the cost matrix for travelling from i to j , and where i_t and j_t are two consecutive points on route t . If the spare capacity of the k vehicles is at least as great as the demand of node p on route $k+1$ and if the term $SAV_k(p)$ is positive, then node p can be eliminated from route $k+1$ and the total routing cost reduced. Where the demand delivered to node p is already split between a number of routes, the k -split interchange procedure examines all split configurations for node p . The demand that node p was receiving from vehicle $k+1$ is divided among the k vehicles by using all the spare capacity of the vehicle with the least spare capacity first, then the spare capacity of the vehicle with the next least spare capacity, etc, until the demand of node p is satisfied.

In some cases, the total routing cost may be reduced by adding a route that eliminates a split delivery to a node. There are $2k^2+1$ different configurations for a single route addition that eliminates a split delivery.

Before investigating the potential of applying split deliveries in any particular problem, Dror and Trudeau improve the initial feasible solution by inter-route node swapping and 2-opt until node interchange produces no further improvement. Then k -split interchange, route addition, node interchange, and 2-opt route improvement is applied until there is no further improvement by k -split. Although this method does not guarantee an optimal SDVRP solution, it is very effective and has implications in real-world routing problems (see e.g. [M97] and [ST98]). The SDVRP and VRP produce equivalent solutions when all nodes have an identical demand of 1 unit, or when all demands require a full vehicle load (assuming a homogeneous fleet), but differ drastically when the demand of each of the n nodes is close to half the vehicle capacity of the homogeneous fleet. Then the minimum number of vehicles in the SDVRP solution is $\frac{n}{2} + 1$, which compares favourably with the VRP solution's n vehicles.

[DT89] found, as have other researchers (see [F93] p. 35, and [F95]), that, assuming no bounds on allowable splits, the advantage of SDVRP over VRP improves when the ratio of demand:vehicle capacity rises above 0.1, for problems in which nodes have identical demands, and increases with the ratio. The procedure constructs routes that consist of nodes from the same geographic area, and the nodes with split deliveries are usually those with above average demand. In addition, splits which result in high distance savings occur relatively further from the depot. Although [DT89] considers only homogeneous fleets in a Euclidean

setting and does not go beyond 3-split configurations, the principles of split delivery can be applied to other frames of reference, greater splitting and vehicle fleets of different capacities.

[M97] extends the split delivery problem in node routing to an arc routing scenario where the inter-route node exchange above is replaced by inter-route edge exchange and may be adapted to include time windows. The k -split generation is extended to investigate whether a delivery made to an arc can be split across k other candidate routes in order to obtain the highest saving. The savings obtained by splitting the demand of arc (i, j) over routes 1 and 2 is given by:

$$SAV_2(i, j) = \left(\bar{C}_{bi} + \bar{C}_{ja} + C_{ij} - \bar{C}_{ab} \right) + \left(\bar{C}_{i_1j_1} - \bar{C}_{i_1i} - C_{ij} - \bar{C}_{j_1j} \right) + \left(\bar{C}_{i_2j_2} - \bar{C}_{i_2i} - C_{ij} - \bar{C}_{j_2j} \right)$$

where $\bar{C}_{xy}$ represents the length of the shortest path from node x to node y , and C_{xy} represents the length of the arc (x, y) . For $k \geq 3$ candidate routes

$$SAV_k(i, j) = \left(\bar{C}_{bi} + \bar{C}_{ja} + C_{ij} - \bar{C}_{ab} \right) + \sum_{t=1}^k \left(\bar{C}_{i_tj_t} - \bar{C}_{i_ti} - C_{ij} - \bar{C}_{j_tj} \right)$$

where arc (i, j) is considered for introduction between consecutive nodes i_t and j_t on candidate route t . The k -split proceeds if $SAV_k(i, j)$ is positive, and the insertion is chosen that produces the highest savings in distance.

The route addition procedure of node routing is similarly extended to k -route addition in arc routing, where arcs whose demand is split among several routes are investigated for consolidation into one new route that may also contain other arcs. Empirical results indicate distance reductions in four out of five cases involving time window constraints with two cases yielding improvements of over 10%, but less obvious benefits in cases where time windows were not employed.

The SDVRP is NP-hard (see [DT90] and [LLM91]), and integer formulations (see [DLT94]), heuristic procedures (see [F95]) and column generation methods (see [ST98]) have been applied to the concept. Problems of up to 50 nodes have been solved to optimality in the Euclidean plane (see [B00]). While most investigations into the SDVRP have been set in the Euclidean plane, the concept has also been applied to problems set in the rectilinear grid (see [F93] and [F95]) with bounds limiting the ways in which demands may be split, introducing a penalty function for splitting, and allowing splitting between depots in multiple depot problems.

2.2.2 Coordinated public transport

Baker, Franz and Sweigart (see [BFS93]) consider a pure pickup vehicle routing problem in the Euclidean plane in which two cooperating agencies, each with its own depot, use the same fleet of two vehicles to transport their clients. Each agency is assigned one vehicle and n passenger stops (nodes), and vehicle routes begin and end at their own depot, with each vehicle travelling a unique route collecting clients for both agencies.

In the scheme involving coordination, based on a multi-depot location-allocation problem, vehicle routes are scheduled so that the vehicles arrive at a coordination point simultaneously after collecting all the passengers, and passengers change vehicles, if necessary, to be transported to their appropriate depot. Suitable sites for this are cul de sacs, church parking lots, parks, one of the agencies, etc. In the scheme that prohibits the use of coordination, each vehicle travels from node to node delivering passengers at the end of a route, and passengers never have to change vehicles. In each scheme, the function to be minimised is the total distance travelled by the entire fleet.

The 16×16 Euclidean plane is divided into 4×4 cells, with no more than 20 nodes and 5 potential coordination sites, where the potential coordination sites are located by considering the contents of each cell: if a cell selected for a site contains a node, the potential site is that node, otherwise, the potential site is located at the node closest to the centre of the cell. Consequently, potential coordination sites are only located in cells that contain nodes. The system included a minimum distance between stops, and a varying distance between depots. These limitations meant that the largest resulting integer linear programming problems did not exceed 250 decision variables and 250 constraints. [BFS93] found that the use of coordination sites generally reduced the total distance travelled more than the system that did not allow coordination, and that each system was far superior to the system of independently operated transportation programs without coordination. In addition, they found that the system involving coordination was more efficient as the inter-agency distance increased from a threshold level, below which the system that prohibited coordination was more efficient.

Russell and Morrel (see [RM86]) investigated a similar problem of two depots and eleven buses, using a savings heuristic modified by a 3-opt improvement, and achieved a reduction in total distance travelled of 11%. They employed a shuttle system where the passengers for both agencies were offloaded at the first node visited and those that needed to travel on were loaded onto the next available bus.

2.2.3 Mid-air refuelling

The mid-air refuelling of aircraft involves a special type of coordination; rather than the payload being transferred, the fuel to transport the payload is moved from one vehicle to another which meet en route. The subsequent consumption of this fuel is irrelevant to the process of coordination. In these circumstances, although the coordination site is mobile, the relative distance travelled by the two aircraft involved during transfer is minute compared with the total distance involved. The process is discussed here because it represents an important real-world application of coordination; the spherical frame of reference is not considered elsewhere in this study.

During mid-air refuelling, a transport aircraft travels from one node to another, while a tanker aircraft originates from and returns to a third node. The objective is to minimise the total fuel consumed subject to restrictions on the range of the transport and tanker. Yamani et al (see [YHM90]) approach this problem using a non-linear objective function in a spherical frame of reference with constraints that depend on the decision variables. The practical constraints are relaxed by making the following assumptions:

- the gross weight of the aircraft should not exceed its maximum takeoff weight,
- there are no load balance or size restrictions associated with the aircraft load,
- the Earth is a perfect sphere,
- the aircraft follows the great circle route (shortest path) in flying from one point to another,
- winds at altitude are assumed to be negligible,
- aircraft refuelling takes a negligible amount of time,
- the region in which fuel transfer occurs is considered to be a point,
- everything takes place at altitude - takeoff, landing and their costs are ignored.

The objective is to find (θ, ϕ) , g , and h that will

$$\text{minimise } V(\theta, \phi) = FC^C(g, w, d_{OR}) + FR^C(w, d_{RD}) + FC^T(h, d_{BR}) + FR^T(d_{BR})$$

subject to

$$d_{OR} \leq R^C(g_{MAX}(w), w) \quad (2.1)$$

$$d_{RD} \leq R^C(G_{MAX}(w), w) \quad (2.2)$$

$$d_{BR} \leq R^T(H_{MAX})/2 \quad (2.3)$$

$$FR^G(w, d_{OR}) \leq g \leq g_{MAX}(w) \quad (2.4)$$

$$FR^T(2 \cdot d_{BR}) \leq h \leq H_{MAX} \quad (2.5)$$

$$h + g = V(\theta, \phi) \quad (2.6)$$

where

θ	=	angle of refuelling point from the Prime Meridian (longitude)
ϕ	=	angle of refuelling point from the equator (latitude)
g	=	initial fuel of the transport aircraft
h	=	initial fuel of the tanker
w	=	cargo weight of the transport aircraft
$g_{MAX}(w)$	=	maximum fuel the transport aircraft can carry in order to be able to take off with a cargo weight w
$G_{MAX}(w)$	=	maximum fuel the transport aircraft can carry when it is in the air given that its cargo weight is w
H_{MAX}	=	maximum fuel capacity of the tanker
$FC^C(g, w, d)$	=	fuel consumed by the transport aircraft when it flies a distance d , its initial fuel is g , and its cargo weight is w
$FR^C(w, d)$	=	fuel required by the transport aircraft to fly a distance d when the cargo weighs w
$FC^T(h, d)$	=	fuel consumed by the tanker when it flies a distance d and its initial fuel is h
$FR^T(d)$	=	fuel required by tanker to fly a distance d
d_{OR}	=	spherical distance between the origin base and the refuelling point
d_{RD}	=	spherical distance between refuelling point and destination
d_{BR}	=	spherical distance between tanker base and refuelling point
$FR^T(d)$	=	fuel required by the tanker to fly a distance d .

Constraints (2.1) - (2.3) represent conditions that must be satisfied by the refuelling point. Constraint (2.1) prevents the distance between the origin base and the refuelling point from exceeding the maximum range of the transport aircraft. Constraint (2.2) indicates that the transport must be able to reach its

destination from the refuelling point. Constraint (2.3) indicates that the tanker must be able to make the round trip from its base to the refuelling point. Constraints (2.4) - (2.5) make similar restrictions on the amount of fuel required. There is only one value of h for every choice of g .

Solution involved using an upper bound function that was improved from one iteration to the next. Investigations revealed that the search for the optimal refuelling point could be restricted to the spherically convex hull of the three base points: the origin, the destination and the tanker base. Since the objective function is spherically convex, the solution procedure will converge to the global optimum solution but there is no guarantee that this will occur in a finite number of iterations.




Figure 2.1: The optimal refuelling point is restricted to the spherically convex hull defined by the origin, destination and tanker base.

2.2.4 Pre-emptive routing

During pre-emptive routing, part of the load of a vehicle is left for later uplift, and can be considered as a case of single vehicle coordination. Charikar, Khuller and Raghavachari (see [CKR98]) report on investigations involving TSPs in the Euclidean plane where a single commodity of unit size is transported from a set of suppliers to a set of demand points using a vehicle of limited capacity. Loads, or part-loads, can be dropped at intermediate locations and collected later in much the same way that a truck leaves a trailer for later uplift. They found that the ratio of the distances travelled by an optimal non-pre-emptive tour versus a pre-emptive tour is bounded by 4. When the tours are forced to start at a particular location the ratio is 5.

A less obvious version of this form of coordination is found in the manoeuvrings of an aircraft carrier battle group (see [RW96]) which detours slightly from its intended route in order for the aircraft to leave and return on short sorties. Reasons for this include maintenance checks on the aircraft, allowing flight crews to improve their training, and defensive missions. The aircraft carrier at the centre of the battle group can be thought of as the vehicle, which travels from source to destination, while the aircraft can be considered as mobile trailers, part of the vehicle, which become detached temporarily from it at intermediate locations and are collected later at a slightly different location. Normally, the carrier is forced to detour from its intended route in order to best utilise favourable environmental conditions which the departing and arriving aircraft need. Rosenthal and Walsh seek to maximise the number of sorties while minimising the detours.

2.2.5 Inventory routing with satellite facilities (IRPSF)

In the IRPSF (see [BHJD98]), temporary storage facilities are created that enable vehicles to replenish stock without having to return to the depot, but are not available for the storage of vehicles. The temporary facilities are coordination sites that exist only for the duration of the work cycle (which in real-world problems may be of the order of a year in length). The IRPSF is discussed in more detail in Chapter 6.

CHAPTER 3

ASPECTS OF COORDINATION

In practice, the re-distribution of payload between depot and customer node is performed by a wide variety of fleet operators seeking to reduce operating costs (by reducing both total distance travelled and fleet size) or to increase their areas of operations (as in mid-air refuelling, fleet re-supply and courier work). Typical practitioners include stock-truck operators, couriers, roading manufacturers (see Figure 3.1), and the military. Factors influencing the application of coordination include

- the predicted operational benefits,
- types of payload being dealt with,
- types of sites at which coordination can occur,
- the method of applying coordination, e.g. one-way sequential multi-vehicle,

each of which will be discussed in more detail below. This chapter considers these factors and gives an indication of how they influence the practical application of coordination. It also discusses the creation of a model for the study of coordination. Theoretical aspects of coordination are dealt with in the following chapters: Chapter 4 deals with coordination in the Euclidean plane, Chapter 5 considers coordination in the rectilinear grid, and Chapter 6 gives an introduction to the application of coordination in the solution of dynamic vehicle routing problems. While many of the factors influencing real-world coordination are not included in the model that is applied later in this study, they are mentioned here to demonstrate the breadth and frequency of their occurrence, and to provide an indication of the practical implications of investigating coordination.



Figure 3.1: These bitumen carriers are transferring part of the load from one vehicle to a second. The first vehicle has travelled from the company's base in the east of the North Island of New Zealand to this siding at an intersection in the Central North Island. Having completed the transfer, the two vehicles will carry on to complete road maintenance in two different areas. The site chosen for the transfer is a layby at an intersection.

3.1 Predicted operational benefits

In addition to being an interesting mathematical phenomenon, coordination sites exist in a variety of physical situations. They may occur under circumstances which include a combination of the following reasons:

(a) Reduction in total distance

Incorporating coordination sites into routing schemes enables the overall distance to be reduced by reducing the common distances traversed by multiple vehicles, and by creating situations in which vehicles are not required to return to the depot to replenish stock other than at the end of a work cycle. These situations will be discussed in more detail in the following chapters.

(b) Reduction in fleet size

When using a coordination site increases the occupancy rate of a vehicle by allowing that vehicle to be used on more than one route in a work cycle, the

size of the vehicle fleet is reduced. The capacity utilisation is increased by having the vehicle more fully laden more of the time, or by re-using the vehicle on smaller routes.

(c) Increasing the area of operations of the concern employing the fleet

By removing the requirement of multiple vehicles to traverse common ground, the establishment of coordination sites permits vehicles to travel further from the depot and, consequently, a company can increase its area of operations to take in a wider client base. In the case of mid-air refuelling of aircraft, the transferred item is aviation fuel, which allows the recipient plane to extend its point of no return. For the situation shown in Figure 3.1, the bitumen supplier deploys one vehicle from its base in Napier, New Zealand, and is able to service customers in other parts of the country by transferring part of the load to other vehicles that cover a wider region than the original vehicle would be able to service in the limited time that hot bitumen is active.

(d) Repair

Repair scenarios provide examples of dynamic vehicle routing because at the time the vehicles are deployed the dispatcher is unaware of various attributes of the customers, such as their demand levels and often their locations. The location that repair vehicles and serviced vehicles meet can be considered a coordination site, and the materiel transferred is the expertise of the repairman and the goods required to perform the repair.

(e) Priority

The establishment of coordination sites can assist dispatchers in enabling high priority loads to return to the depot in pickup scenarios, or to be delivered to customers in delivery scenarios, more rapidly than other loads by transferring higher priority loads to vehicles which travel directly to the destination or to vehicles which travel more quickly because they visit only the high priority customers. Express couriers use this type of coordination site extensively.

(f) Vehicle modification

With the introduction of road-user charges, fleet controllers are encouraged to limit the number of axles and the mass of vehicles, and many find it advantageous to detach trailers that form part of the delivery vehicle, so that sections of the route are traversed by smaller and/or lighter vehicles. Heavy unladen trailers are expensive to transport and may hinder access in difficult terrain. Similarly, trailers which are laden with goods that will not be needed in a specific region have no requirement to enter that region if they

can be transferred to another vehicle or if they can be unhooked and later retrieved by the returning vehicle. Stock vehicle drivers use this method widely to avoid taking large trailers into some areas. Also, trailers can be unhooked to allow access for the main vehicle when it would not be able to use certain roads because of tight turns, small bridges and other hazards to navigation. Charikar [CKR98] has shown that such single-vehicle coordination can reduce the total distance travelled.

(g) Reduction in operating costs

Operating costs, which include running costs of vehicles, purchase costs, rental costs and special taxes, can be reduced by reducing the fleet size, reducing the total distance travelled, reducing the total time required for vehicles to be used, and increasing the utilisation of vehicles and drivers.

(h) Accommodation of split deliveries/pickups

Splitting demands [DT89] allows loads to be carried which can not fit into one vehicle. This form of coordination will be discussed in more detail below.

(i) Accommodation of demands that become available at different times

Individual customers sometimes have more than one request during each planning period. These requests do not necessarily occur at the same time; some happen after the vehicle has been despatched to deal with the first one. Coordination sites allow the customer to be serviced by fewer vehicles than the number of requests when some of those requests can be combined en route. Coordination sites also allow later vehicles to deviate from their existing routes to service such customers. The benefits to some customers must be weighed against inconveniencing others.

3.2 Payload requirements

Payloads can be classified into broad groupings called commodity classes by the attributes of their constituent goods, and commodity classes can be defined by considering either similarities or differences. For example, classes could be defined by

Definition 3.1: Goods that have the same carriage requirements fall into the same commodity class.

or by

Definition 3.2: Goods that cannot be carried together fall into different commodity classes.

Clearly, definitions 3.1 and 3.2 lead to different classes for given loads. A more rigorous definition of commodity classes must consider several attributes, including

- compatibility
- volume
- mass
- divisibility
- the need to return to the depot
- the need for specific vehicle type
- social acceptability
- legal acceptability
- support vehicles (police escort, wide load escort)
- en route processing
- urgency.

Each is discussed in more detail below with an indication of practical situations that pertain to it.

3.2.1 Compatibility

Some loads cannot be carried together solely because they are incompatible due to physical characteristics of the goods carried, political requirements of the customer or personalities of the live payload.

Some strongly smelling materials impart their odours to neighbouring matter, which can be unacceptable to customers. Raw fish is such an example.

Occasionally, customers will not use vehicles carrying competitors' products. This could be to prevent the opposition knowing the nature and timing of special product launches, or to inhibit the marketability of the competitor's product.

When carrying livestock, animals are sometimes disturbed by the presence of a particular species. Sometimes males of the same species are provoked by close proximity to each other, even when they are in the foreign environment of a stock truck. This problem can be overcome by modifying the vehicle (see below) but, in practice, the incompatibility of goods is a greater factor in the assignment of vehicles than the compatibility.

3.2.2 Volume

Loads must be able to fit onto a vehicle without overhanging beyond a reasonable length for reasons of safety, legality and manoeuvrability. In addition, the choice of route may be restricted by the carriage of a large volume. If the bulk of a load is too great, access may be limited by obstacles such as low overhead wires or overbridges. As a specific example, trailers carried on the Inter-Island ferries between the North and South Islands of New Zealand are often built to the maximum height permitted by the shipping company in order to carry as much payload as possible, even though the vehicles may travel the length of the country without such height restrictions.

3.2.3 Mass

When the mass of a load is excessive, the maximum laden weight of the vehicle on which it travels will restrict the vehicle to a particular type and the route may have to be modified to avoid weak road surfaces and small bridges, thus increasing operating costs. In some situations, such as supplying large engineering operations in remote locations, the roads themselves may have to be strengthened, or even constructed, to accommodate the heavy loads that must be transported to the sites.

Since some road licensing authorities tax vehicles according to the maximum laden weight, it is sometimes advantageous for operators to use larger fleets of lighter vehicles than smaller fleets of heavier vehicles, in order to reduce operating costs.

3.2.4 Divisibility

Split deliveries [DT89], where they are permitted, allow greater flexibility in routing but not all goods are capable of being split. Separable loads may also have limits put on the minimum size of a split and on the number of divisions that the customer will accept.

3.2.5 Need to return to the depot

Under some circumstances, not all loads actually need to return to the depot. Those that do not need to return to the depot could be grouped together and treated in a similar fashion, thus creating different versions of the routing problem.

3.2.6 Need for specific vehicle type

Some types of load require specific vehicle types. For example, the carriage of bulk flour requires large reinforced spheroidal metal containers, whereas 5kg bags of flour can be carried in a van. Some loads require refrigerated containers; special rules, e.g. [NZS88], normally govern the carriage of hazardous substances by determining the type of vehicle which can transport such goods. Secure money transportation from banks and EFTPOS machines is handled by armoured vehicles. Armoured vehicles used in the transport of prisoners are sometimes divided into two prisoner compartments so that incompatible passengers can be segregated.

3.2.7 Social acceptability

Some loads cannot be carried with others because of a social repugnance. An example of this is the carriage of food with coffins, which is not illegal and is not prevented under other criteria but which is not socially acceptable.

3.2.8 Legal acceptability

Legal acceptability involves the mixing of loads, the routes along which particular loads are prevented from travelling, and the times at which some goods must travel.

In New Zealand, as in other countries, several laws (e.g. [NZE57], [NZS88]) prevent certain types of load being carried with other types of load. For example, fireworks must not travel in the same vehicle as paying passengers.

Load toxicity [USHM90] prevents some loads being carried through residential or commercial areas. This includes household garbage, industrial waste and nuclear waste; local bylaws prevent loads such as uncovered offal from being transported on some roads. This prevents certain payloads from being transported through some areas.

Explosives [NZE57] are not allowed to be carried at night. In New Zealand, explosives and fireworks fall into seven classes, most of which are further classed into divisions. The classes are: gunpowder, nitrates, nitro-compounds, chlorates, fulminates, ammunition and fireworks. The United Nations Committee of Experts on the Transport of Dangerous Goods classifies "dangerous goods" into: gases when compressed, liquefied or dissolved under pressure; ethane, ethylene, hydrogen, methane, and any other flammable gas; acetylene, compressed or dissolved, and

contained within a porous substance; liquefied petroleum gas and any other liquefied flammable gas; chlorine; anhydrous ammonia; and liquid oxygen.

Where vehicles cross international or federal boundaries, the laws governing the carriage of goods for each state entered must be adhered to.

3.2.9 Support vehicles

The types of load that require escort vehicles include wide loads, high-security goods, naval fleet re-supply and royalty, and the routes used by such convoys are limited by physical characteristics as well as by disruption to normal traffic. They are also limited by the speed and manoeuvrability of the least able member of the convoy. Coordination problems involving convoys must include constraints on the number of support vehicles that can be removed from escort duty at any time.

3.2.10 En route processing

Some types of load require servicing while en route. For example, livestock needs to be watered, and long-haul passengers often require comfort stops. Such servicing is an integral part of the carriage of the goods and is sometimes accommodated by modifying the vehicle, such as including bathroom facilities on planes, trains and buses.

3.2.11 Urgency

Loads might not need to be collected or delivered in the current planning period; some may be high priority and need to be dealt with urgently, while others may be delayed in their collection or delivery.

3.2.12 Commodity classes in this study

In a multi-commodity coordination problem, as in single commodity systems, the coordination process is an improvement procedure. Where more than one commodity class is involved, coordination involving incompatible classes is prohibited. Consequently, there is no loss of generality in considering single commodity coordination. The inclusion of multi-commodity routing involves the addition of another variable in the decision-making process and can be considered as two problems with similar attributes occupying the same coordination space.

By defining commodity classes by their differences, rather than by their similarities, the question that needs to be answered when deciding whether vehicles can participate in coordination is "*Are these payloads proscribed?*" As this discussion has shown, the usual answer is that they are not, and that it is special circumstances that create most of the classes discussed above.

The models used in this study involve single commodity goods travelling on identical vehicles. The situation involving the divisibility of goods is considered later, but even there, all goods are divisible (or not) to the same extent.

3.3 Coordination sites

Definition 3.3: A coordination site exists in space and time for the transfer of materiel and/or personnel and/or information between vehicles.

The transfer at a coordination site may be initiated to accommodate payload, driver or vehicle requirements, in contrast to transfer at a node, in which goods alone are transferred between customer(s) and/or depot(s).

Coordination sites can be located at existing nodes, on part of a route or somewhere "close to" other routes. The suitable proximity of nodes and routes is discussed later. Load re-deployment at coordination sites is commonly used by couriers swapping loads on the side of the road; same-vehicle coordination [CKR98] sites are used by stock vehicles that drop a trailer somewhere and collect it later. This not only reduces the overall distance travelled by the complete vehicle, it may improve access and allow the vehicle to use a track that the complete vehicle and trailer cannot. The advantages have to be weighed against the extra distance incurred in retrieving the trailer.

Coordination sites are not limited by size, location or movement. They have spatial and temporal characteristics that play a role in determining their classification, and can be serendipitous, pre-determined or, in the case of dynamic coordination, post-determined.

Coordination sites can be as large as a sub-depot or as small as a single letter drop-point; they can be strips of roads, areas of seas or deserts, or volumes of air-space. They can be stationary, as in most land-based applications, or mobile as in the mid-air refuelling of aircraft.



Figure 3.2: Mid-air refuelling is a form of mobile coordination. Hercules KC-130B 720 refuelling two F-5 Tigers. (Image courtesy US Information Office)

Coordination has temporal as well as spatial aspects, which will be shown in Chapters 4 and 5 to be inter-dependent. Temporal considerations help define the type of coordination site as either “drop-off” or simultaneous; in either case, the recipient must not leave the coordination site before the donor arrives.

The concept of coordination sites includes all those sites known as transfer points, swap sites, exchange points, meeting places, transient nodes and temporary dump sites (used for short term storage and later retrieval).

Figure 3.3 indicates the breadth of categories which can be used to classify coordination. Individual instances of coordination can belong to more than one branch of the diagram, such as “one-way pre-determined simultaneous coordination at a node”. This discussion creates a hierarchy of coordination types by considering initially the type of meeting place at which they occur. This includes

1. satellite facility,
2. node,
3. existing route,
4. not on existing route,
5. mobile,
6. driver-defined, and

7. single vehicle.

Items 1-6 can be classified further by the direction of transfer. These broad groupings are discussed in more detail below.




Figure 3.3: Coordination can be categorised in many ways.

3.3.1 Satellite facility

In real-world scenarios, it is common practice for vehicles to begin and end their routes at a depot (usually the same depot from which they were deployed), although they may visit the depot during the work cycle to replenish stock. Consequently, each routing plan contains at least one depot. Where an operation involves the frequent need for replenishment, some strategies establish temporary satellite facilities as “sub-depots” - short-term storage facilities that can be visited many times by any number of vehicles but that are not used for between-cycle vehicle storage. Such a satellite facility is a coordination site because it exists solely for the transfer or replenishment of goods and need not involve contact with customers, and for routing and scheduling purposes can be considered as a node that has zero demand. The establishment and use of satellite facilities as coordination sites will be discussed in more detail in Chapters 5 and 6.

3.3.2 Node

In a non-coordinated system, a node is a customer location that is associated with a (possibly zero) demand and that is visited by vehicles as they collect and/or deliver goods for that node. In a coordinated system, nodes have the added characteristic that they can be used to store goods for other customers, acting as temporary depots, in which case the node and coordination site are said to be co-located. This study considers all nodes as capable of becoming co-located with coordination sites, although in practice nodes are acceptable only if they have appropriate attributes such as accessibility, space for the storage of goods, and agreement by the occupiers. In this study, goods can be left at nodes without the provision of support for the node although in real-world situations, the creation of a coordination site carries with it the requirement for securing the location against incursion by competitors. A typical example of this is the protection of temporary ammunition dumps.

The existence and identity of a node suitable to act as a coordination site is discussed in §3.8.

3.3.2.1 Sequential multiple vehicle

Sequential multiple vehicle visits to a node can occur when at least one of the vehicles does not completely service the demand of a node which occurs later in its route and which uses this multiply visited node as a temporary storage facility. Both vehicles servicing the later node are obliged to visit the temporary storage facility – one to deliver a partial demand, the other to collect that demand and deliver the combined demand for the later node. This incorporates the concept of split delivery in which a series of vehicles carry part-loads to and from customers. The node that acts as a temporary storage facility for the goods which are to be transferred becomes a coordination site. It is assumed that nodes which act as temporary storage facilities have the capacity to accommodate such transient loads.

Sequential multiple vehicle visits can also occur during dynamic vehicle routing when the requirements of the customer at a node change with time. As shown in Figure 3.4, while one vehicle (red route) is en route to servicing the advance request of a node, the demand is increased (at some time t), and the surplus is delivered by a second vehicle (blue route) to another node (filled disc) that the first vehicle later encounters, at which the part-loads for the target node are combined and eventually delivered by the vehicle originally assigned to the task. The node indicated by the filled disc is a coordination site and, because the

transfer is one way, there is no requirement for the vehicles to visit the coordination site at the same time, only that the vehicle collecting the load leaves the site after the delivery vehicle arrives there. This is discussed in more detail in the following chapters.




Figure 3.4: Sequential multiple vehicle at node.

When a prioritised delivery/pickup system is used, some items of payload are considered more important than others and are handled with greater urgency. Under these circumstances, the higher priority items are carried on vehicles that can travel more quickly between the depot and the appropriate nodes. Because of the time differences in vehicle departures from the depot, and the differing travel times to nodes because of distance, a high priority parcel may be able to be transferred from a high priority vehicle to a slower vehicle if that slower vehicle can actually attend the node more quickly. The transfer from the higher priority vehicle to the lower priority vehicle can take place at another node, which precedes the node with multiple priority demand, which then becomes a coordination site.

Where payload types can be divided into commodity classes, items fall into a single commodity class if they are able to be transported on the same vehicle at the same time. Consider item a and item b, each of unit size, belonging to commodity classes A and B, respectively, which must not be carried together but which do not require special vehicles to transport them. Now consider item c, which can be carried with items A or B. Once the vehicle (of maximum capacity 2 units) carrying a and c has delivered A, it can accommodate b. So item b can be left at a node which will later be visited by the vehicle carrying item c. Such a node then acts as a coordination site.

Sequential multiple vehicle coordination at nodes involving single commodity payloads in a pure delivery setting is discussed in more detail in the following chapters.

3.3.2.2 Simultaneous multiple vehicle

During simultaneous multiple vehicle visits to a node, many-way transfer can be made between vehicles. The node need not act as storage since the payloads are transferred directly between vehicles, but it must be able to accommodate all vehicles involved in the exchange at the same time. The simultaneous attendance of multiple vehicles at a node occurs under circumstances of split delivery, dynamic demand, multiple commodity deliveries and priority payloads.

During split delivery, the demand for a particular node may be divided among more than one vehicle. If both (or all) vehicles carrying such loads arrive at the customer at the same time, their remaining payloads may be re-distributed among them. In that case, the node has not acted as storage but has been a coordination site.

Under similar circumstances, the vehicles may be able to re-distribute their loads earlier in their routes if one vehicle is able to accommodate fully the previously split delivery by increasing its carrying capacity more quickly. This happens when such a vehicle delivers the equivalent of the remaining split load prior to reaching the intended recipient of the previously split load. The vehicles involved can then meet earlier in their routes at another node. This earlier node has not acted as storage but has been a coordination site.

During dynamic vehicle routing, a customer can have more than one demand, occurring at different times. If the vehicles that are allocated each delivery meet at the node at the same time (or close enough together in time that one vehicle does not have to wait too long), they can use this node as a coordination site and re-distribute their payloads. They might also be able to meet at a node earlier in the route of one of them, if capacity constraints permit transfer.

The coordination sites that occur under circumstances involving different commodity classes and of prioritised delivery/pickup discussed above in relation to sequential multiple vehicle visits to a node can also occur simultaneously.

Simultaneous multiple vehicle coordination at nodes involving single commodity payloads in pure delivery static and dynamic scenarios is discussed in more detail in later chapters.

3.3.2.3 Single vehicle

Single vehicle multiple visits to a node occur due to capacity constraints of the vehicle, time constraints on the customer (part of the load was not ready for uplift) or scheduling constraints.

Any demand that exceeds the carrying capacity of a vehicle must be carried either on a larger vehicle or in smaller portions; the latter form of delivery is known as single vehicle split delivery. During such a delivery, part of the load that is not intended for that customer may be left at the node while the vehicle returns to the depot for the remaining load. Consider the scenario in which a vehicle is to carry payload for nodes A and B, and in which the payload for node A far exceeds the capacity of the vehicle. On the first trip from the depot, the vehicle carries all the demand for B and part of the demand for A. It deposits both at node A and returns to the depot for the remaining demand for A. This demand is then delivered to A in one or more trips. On the last trip, after the final delivery to A, the payload for B is collected from its temporary storage at A and delivered to B. Node A has acted as a coordination site in a single vehicle multiple visit. The same node has appeared in both routes of a doubly-routed vehicle.

When customer requirements arrive at the depot in a dynamic fashion, an extra demand for a particular node might arrive after the vehicle servicing that node has departed. If the demand is to be satisfied, it is possible that the same vehicle will service the node more than once in a planning period. As in the single vehicle split delivery mode above, this single vehicle multiple delivery mode allows the node to act as a coordination site for payloads belonging to later nodes. Consider the scenario in which a vehicle is to carry multiple payloads for node A and a single payload for node B. On the first trip from the depot, the vehicle carries all the demand for B and the first delivery for A. It deposits both at node A and returns to the depot for the next demand for A. This demand is then delivered to A in one or more trips. On the last trip, after the final delivery to A, the payload for B is collected from its temporary storage at A and delivered to B. Node A has acted as a coordination site in a single vehicle multiple visit.

Similarly, when a prioritised delivery system is employed, the vehicle may be required to return to the depot before delivering its complete payload in order to collect higher priority demand that only becomes available after the vehicle has already begun its journey. Lower priority payload can be left at an earlier node and collected when the higher priority payload has been dealt with. The node at which such payload has been temporarily deposited acts as a coordination site.

3.3.3 Existing route

An *existing route* is a route that has been accepted into a feasible scheme; it may be adjusted or removed by subsequent improvements. Where a section of such a system is traversed by two or more vehicles, coordination may be used to improve the scheme by allowing the common juncture of these routes to be re-designed.

3.3.3.1 Chance meeting on an existing route (serendipity)

When two or more vehicles happen to traverse the same part of an existing route at the same time, transfer of goods can be made between vehicles. The meeting place becomes a coordination site.

3.3.3.2 Pre-determined meeting on existing routes

When two or more vehicles traverse routes that share a common region (see Figure 3.5), and are able to transfer stock feasibly, such interchange may be predicted and incorporated into the routing scheme. As in §3.3.2, the transfer may be simultaneous or sequential. If the intersection is simultaneous, the goods can be transferred directly from one vehicle to another; sequential intersection means that goods must be left unattended until the recipient arrives to collect them. In either case, the donor vehicle must arrive within a reasonable time of the recipient vehicle, otherwise the waiting time accrued by the waiting vehicle may diminish any advantage in travel time that such a transfer would confer.




Figure 3.5: Meeting on an existing route. The dotted area is able to act as a coordination site.

In Figure 3.5, two routes intersect at a point that is not co-located with a node, and the temporal characteristics are such that stock may be transferred between them. The intersection point, indicated by a dotted region, becomes the coordination site and, for planning purposes, is treated as a node with zero demand

that ceases to exist after all stock is removed from it. Coordination at such a site is identical to coordination at a node, and comments in §3.3.2 about coordination at a node are also applicable here.

The establishment and use of coordination sites on existing routes is discussed in more detail in later chapters.

3.3.3.3 Post-determined meeting on existing routes

A post-determined meeting is a meeting that is scheduled after at least one of the participating vehicles has left the depot, and coordination under such circumstances is applied in response to the changed nature of the system. The involvement of coordination in dynamic vehicle routing is discussed in Chapter 6.

3.3.3.4 Repair

Repair is a special case of a meeting on an existing route. It is related to the post-determined meeting in that the coordination must be arranged after the donor vehicle has left the depot but differs because such instances are highly undesirable, rare and handled in a standard manner. The coordination site may consist of a repair vehicle providing maintenance to restore the active vehicle to operation if such maintenance could be achieved without prejudicing the schedules, or it may consist of a replacement vehicle accepting all the remaining duties of the broken vehicle, in which case, the broken vehicle would no longer be active. The coordination site remains in existence until the broken vehicle is collected or repaired, the payload is reassigned and the driver leaves the site.

3.3.4 Not on existing route

Existing (standard) routes are not the only places available for travel; detours may also be possible. Coordination sites that occur away from existing routes cannot happen by chance since, by definition, vehicles travel on existing routes unless involved in coordination. Since the only manner in which vehicles involved can recognise the potential for such a meeting is by way of on-board communications, such meetings must be either pre-determined or post-determined.

3.3.4.1 Pre-determined meeting involving detour

Before the vehicles to be involved in such a meeting leave the depot, they are scheduled to deviate from the existing routes. Each vehicle then travels further

than the shortest distance between existing customers to facilitate a transfer which results in an overall benefit to the routing scheme.




Figure 3.6: Detour. Upper: Both vehicles detour. Lower: The red vehicle detours.

In Figure 3.6, coordination sites are established that require at least one vehicle to divert from its existing route and detour to the site. Because nodes usually take up only a minute area in any region of operations, it is more likely that the use of coordination sites will involve vehicles in deviating from their non-coordinated paths. This application of coordination will be discussed in more detail in later chapters.

3.3.4.2 Post-determined meeting involving detour

The meeting is scheduled after at least one of the vehicles has left the depot. The vehicle which has already departed the depot must be in communication with the depot for such a meeting to be arranged, either by on-board radio or by messages left at customer nodes visited earlier in the route. The incorporation of such sites into routing schemes is discussed in detail in Chapter 6.

3.3.5 Mobile

Coordination sites do not need to be fixed in space or in time; they can be mobile in any dimension.

3.3.5.1 Spatially mobile

Spatially mobile coordination sites exist when the transfer takes place “on the run”, and at least one of the vehicles does not stop. Such transfers exist in mid-air refuelling of aircraft, in which both the donor and recipient vehicles keep moving although are stationary with respect to each other (see [YHM90]).

3.3.5.2 Temporally mobile

Temporally mobile coordination sites exist when the transfer can occur at any time. Such transfers take place when, for example, a collection vehicle that does not necessarily interact with customer nodes meets a potential donor vehicle. The collection vehicle travels between depot(s) and donor vehicles. Meetings need not be scheduled at particular times or even at particular places. The collection vehicle “mops up” payload from other vehicles and can be thought of as a form of mobile sub-depot.

3.3.6 Driver-defined

A driver-defined coordination site exists because one or more drivers want to create it, rather than because of any requirement by the fleet router. Such a coordination site can be single vehicle or multiple vehicle, on or off existing routes, and permanent or temporary. A coordination site is permanent if it is used on a regular basis; otherwise it is temporary. Driver-defined locations can be incorporated into a routing scheme if they do not adversely affect the scheme by more than an acceptable amount.

3.3.6.1 Single vehicle driver-defined

Single vehicle driver-defined coordination sites may be created as comfort stops or to comply with legal requirements for maximum driving hours. Such sites are more likely than corresponding multiple vehicle stops to be on or near existing routes.

3.3.6.2 Multiple vehicle driver-defined

Multiple vehicle driver-defined coordination sites involve the synchronous meeting of more than one vehicle in a place chosen by at least one of the drivers. At some point in the meeting, more than one vehicle is present even though not all vehicles may reach the meeting point at the same time or leave at the same time. In addition to the uses of single vehicle driver-defined coordination sites, multiple vehicle stops can be used for the dissemination of information and for mutual support.

3.3.7 Single vehicle

A coordination site can be created by a vehicle not involved in a meeting with any other vehicle; the donor and recipient vehicles are the same. In this case, either part of the load is left for later uplift or part of the vehicle is left for later uplift. Single vehicle coordination sites can be created at nodes (see above), on existing routes or at detours between routes. Trailed vehicles and containerised loads are particularly suitable for use in single vehicle coordination sites because of the cost differentials involved in carrying trailers and accessibility problems associated with larger or heavier vehicles.



Figure 3.7: Single vehicle coordination. This logging trailer has been left at a layby on an existing route.



Figure 3.8: Single vehicle coordination involving detour. The freight operator in charge of this trailer has diverted from his existing route to deposit the trailer down a quiet side street in order to reduce running costs. The trailer is collected on the return journey.

3.3.8 Direction of transfer

Since coordination sites exist to facilitate the transfer of some entity, they can be classified further by the direction of transfer. This can be one-way, two-way or, in the case of more than two vehicles meeting, many-way.

3.3.8.1 One-way transfer

During one-way transfer, one vehicle is the donor and another is the recipient. The donor vehicle must arrive at the coordination site before the recipient leaves; the recipient must not leave until after the transfer has occurred. The carrying capacity of the donor is increased, while the recipient has reduced carrying capacity as a result of the transfer.

3.3.8.2 Two-way transfer

Two-way transfer can be thought of as a series of one-way transfers performed either in a synchronous or in a consecutive manner.

During synchronous two-way transfer, the entity to be transferred from one vehicle is passed at the same time as the entity to be transferred to that vehicle. Each of the two vehicles must be capable of carrying both entities since there is no buffer.

Consecutive two-way transfer is one one-way transfer followed immediately by another one-way transfer in the opposite direction. Thus, the initial donor becomes the second recipient and the initial recipient becomes the second donor. Each transfer can be directly into the donor vessel or it can be into a temporary storage facility at the coordination site in which case such a temporary storage facility becomes a buffer for the transfer. Donor and recipient vehicles need not change their carrying capacity after two-way transfer, since the volume of goods received may equal the volume donated.

3.3.8.3 Many-way transfer

Many-way transfer occurs when more than two vehicles meet at a coordination site. It is a combination of one-way and two-way transfer.



Figure 3.9: Three-way transfer between courier vans. The coordination site is the parking lot of an orchard and lies at the intersection of two major provincial roads.



Figure 3.10: Many-way transfer. A heterogeneous fleet of buses carries pupils from several small schools to a coordination site at one of the schools where the children redeploy themselves to the appropriate bus that then takes them from the coordination site to their home area. If coordination had not been employed here, each bus would have had to travel to every school and then to each home area.

3.4 Classification

Practicable coordination classifications are:

- sequential multiple vehicle at node
- simultaneous multiple vehicle at node
- serendipitous meeting on existing route
- pre-determined meeting on existing route
- post-determined meeting on existing route
- pre-determined detour
- post-determined detour
- mobile mop-up
- satellite depots
- single vehicle off-route
- single vehicle drop-off (at node or existing route).

This study does not approach the investigation of satellite depots or single vehicles. Chapter 4 considers both sequential and simultaneous multiple vehicle coordination at nodes in a Euclidean setting, and mentions briefly mobile mop-up. Chapter 5 considers the same in a rectilinear setting as well as pre-determined meetings on an existing route and pre-determined detours. Chapter 6 considers post-determined meetings on an existing route and post-determined detours in a dynamic rectilinear setting.

3.5 The coordination model

In order to investigate coordination in an analytical manner, a coordination model is developed that allows many of the aspects of coordination mentioned above to be studied. The remainder of this study is concerned with investigations based on the following coordination model.

In the static case, the system consists of a single depot with a known number of nodes, each with known location and delivery demand and each with an infinite capacity for the storage of goods. The fleet is made up of identical vehicles, each with a fixed payload capacity and maximum distance, which begin and end their routes at the depot. Vehicles are in constant communication with the depot and are able to vary their routes when required. The model has been developed for a pure delivery scenario but this can be modified to incorporate pure pickup or a combination of delivery and pickup. The payload consists of items of known (single) commodity class. The objectives are to minimise total distance travelled and to minimise the fleet size, although neither objective shall be achieved at the expense of the other unless specified. That is, fleet size should not be reduced if there is a concomitant increase in total distance travelled, and total distance should not be reduced if there is an accompanying increase in the number of vehicles. Performance is measured by the percentage reduction in total distance and the reduction in number of vehicles of the coordinated scheme over the scheme that exists without coordination.

In the dynamic case, at least one of the attributes of location, demand, and time for at least one node is unknown at the start of the work cycle. Otherwise, the dynamic model has the same assumptions, objectives and constraints as the static model.

In each case, coordination will be considered to be between multiple vehicles unless specified otherwise.

3.6 Performance indicators

Throughout this study, preference is given to non-coordinated solutions, and coordination is applied only where its use produces a solution with a stronger result. In determining the value of solutions, the following performance indicators are applied:

1. fleet size
2. total distance travelled by the fleet
3. number of coordination points required
4. number of customers affected by coordination
5. balanced routes

of which (1) and (2) are far more important than the rest. A slight increase in total distance is acceptable if fleet size is reduced, and zero reduction in fleet size is acceptable if the total distance is reduced. Where more than one solution is available, that solution involving the smallest number of coordination points is accepted. Similarly, solutions are sought that involve fewer customers in coordination. Although payloads are involved in each coordination event, a customer is involved only if its demand arrives split. Route balancing is not applied at the expense of fleet size and distance.

A highly desirable solution is one where a low total distance is travelled by a small fleet of vehicles, each travelling approximately the same distance, with few coordination points affecting few customers.

3.7 Spatial and temporal inter-dependence

Since multiple-vehicle coordination involves the transfer of materiel from one vehicle to another, the collection vehicle must not leave the coordination site before the delivery vehicle arrives. Several successful scenarios exist:

(a) Drop and go

The delivery vehicle arrives at the site before the collection vehicle; it drops the goods and leaves. The collection vehicle arrives later and collects the goods from the site without physical interaction with the delivery vehicle.

(b) Delivery vehicle waits

The delivery vehicle arrives before the collection vehicle and waits for the delivery vehicle to arrive. The goods are transferred directly from the delivery vehicle to the collection vehicle.

(c) Collection vehicle waits

The collection vehicle arrives before the delivery vehicle and waits for the delivery vehicle to arrive. The goods are transferred directly from the delivery vehicle to the collection vehicle.

(d) Simultaneous arrival

The delivery vehicle and the collection vehicle arrive at the site at the same time. The goods are transferred directly from the delivery vehicle to the collection vehicle.

The spatial and temporal characteristics are thus interdependent. This interdependence is exhibited by the importance of route order and orientation.

3.8 Route order and orientation

The order in which a particular vehicle services nodes is determined by two factors:

1. the route order travelled by the vehicle;
2. the orientation of the route.

Route order is the sequence in which a multi-routed vehicle services each route. For example, a vehicle that services routes 1, 2 and 3 can service them in any one of the following orders:

123	213	312
132	231	321

Table 3.1: Route order for a three-routed vehicle.

If the coordination occurs on route 2 (say), the time at which this can occur will vary depending on the route order being used. The other vehicle involved in the coordination will not be in the same place in every option in Table 3.1 and there is no guarantee that it will be in a position to interact successfully with the vehicle servicing route 2.

Orientation refers to the direction in which an individual route is travelled. All multi-noded routes have two orientations. For example, if route 2 consists of the depot (node 0) and nodes 3, 7, 12, 1 and 27, the two orientations are 0 3 7 12 1 27 0 and 0 27 12 1 7 3 0. If the coordination occurs near node 12, the time at which this can occur will vary depending on the orientation being used.

For the purposes of this study, orientation is defined with respect to the node identifiers at either end of the route and assumes that all routes are traversed as easily in either direction. In practice, this assumption is not always valid since it ignores one-way roads, difficult intersections and the effect of hills on running costs.

$$\text{orientation} = \begin{cases} -1 & (\text{finalnode} - \text{initialnode}) > 0 \\ 0 & \text{single-noded route} \\ +1 & (\text{finalnode} - \text{initialnode}) < 0 \end{cases}$$

The number of possible route order/orientation combinations is determined by Theorem 3.1.

Theorem 3.1

The number of possible schedules for a routing scheme with one depot, m vehicles of which n are multi-routed, servicing k multi-noded routes is

$$2^k(j_1!)(j_2!)\dots(j_n!).$$

Proof:

Consider a vehicle routing scheme consisting of one depot and m vehicles, of which $n < m$ are multi-routed, servicing $k < m$ multi-noded routes. Let j_i be the number of routes serviced by the i th multi-routed vehicle, $i=1, \dots, n$.

Then the number of orientations, x , of the k multi-noded routes is

$$x = 2^k.$$

The number of different route combinations, y , serviced by the multi-routed vehicles is

$$y = (j_1!)(j_2!)\dots(j_n!).$$

The number of possible route order and orientation combinations is

$$xy = 2^k(j_1!)(j_2!)\dots(j_n!).$$

■

Any of the resulting routing schemes might be suitable for coordination. Consider the routing scheme shown in Figure 3.11 of a single depot, ten node Euclidean system, where each colour represents a different vehicle. Two vehicles (purple and pale blue) are multi-routed, each route being single-noded; one vehicle

(red) is single-routed, single-noded; two vehicles (green and dark blue) are single-routed, multi-noded.




Figure 3.11: A single-depot, ten-node system in the Euclidean plane.

A snapshot of the system in Figure 3.11 will show the location of each vehicle at a particular point in time and space as the end of a coloured line. Two possible snapshots are shown in Figures 3.12 and 3.13, each relating to a different combination of route order and orientation. For the sake of clarity, the nodes have been removed and the traversed paths to single nodes have been exaggerated.




Figure 3.12: One snapshot of the system from Figure 3.11.

The vehicle routes for the scenario in Figure 3.12 are shown in Table 3.2. The snapshot is taken when the purple vehicle is en route to node 6 having completed the first of its two routes; the green vehicle has serviced nodes 10 and 2 and is returning to the depot; the red vehicle has completed its route and is back at the depot; the black vehicle has completed the first of its two routes and is on the second route, en route to node 4; the dark blue vehicle has serviced all the nodes on its only route and is returning to the depot.

Vehicle	Route
purple	0 8 0 6 0
green	0 10 2 0
red	0 7 0
black	0 9 0 4 0
dark blue	0 3 5 1 0

Table 3.2: Routing scheme for the system in Figure 3.12.

If the snapshot had been taken at the same time but with a different route order and orientation, the result might have been that shown in Figure 3.13.




Figure 3.13: Another snapshot. Same time, different route order and orientation combination.

The vehicle routes have changed in Figure 3.13 to become those in Table 3.3. At the instant of the snapshot, some of the vehicles are in different locations. In accordance with the theorem above, there are sixteen possible route order / orientation combinations for the system shown in Figure 3.11 of which only two have been shown here. Any, or none, of these might lead to a condition suitable for coordination to occur.

Vehicle	Route	Comparison
purple	0 8 0 6 0	No change
green	0 2 10 0	Other orientation
red	0 7 0	No change
black	0 4 0 9 0	Different route order
dark blue	0 1 5 3 0	Other orientation

Table 3.3: Changes to the routes caused by using a different combination of route order and orientation.

3.9 Existence and identity of a suitable node

Whether or not coordination can occur at a node depends on several factors. The two main approaches (node-based and route-based) take advantage of the neighbourhood structure of nodes and routes and depend on a node being close to another route. The definition of "close" is determined to be the value of the near radius, which is discussed later.

Having identified candidates for a coordination site, their suitability is determined according to the type of coordination being investigated. This is covered in detail in Chapters 4 and 5, but the aspects that are common to both Euclidean and rectilinear scenarios are mentioned here.

3.9.1 Node-based approach

In the node-based approach, each node is investigated to determine whether any routes come close to it. This approach is based on the assumption that, because the target node lies close to another route, a vehicle on that route would not have to detour very far to meet it. Where nodes are determined to be close to routes, they become candidates for coordination sites.

The approach is illustrated in Figure 3.14, in a Euclidean setting. The near radius (edged in pink) of the red target node encompasses part of a (green) link on another route. The link has been identified as interesting because the nodes on the other route that are closest to the target are contiguous and thus the environment for feasible insertion is favourable. This is discussed in more detail in Chapters 4 and 5.




Figure 3.14: Node-based approach.

Node-based approach	
Step One	Choose a route.
Step Two	Consider a node in that route (target node), that has not already been a target node.
Step Three	For each other route:
Step 3(a)	Find nodes on that other route that are closest to the
Step 3(b)	target.
	Are the close nodes consecutive?
	If so, find path joining them; go to Step 3(d).
	Otherwise, go to Step 3(a) for another route.
Step 3(c)	Repeat from Step One with another route that has not already been chosen.
Step 3(d)	Find the distance between the target node and the lines.
	Is the distance acceptable?
	If so, target may be a coordination site.
	Otherwise, go to Step Two for another node.

Procedure 3.1: Node-based approach for determining a coordination site at a node.

3.9.2 Route-based approach

In the route-based approach, each combination of routes is investigated to find the node that comes close to at least two routes.

Consider the scenario in Figure 3.15 in which the purple dashed line indicates the near radius for each route in a rectilinear setting. Every point inside the green region, is close to both routes and represents a collection of points of close approach. One node is close to this region and becomes a candidate for coordination site.




Figure 3.15: Route-based approach.

Route-based approach	
Step One	Find the list of close approaches for each orientation option.
Step Two	Order these close approaches in increasing order.
Step Three	Look at each close approach. Is the point of close approach near a node? If not, reject that point; repeat Step Three with the next on the list. Otherwise, go to Step Five.
Step Four	If no more points of close approach, repeat from Step One with another route order / orientation combination. If no more combinations, STOP.
Step Five	The node is a possible coordination site.

Procedure 3.2: Route-based approach for determining a coordination site at a node.

Occasionally, the closest node to a close approach lies on a route that is not involved in forming the close approach, as shown in Figure 3.16. By dividing the work cycle into discretised time steps and determining the vehicle positions and system status at each time, the path traversed by each vehicle is represented in the model by a series of discrete positions, and, because of this step-wise nature of creating the path from one node to the next, a close approach might be represented as a different point on each route.




Figure 3.16: The situation prior to exchange. The closest node lies on a third route.

At the coordination site, six part-routes exist. They are the part-routes representing the routes before- and after-coordination-site. The outcome depends on two factors:

- 1. the vehicle that carries the demand for the coordination site, and
- 2. the way in which the part-routes re-combine,

and the combination of these results in several options, shown in Table 3.4.

Option	Carrier of demand for coordination point	Part-route combination
C1	Vehicle A	Original (except Vehicle C).
C2	Vehicle B	Original (except Vehicle C).
C3	Vehicle A	Part-routes swap vehicles.
C4	Vehicle B	Part-routes swap vehicles.
C5	Vehicle C	Original.
C6	Vehicle C	Part-routes swap vehicles.

Table 3.4: Combination options.

In options C1 - C4, the coordination site is removed from its original route and the demand is carried by Vehicle A or Vehicle B; Vehicle C has a reduced payload.




Figure 3.17: Options C1 and C2.

Options C1 and C2 differ only in the carrier of the demand for the coordination site. In option C1, the demand is carried by Vehicle A; in C2, it is carried by Vehicle B.




Figure 3.18: Options C3 and C4.

Options C3 and C4 differ only in the carrier of the demand for the coordination site. In option C3, the demand is carried by Vehicle A; in C4, it is carried by Vehicle B. Options C1 and C3 differ in the post-coordination routing. In C1, the original route is adhered to; in C3, Vehicle A continues with the latter part of Vehicle B's original route. Options C2 and C4 differ in a similar fashion.

In options C5 and C6, the demand for the coordination site is carried by Vehicle C, which would have carried it had coordination not been applied. Vehicles A and B still visit the coordination site because that is the node closest to each of

their routes. Thus, the coordination site is attended by three vehicles, not necessarily at the same time.




Figure 3.19: Option C5.

Options C5 and C6 differ only in the post-coordination parts of their routes. In option C5, the original routes are adhered to; in C6, the latter parts are swapped.




Figure 3.20: Option C6.

3.9.3 Intersection-based approach

In the intersection-based approach, the aim is to find a node close to an intersection. This approach differs from the route-based approach in that the orientation and route order are determined *a posteriori*. The routing scheme such as that shown in Figure 3.21 is investigated to determine whether intersections exist without regard to orientation or to route order. If any intersections exist, each is considered in turn. When a suitable node has been located, the orientation and route order combinations that allow that intersection to occur are found. The resulting part-routes are the same as those in the route-based figures above.




Figure 3.21: Intersection-based approach.

In the scheme shown in Figure 3.21, the purple dotted line indicates the near radius of an intersection. Two nodes lie within this limit and thus, are close to the intersection. One of the nodes is on a route that forms the intersection. The other node lies on a third route.

3.9.4 Limitations

It should not be assumed that a node suitable for acting as a coordination site actually exists. In each of the approaches above, there is no guarantee that a node will be capable of hosting a coordination site.

This study considers coordination in both the Euclidean (Chapter 4) and rectilinear grid (Chapter 5) scenarios, each of which presents slightly different characteristics for the identification of suitable nodes. In the Euclidean plane, coordination sites are restricted to existing only at the locations of existing nodes, and the paths between nodes are unique straight lines, thus limiting the number of interactions that can exist between routes to one location per intersection. In contrast, the rectilinear grid allows paths between nodes to take the form of a

contiguous series of small links in the grid, thus allowing much greater choice in the actual path taken by a vehicle. With the increased choice of path there is a concomitant increase in the number of potential interactions, each of which is considered when investigating the existence of suitable locations for coordination. This is illustrated in Figure 3.18 which demonstrates the effect of having just two possible paths for each of two part-routes. Two of the instances contain no intersections; one of the instances has vehicles sharing a common section.




Figure 3.22: Paths in the rectilinear grid. Even with only two possible paths for each of the two routes, the environment for coordination is altered.

Also, because of the nature of the grid which allows nearby paths to contain parallel sections, interactions may occur over a region rather than at a point as in the Euclidean plane.

Consequently, the existence and identity of coordination sites varies depending on whether the scenario occurs in a Euclidean plane or rectilinear grid. For this reason, a discussion of the types of coordination that can be employed is dealt with in the appropriate chapters below.

CHAPTER 4

COORDINATION IN THE EUCLIDEAN PLANE

The aim of this chapter is to investigate the effect of coordination as a route improvement method in the Euclidean plane. This coordinate system has been chosen because it presents no limit on the location of points or paths, and to provide a basis for comparison with Euclidean vehicle routing problem schemes that do not use coordination. By allowing complete freedom in the location of nodes and paths, coordination sites can be created anywhere within the plane. On the other hand, the lack of restriction on vehicle paths greatly limits the number of direct path interactions since paths are less likely to have coincident sections, reducing chances of path overlap. Consequently, regions around paths are investigated (once path locations are known) for potential interactions with other paths. The next chapter deals with the rectilinear grid which forces path loci into restricted configurations and enables path interactions to be investigated more fully.

4.1 Node location

Nodes are located randomly at integer positions within a plane measuring 100 by 100 distance units with the exception that

- all coordinates are integral,
- no two nodes occupy the same location, and
- no node is co-located with a depot.

Restricting the coordinates to integers permits them to be represented relatively simply by matrices, and allows for straightforward implementation in the MATLAB programming language.

4.2 Euclidean distances

The standard Euclidean metric has been used to calculate symmetric inter-nodal distances. The distance between node A , located at (x_A, y_A) , and node B , located at (x_B, y_B) is given by

$$\text{dist}(A, B) = \sqrt{(x_A - x_B)^2 + (y_A - y_B)^2}.$$

4.3 Depot location

The depot can be located at any point within the Euclidean framework but it has been restricted initially to a constant position for consistency. Unless otherwise stated, systems being compared have the same depot location in order to prevent such location from posing an obfuscatory influence on the outcome of a routing improvement scheme.

To investigate the effect of depot location on route structure, twenty datasets were examined, each at five depot locations (results are shown in Table 4.1). When the depot is located at or near the centroid of nodes, routes are evenly balanced typically with one or two single-noded routes, several routes of four to six nodes, and with the fleet containing two to four multi-routed vehicles. When the depot is located outside the convex hull defined by the nodes, the distance between the depot and the first node on a route increases, resulting in reduced manoeuvrability for coordination. When the depot is located in a corner, well away from almost all nodes, the routes are unbalanced with few nodes per route and few (if any) multi-routed vehicles; most of the distance allowance is used just to travel between the depot and the first node of a route.

Depot location	Nodes per route	Fleet size	No. multi-routed vehicles
Centroid	[1,6]	[5,6]	[2,4]
Near centroid	[1,6]	[5,6]	[2,4]
Just outside hull	[1,5]	[6,9]	[1,2]
Near edge of plane	[1,3]	[8,17]	[0,1]
Near corner of plane	[1,3]	[8,20]	[0,1]

Table 4.1: Effect of depot location on routes for a single depot system of 25 nodes located randomly across the plane.

Since some coordination requires the existence of multi-routed vehicles, the placement of a depot at an extreme position (edge or corner) away from the nodes

was not pursued. A default location was adopted near the centroid to allow multi-routed vehicles to participate in coordination. As a bonus, this adds an air of realism since, in practice, it is not uncommon for depots to be situated slightly off-centre with respect to their clientele.

4.4 Customer demand

Each customer is assigned an integer-valued demand generated from a normal random distribution, which has mean 0 and variance 1, by the equation

$$Demand = \text{floor}(\text{abs}(Random \times 100)).$$

Where the demand is greater than the maximum vehicle capacity and load splitting is not permitted, the demand is reduced to the value of the maximum vehicle capacity. This allows a feasible solution to be reached and allows slightly more loads of larger value to be generated. An alternative scheme of generating only demands of value less than or equal to the vehicle capacity would also have been acceptable. A normal distribution is chosen, in preference to a uniform distribution, to simulate more closely the actual demands received by a courier company. Demands belong to the same commodity class and have equal priority.

4.5 Delivery times

The time required to deliver a customer's demand is a linear function of the demand and the distance travelled to reach the customer's node. To deliver x units of demand made up of y different items, the delivery time is

$$Delivery\ Time = A + Bx + Cy$$

where A represents the fixed time to make a delivery, and B and C are user-defined constants, typically of the order 0.1 and 0.05. This function acknowledges both the longer unloading time of loads that involve more items and the economies of scale of unloading bulky items, and the values of B and C have been chosen to reflect that.

4.6 Coordination opportunities in the Euclidean plane

The structure of a Euclidean plane does not limit the positioning of coordination sites. They can occur at existing nodes, on existing non-coordinated routes between nodes and at points away from such routes. Travel can occur anywhere within the plane, and in a non-coordinated system vehicles travel along edges joining pairs of nodes because this results in shortest paths. In a coordinated system, the smallest total distance travelled may include some edges that do not form the shortest paths between a pair of nodes, since vehicles may have to detour. The intersection will be within a convex hull defined by the depot and node locations as shown in Figure 4.1.




Figure 4.1: Euclidean convex hull. On the left, two pairs of nodes with connecting edges. On the right, a possible coordination site within the convex hull created by the four nodes.

Coordination sites at nodes have the advantage that their locations are known before coordination is attempted. An existing node can act as a coordination site in two ways:

1. The demand for that node can be carried by more than one vehicle. This is known as *split delivery*.
2. The node can act as a temporary distribution centre, where loads can be deposited for later uplift to other nodes.

The two mechanisms are quite different and not mutually exclusive. Because both temporal and spatial coordination between vehicles is important, both the routing and scheduling phases are affected. For the purposes of this study, any node can function as either (or both) type(s) of coordination site.

The assumption is that nodes already have the facility to accept goods, whether those intended for them or for other nodes. Access and availability are not

limited; there is a finite number of existing nodes. In contrast, there is an infinite number of possible coordination sites on existing routes.

Normally, the function of a node in a pure pickup or delivery system is to provide a vehicle or vehicles with goods or to accept goods from a vehicle; such goods are to be delivered to or from the depot. It is not the standard function of a node to store vehicles or to store goods for delivery to other nodes. In contrast, the node in a coordinated system can act as a temporary depot. The *nodal zone of influence* is the area defined by a minimum distance from a node within which coordination has no significant effect. That is, the region in which the act of coordination might as easily occur at the node. The value of a "significant effect" is determined as being within one unit in the smallest decimal place in the measurement of total distance. For example, a transfer of goods directly adjacent to a node could have occurred at that node with no loss of distance advantage. Without zones of influence, nodes would exist merely as infinitely small points in a plane and no path would intersect them other than to service them. The use of zones of influence confers finite areas on nodes and enables continuous edges to be considered as objects, which can be divided into discrete quanta.

The nodal zone of influence can be represented as a spatial and a temporal measure; each can be fractional or absolute. As a fractional spatial measure, the radius of the zone is a fraction of the length of the edge under consideration. As a fractional temporal measure, the radius is a fraction of the travel time along an edge. Two intersecting edges will have different fractional measures but the same absolute values. The spatial zone of influence is referred to as an horizon; the temporal zone is a threshold. During coordination, both horizon and threshold are important.



Figure 4.2: Zones of influence. Each node has a spatial zone of influence of (say) 0.3 unit radius. For the longer edge, the zone of influence has fractional value 0.08, and for the shorter edge 0.23. Activities occurring within the zone of influence can be considered as occurring at the node itself.

Similarly, in multiple depot systems, there is a depot zone of influence. However, inter-depot coordination and transfers are not considered in this study since it deals with single depot systems.

There is no loss of generality in considering coordination at nodes and coordination on existing routes together since routes are defined by the nodes through which they pass. Any position on an edge of an existing route can be defined by a value combining

- the nodes at either end of that edge, and
- the distance from one end of the edge.

Other coordination sites can be found by pre-determining the possible coordination points that are neither at nodes nor on existing routes. These become, in effect, nodes that have zero demand and with no requirement to be allocated to routes. In practice, pre-determined off-route sites are selected for their accessibility, size and convenience.

Coordination in the Euclidean plane can result in

- allowing load splitting to occur,
- allowing load swapping to occur,
- obviating or deferring the return to the depot for multi-routed vehicles, and
- the construction of a feasible solution that would not have been achieved by the route construction method without coordination, by using coordination as an interim step,

each of which is discussed in more detail below.

4.7 Model construction

The model construction consists of

- determining the initial data values, i.e. the numbers of nodes and depots, their locations and service requirements, and setting constraints, then
- constructing routes without the aid of coordination, followed by
- assigning routes to vehicles, followed by
- applying coordination as an improvement procedure.

These are discussed in more detail below.

4.7.1 Initial data

A number of customer nodes are created at random sites, each with known delivery demand. The distance between nodes is independent of direction of travel.

For illustrative purposes in this study, the fleet consists of identical vehicles, each with a maximum vehicle capacity, maxVehCap , (set at an arbitrary value of 200) and with a maximum distance, maxVehDist , (set at an arbitrary value of 220). These default values have been chosen empirically so that they force routes to contain no more than about seven nodes each and allow vehicles to service up to three routes. This size of problem allows vehicle and route interactions to occur with some consistency and to be identified easily. The effect of varying these default values is discussed in more detail later. Maximum distance refers to the total distance that may be travelled by a vehicle during the work cycle and may include more than one route.

Vehicles begin and end the work cycle at the single depot. Although vehicles leave the depot at the start of the work cycle, their time of return depends on the routes used. The effect of a staggered departure of vehicles is discussed later. Since all parts of the plane suffer identical conditions, travel times depend only on length of routes.

Load splitting is permitted with no limitation on the number of times that loads may be divided or on the size of partition. The demand of each customer must be fully satisfied.

4.7.2 Route construction

Nodes are allocated feasibly to routes in a route construction phase. A variety of route construction methods has been investigated, the effects of which will be discussed later. At the allocation stage, the vehicle capacity is more significant than the distance constraint since most vehicles are filled before their distance allocation is satisfied. In practice, this is common among carriers of bulky items such as bitumen and crates of soft drink bottles. The effect of varying the default values is discussed later. All nodes with non-zero demand must be serviced.

Where split deliveries are not permitted in the initial routes, once the initial routes are constructed they are improved to become reasonable examples of non-coordinated routes, allowing any subsequent improvement to be caused clearly by the incorporation of coordination and not due to the introduction of some other

improvement method. While metaheuristic implementations of these improvement strategies could be used, they were not considered here and may provide a fertile area for future research. The pre-scheduling improvement applied in the route construction phase might involve one or more routes:

- multiple route improvement
- inter-route node transfer
- inter-route node exchange
- inter-route path exchange
- split delivery
- single route improvement
- 3-opt exchange.

4.7.2.1 Route improvement by intra-route 3-opt exchange

During 3-opt exchange, three edges are broken and three are made, one edge is traversed in the opposite direction, and no routes are created or deleted. The 3-opt exchange mechanism incorporates 2-opt exchanges and has been found to give much better results than 2-opt exchange and to be more efficient with respect to computation time than 4-opt exchange. As an intra-route improvement, there is no adjustment to load.

In Figure 4.3, the three made edges are between node pairs (A,B), (C,D) and (E,F), and the three broken edges are between node pairs (A,D), (B,F) and (C,E). The exchange is acceptable if

$$\text{dist}(A,B) + \text{dist}(C,D) + \text{dist}(E,F) - \text{dist}(A,D) - \text{dist}(B,F) - \text{dist}(C,E) < 0$$




Figure 4.3: 3-opt exchange. Arrows indicate direction of travel.

4.7.2.2 Route improvement by inter-route node transfer

Inter-route node transfer consists of removing a node from the route to which it was assigned during route generation, and either inserting it into another existing route or creating a route for which it is the only node. This improvement method, therefore, can create and delete routes.

The method consists of selecting each node in turn and attempting to insert it into each other route and into its own individual route. Clearly, vehicle capacity constraints must be taken into account; it may be that demand for a node is so high as to preclude its entry into particular routes, since the vehicles servicing those routes would then be overloaded.

The arrangement is accepted which results in the lowest total distance travelled.

Consider the situation depicted in Figure 4.4 in which node A is examined for transfer from Route 1 to Route 2'.




Figure 4.4: Inter-route node transfer. Node A is transferred from Route 1 to Route 2'.

Inter-route node transfer is accepted subject to the usual constraints of capacity and distance as well as

$$\begin{aligned}
 \text{Change in total distance} &= \text{distance of made edges} - \text{distance of broken edges} \\
 &= \text{dist}(A,D) + \text{dist}(A,E) + \text{dist}(B,C) - \text{dist}(A,B) - \text{dist}(A,C) - \text{dist}(D,E) \\
 &< 0 \quad \text{for transfer}
 \end{aligned}$$

and

$$\text{Load}(\text{Route } 2) + \text{Demand}(\text{Node } A) \leq \text{maxVehCap.}$$

Inter-route node transfer
<pre> check each node as TargetNode for insertion into other route check each route as DestinationRoute IF the destination vehicle will hold TargetNode attempt to insert TargetNode into this route accept this insertion if favourable END IF NEXT DestinationRoute NEXT TargetNode remove empty routes check each node as TargetNode for creation of a new route attempt to create this new route accept this route if favourable NEXT TargetNode recalculate route distances redraw routes </pre>

Procedure 4.1: Inter-route node transfer procedure.

4.7.2.3 Route improvement by inter-route node exchange

Inter-route node exchange consists of exchanging a node on one route with a node on another route. The procedure is performed only if the total distance travelled is reduced. During exchange, four edges are removed and four are created; no routes are created or deleted. Consider the exchange shown in Figure 4.5 in which Target A on Route 1 and Target B on Route 2 are examined for exchange.




Figure 4.5: Inter-route node exchange. Target A and Target B on different routes are exchanged.

The four edges broken are:

- PrevA – Target A
- NextA- Target A
- PrevB – Target B
- NextB – Target B

and the four edges made:

- PrevA – Target B
- NextA – Target B
- PrevB – Target A
- NextB – Target A.

As in §4.7.2.1, an exchange is favourable if both distance and capacity constraints are observed:

$$\begin{aligned}
 &dist(PrevA, Target B) + dist(NextA, Target B) + dist(PrevB, Target A) + \\
 &\quad dist(NextB, Target A) - (dist(PrevA, Target A) + dist(NextA, Target A) \\
 &\quad + dist(PrevB, Target B) + dist(NextB, Target B)) < 0 \\
 &Load(Route 1) - Demand(Target A) + Demand(Target B) \leq maxVehCap \\
 &Load(Route 2) - Demand(Target B) + Demand(Target A) \leq maxVehCap.
 \end{aligned}$$

The exchange is accepted that most greatly reduces the total distance travelled.

Inter-route node exchange
<pre> BestDiff := 0 FOR each Target A node FOR Target B = Target A + 1 to numNodes IF Route A <> Route B then Made := total distance of the four made edges Broken := total distance of the four broken edges NewLoad1 := load(Route 1) - demand(Target A) + demand(Target B) NewLoad2 := load(Route 2) - demand(TargetB) + demand(Target A) Difference := Made - Broken IF (Difference<0) AND (NewLoad1<maxVehCap) AND (NewLoad2<maxVehCap) AND (Difference<BestDiff) then swap nodes Target A and Target B BestDiff := Difference take another look at all routes since some have changed END IF NEXT Target B NEXT Target A recalculate route distances </pre>

Procedure 4.2: Inter-route node exchange procedure.

4.7.2.4 Route improvement by inter-route path exchange

Inter-route path exchange involves the removal of a section of path from one route and its subsequent insertion either in another route or in its own newly created route. It differs from inter-route node exchange in both the length and inconsistency of paths being exchanged. Where node exchange involved only individual nodes, there is no requirement for the paths being exchanged to contain the same number of nodes. Four edges are removed and four are created. Consider the exchange shown in Figure 4.6 in which a path FirstA---LastA in Route 1 is examined for exchange with the path FirstB---LastB in Route 2. The four edges broken are:

- PrevA – FirstA
- NextA – LastA
- PrevB – FirstB
- NextB – LastB

and the four edges made are:

- PrevA – FirstB
- NextA – LastB
- PrevB – FirstA
- NextB – LastA.

An exchange is favourable if both distance and capacity constraints are observed:

$$\begin{aligned} & \text{dist}(\text{PrevA}, \text{FirstB}) + \text{dist}(\text{NextA}, \text{LastB}) + \text{dist}(\text{PrevB}, \text{FirstA}) + \\ & \text{dist}(\text{NextB}, \text{LastA}) - \text{dist}(\text{PrevA}, \text{FirstA}) - \text{dist}(\text{NextA}, \text{LastA}) - \\ & \text{dist}(\text{PrevB}, \text{FirstB}) - \text{dist}(\text{NextB}, \text{LastB}) < 0 \end{aligned}$$

or

$$\begin{aligned} & \text{dist}(\text{PrevA}, \text{LastB}) + \text{dist}(\text{NextA}, \text{FirstB}) + \text{dist}(\text{PrevB}, \text{LastA}) + \\ & \text{dist}(\text{NextB}, \text{FirstA}) - \text{dist}(\text{PrevA}, \text{LastA}) - \text{dist}(\text{NextA}, \text{FirstA}) - \\ & \text{dist}(\text{PrevB}, \text{LastB}) - \text{dist}(\text{NextB}, \text{FirstB}) < 0 \end{aligned}$$

or

$$\begin{aligned} & \text{dist}(\text{PrevB}, \text{FirstA}) + \text{dist}(\text{NextB}, \text{LastA}) + \text{dist}(\text{PrevA}, \text{FirstB}) + \\ & \text{dist}(\text{NextA}, \text{LastB}) - \text{dist}(\text{PrevB}, \text{FirstB}) - \text{dist}(\text{NextB}, \text{LastB}) - \\ & \text{dist}(\text{PrevA}, \text{FirstA}) - \text{dist}(\text{NextA}, \text{LastA}) < 0 \end{aligned}$$

or

$$\begin{aligned} & \text{dist}(\text{PrevB}, \text{LastA}) + \text{dist}(\text{NextB}, \text{FirstA}) + \text{dist}(\text{PrevA}, \text{LastB}) + \\ & \text{dist}(\text{NextA}, \text{FirstB}) - \text{dist}(\text{PrevB}, \text{LastB}) - \text{dist}(\text{NextB}, \text{FirstB}) - \\ & \text{dist}(\text{PrevA}, \text{LastA}) - \text{dist}(\text{NextA}, \text{FirstA}) < 0 \end{aligned}$$

$$\begin{aligned} & \text{Load}(\text{Route } 2) - \text{sum}(\text{Demand}(\text{FirstB}) \dots \text{Demand}(\text{LastB})) + \\ & \text{sum}(\text{Demand}(\text{FirstA}) \dots \text{Demand}(\text{LastA})) \leq \text{maxVehCap} \end{aligned}$$

$$\begin{aligned} & \text{Load}(\text{Route } 1) - \text{sum}(\text{Demand}(\text{FirstA}) \dots \text{Demand}(\text{LastA})) + \\ & \text{sum}(\text{Demand}(\text{FirstB}) \dots \text{Demand}(\text{LastB})) \leq \text{maxVehCap} \end{aligned}$$

Since there is no restriction on the number of nodes in the paths being exchanged, FirstA and LastA might be co-located (similarly, FirstB and LastB), and either or both might be the depot. In order to preserve the presence of the depot in each resulting route, whenever the depot participates in path exchange, it must be included in each exchanged path. That is, if the depot is in the interior of FirstA...LastA (or FirstB...LastB), it must be included somewhere in FirstB...LastB (FirstA...LastA); if the depot is co-located with FirstA or LastA (FirstB or LastB), it

must be included somewhere in FirstB...LastB (FirstA...LastA), not necessarily as FirstB or LastB (FirstA or LastA).


Both orientations of each path must be considered and the orientations adopted that minimise the sum of the distances between the endpoints of the paths being exchanged and the connecting points of the remnants. The orientation of Route 1' is determined by

$$\min(\text{sum}(\text{dist}(\text{PrevA}, \text{FirstB}), \text{dist}(\text{NextA}, \text{LastB})), \text{sum}(\text{dist}(\text{PrevA}, \text{LastB}), \text{dist}(\text{NextA}, \text{FirstB})))$$


and the orientation of Route 2' is determined by

$$\min(\text{sum}(\text{dist}(\text{PrevB}, \text{FirstA}), \text{dist}(\text{NextB}, \text{LastA})), \text{sum}(\text{dist}(\text{PrevB}, \text{LastA}), \text{dist}(\text{NextB}, \text{FirstA}))).$$

These calculations, must be performed for all paths, which also vary in length from one node to (route length-1), making the process computationally expensive. The resulting routes are subjected to intra-route adjustment by 3-opt exchange.



(a) Before exchange



(b) After exchange

Figure 4.6: Inter-route path exchange. The path FirstA...LastA on Route 1 is exchanged with the path FirstB...LastB on Route 2, producing four possible arrangements of Routes 1' and 2'.

Inter-route path exchange

```

FOR TargetRoute1 := 1 to NumRoutes-1
  FOR i := 1 to length(TargetRoute1)
    TargetNode1 := Routes(TargetRoute1,i)
    RouteLength1 := length(Routes(TargetRoute1,*))
    FOR PathLength1 := 1 to RouteLength1-1
      FirstA := TargetNode1
      IF (i+PathLength1) < RouteLength1
        LastA := Routes(TargetRoute1,i+PathLength1)
      ELSEIF PathLength1 == RouteLength1-i+1
        LastA := depot
      ELSE LastA := Routes(TargetRoute1,PathLength1-1-i)
      ENDIF
    TargetRoute2 := TargetRoute1 to NumRoutes
    FOR j := 1 to length(TargetRoute2)
      TargetNode2 := Routes(TargetRoute2,j)
      RouteLength2 := length(Routes(TargetRoute2,*))
      FOR PathLength2 := 1 to RouteLength2-1
        FirstB := TargetNode2
        IF (j+PathLength2) < RouteLength2
          LastB := Routes(TargetRoute2,j+PathLength2)
        ELSEIF PathLength2 == RouteLength2-i+1
          LastB := depot
        ELSE LastB := Routes(TargetRoute2,PathLength2-1-j)
        ENDIF
        Calculate loads
        IF loads OK
          calculate the four distances
          choose smallest distance
          IF distance OK
            store result
          ENDIF
        ENDIF
      NEXT PathLength2
      choose best result for this TargetNode2
    NEXT j
  choose best result for this PathLength1

```

```
    NEXT PathLength1
    choose best result for this node on route1
  NEXT i
NEXT TargetRoute1
take best overall result
IF best overall result exists
  make the exchange
  update Routes
  reset TargetRoute to 1
  re-start
ENDIF
```

Procedure 4.3: Inter-route path exchange

4.7.2.5 Route improvement by split delivery

Splitting the demand for one node enables the demand for a second node to be serviced by a route other than that which would have serviced it in an un-split environment, thus enabling a possible reduction in total distance travelled and, where the second node was on a single-noded route, reducing the number of routes. Since such routes might be serviced by single-routed vehicles, the prospect of fleet reduction is inherent in split delivery routing even if this is not part of the acceptance criteria.

When split deliveries are permitted in the initial routes, a demand is split in such a way that the largest part of the demand that can be carried on the target vehicle is assigned to that vehicle. Hence, the target vehicle is filled by accepting part of a split delivery. The remaining portion of the demand is available for assignment to another vehicle.

A split is not performed if it would result in a split demand being carried on a route servicing only one node since this would mean that more than one vehicle visited that node when one was compelled already to do so. In such a case, the total distance is reduced not by splitting the demand but by assigning a new vehicle to carry that demand to the node.

Where there is a choice of which demand to split, the larger load is accepted intact and the smaller load is split. This inherently limits the number of times a load will be split. If there were customer-determined limits on splitting, they would be accommodated in preference to this stratagem.

Several methods can be used for applying splitting. Some researchers [HR85] consider the splittable demand of size m as m demands of unit size and route these as m individual demands originating at the same node. This is computationally expensive and inefficient, since most of these will be recombined on delivery.

In this study, splitting is performed in the route construction phase only in order to fill a vehicle. The minimal fleet size is determined in advance by

$$FleetSize \geq \text{ceil}[\text{sum}(\text{Demand})/\text{maxVehCap}].$$

Vehicles are filled sequentially, each feasibly allocated unsplit demands and, where no intact demand can fit into the remaining free payload capacity, filled by splitting the next best feasible demand. This ensures that most vehicles are fully laden. Thus, no two routes will have more than one split demand point in common.

The fleet sizes of split and non-split feasible solutions coincide [DT90] when all nodes have identical demand of either 1 unit or maxVehCap . The greatest dissimilarity in size between split and non-split systems occurs when nodes have demands slightly greater than $0.5 \times \text{maxVehCap}$, in which case the split delivery fleet has a size one greater than half the non-split size.

$$|V_{split}| = \begin{cases} |V_{non-split}|, & d_i = 1, \text{maxVehCap} \\ \frac{|V_{non-split}|}{2} + 1, & d_i = \frac{\text{maxVehCap}}{2} + \varepsilon \quad \text{small } \varepsilon > 0 \end{cases}$$

Split delivery during route construction

```

initialise variables
assign minimum fleet size, minFleet
List := {nodes ordered by current method (eg savings, distance, angle)}
FOR route r := [1..minFleet]
  freeSpace := 0
  routeLoad := 0
  unassignedNodes := [List(1)..List(numNodes)]
  assignedNodes :=  $\phi$ 
  WHILE routeLoad  $\leq$  maxVehCap
    FOR each node i  $\in$  unassignedNodes
      IF demand(node i)  $\leq$  freeSpace AND distance OK
        assign node i to route r
        append node i to assignedNodes
        remove node i from unassignedNodes
        freeSpace := freeSpace - demand(node i)
        routeLoad := routeLoad + demand(node i)
      ELSE IF demand(node i) can be split AND distance OK
        demand1(node i) := freeSpace
        demand(node i) := demand(node i) - demand1(node i)
        assign node i to route r with demand = demand1(node i)
        freeSpace := 0
        routeLoad := routeLoad + demand1(node i)
      ELSE
        start new route
    END IF
  NEXT node
END WHILE
NEXT route
apply improvement (k-opt, exchange)
FOR each single-noded route
  IF node on single-routed node has split demand
    un-split it
    give it its own route (removes multiple travel to a node)
  ENDIF
NEXT single-noded route

```

Procedure 4.4: Split delivery during route construction

Consider the system shown in Figure 4.6 in which the demand for node C is capable of being split into $Demand^*(C)$ and $Demand^{**}(C)$. Should the division of $Demand(C)$ proceed, the demand for node A will be carried on a route other than that which would have serviced it in an un-split environment. The indication of nodes between the depot and nodes A and B does not imply the existence of those nodes; that is, nodes A and B could have been adjacent to the depot. Similarly, node C need not be contiguous with the depot. Four edges are broken and four are made; if node A had been on a single-noded route, the number of routes would have been reduced by one. The four broken edges are:

- A – PrevA
- A – NextA
- B – NextB
- PrevC – NextC

and the four made edges are:

- PrevA – NextA
- A – NextB
- C – PrevC
- C – NextC.

A comparison of these distances forms part of the acceptance criteria:

$$dist(PrevA, NextA) + dist(A, NextB) + dist(C, PrevC) + dist(C, NextC) - \\ dist(A, PrevA) - dist(A, NextA) - dist(B, NextB) - dist(PrevC, NextC) < 0$$

$$Demand^*(C) + Demand^{**}(C) = Demand(C)$$

$$Load(GreenRoute) + Demand^*(C) < maxVehCap$$

$$Load(BlueRoute) + Demand^{**}(C) < maxVehCap$$





Figure 4.7: Split delivery.

Like the form of split delivery implemented by Dror and Trudeau (see [DT89]), who used it as an improvement method without explicitly recognising it as a form of coordination, the purpose of splitting deliveries is to reduce the total distance without violating capacity or distance constraints.

This implementation of split delivery differs from that used by Frizzell and Giffin (see [F92]) in which upper and lower bounds were set on the size of the allowed split. Here, with no restrictions on the size of the division of any demand, vehicle carrying capacity can be fully utilised. The possible future addition of bounds on the size of demand division should not interfere with the concepts established here.

In this study, the use of split deliveries is allowed to occur between the depot and any node or between any two nodes, unlike other studies in which the splitting occurs only between depot and node, subject to the constraints of feasibility, capacity and distance.

Since split deliveries can be used in both route construction (Procedure 4.4) and route improvement (Procedure 4.5) phases, a comparison of the respective efficiencies is in order. The application of split deliveries in route construction occurs as a two-step process:

1. demands are split initially without regard to the number of nodes on the route to which the node with the split demand is being allocated,
2. the splitting may be reversed if the node with split demand is later found to be on a single-noded route.

In between these steps is an improvement phase, which seeks to refine the initial routes. During this improvement-in-construction phase new single-noded routes might be formed that later will be considered for removal or acceptance. Because of the existence of this improvement-in-construction phase, there is no practicable difference in applying coordination in the construction or improvement phases. The shift away from pure construction to a hybrid construction-improvement destroys any difference between applying split deliveries in an initial route construction or as the first improvement in a formal improvement phase. The effect of the order of application of different types of coordination is discussed later. In order to test the effect of splitting deliveries in the construction and improvement phases, fifty problems were compared, results of which are shown in Table 4.2. The construction techniques were:

- modified savings, Sav
- modified sweep, Swp
- nearest insertion, NI
- furthest insertion, FI
- modified savings including split delivery, SavSD
- modified sweep including split delivery, SwpSD
- nearest insertion including split delivery, NISD
- furthest insertion including split delivery, FISD

each of which took into account the loading and distance constraints. The insertion Procedures used a neighbourhood size of three. Each technique was applied to

every problem, followed by the improvement by split delivery. Three node densities were investigated: sparse (up to 15 nodes), medium (up to 30 nodes) and dense (more than 30 nodes). The results are given in Table 4.2 and in Figure 4.8. Average values include those cases where no splitting occurred.

Split delivery during route improvement
<pre> determine splittable demands FOR each route feasibly construct route using unsplit demands until no more nodes can be fully accommodated on current route NEXT route apply improvement (eg k-opt, exchange) assign minimum fleet size, minFleet IF fleet size > minFleet apply splitting FOR route r := [1..numRoutes] freeSpace := maxVehCap - routeLoad(route r) IF routeLoad(route r) < maxVehCap find best node to insert into route r by splitting calculate new total distance IF new total distance < old total distance accept new arrangement update variables ENDIF NEXT route ENDIF FOR each single-noded route IF node on single-noded route has split demand un-split it give it its own route ENDIF NEXT single-noded route apply improvement </pre>

Procedure 4.5: Split delivery during route improvement

Node density	Construction technique	Initial distance (min, med, max)	Improvement (%)
sparse	Sav	336,446,544	1.4
	Swp	351, 482, 544	1.5
	NI	332, 491, 591	1.5
	FI	336, 504, 608	1.6
	SavSD	334,440,529	0
	SwpSD	335, 474, 536	0
	NISD	327, 484, 541	0
	FISD	330, 496, 543	0
medium	Sav	765,1043,1361	3.1
	Swp	791,1226, 1350	3.4
	NI	767,1151, 1385	3.4
	FI	765,1181,1402	3.5
	SavSD	753,1011,1250	0
	SwpSD	741, 1184, 1247	0
	NISD	759, 1112, 1252	0
	FISD	761, 1140, 1256	0
dense	Sav	1232,1428,1623	4.1
	Swp	1232, 1555, 1625	4.4
	NI	1232, 1570, 1640	4.5
	FI	1232, 1612, 1656	4.6
	SavSD	1123,1369,1615	0
	SwpSD	1121, 1487, 1610	0
	NISD	1131, 1499, 1617	0
	FISD	1128, 1538, 1621	0

Table 4.2: Effect of route construction technique on average values for fleet size and total distance. All problems had the same value of maxVehCap and maxVehDist.

Generally, construction based on a modified savings approach results in an initial routing scheme with lower total distance than the other construction techniques (swapping, nearest insertion, furthest insertion). Each method found the lowest distance at least once and several methods found the same result using split delivery. The final routing schemes involving split deliveries showed far less variation across the techniques than did the initial schemes.




Figure 4.8: Effect of split delivery during construction and improvement using eight different construction methods. The lines indicate minimum and maximum values, colours denote node density: red –sparse; green – medium; blue - dense.

The inclusion of split deliveries in the construction phase has no benefit over the application of this form of coordination in the improvement phase because of the improvement that occurs within the construction phase as part of the split delivery process.

Initial routes that were dispersed did not respond well to the application of splitting since the distance involved in travelling between the established routes and the coordination sites was excessive. For this reason, schemes involving sparse node density had lower improvement rates than the denser schemes and, when improvement was achieved, the percentage reduction in distance was less.

The greatest average reduction in distance was achieved by systems involving many nodes. This effect is discussed later.

Greater distance reduction was achieved when the splitting involved nodes geographically distant from the remainder of their route and with nodes distant from the depot.

Splitting tended to occur with nodes having high demand. Although this trend was not searched for specifically, it was obvious in the results. This is in accordance with the results of other researchers (e.g. [DT89], [B93]) who have expressly investigated this behaviour.

The following figures show the effect of applying split deliveries to systems from each of the three node densities (load plans are given in Appendix C).




Figure 4.9: Sparse node density (10 nodes) without (upper) and with (lower) split deliveries. Distance is reduced by 3%.




Figure 4.10: Medium node density (30 nodes) without (upper) and with (lower) split deliveries. The split delivery solution has a 9% reduction in total distance travelled and a reduction in routes from 13 to 11. Route distances are more equitable.




Figure 4.11: Dense node density (75 nodes) without (upper) and with (lower) split deliveries. In the split delivery solution, the total distance decreases from 835.3 to 824.8 (1.3%). Node identifiers have been omitted for clarity.

4.7.3 Route scheduling

Once routes have been constructed, they are assigned to vehicles in a route scheduling phase in such a manner as to minimise the fleet size. The sum of all route distances assigned to a vehicle must not exceed the maximum distance that any vehicle may travel, maxVehDist . There is no other maximum distance for an individual route.

Initially, each route that is longer than a value equal to half the maximum vehicle distance, maxVehDist , is assigned to a new vehicle. This is done because any route longer than $\text{maxVehDist}/2$ will not be able to share with any other route that also is longer than $\text{maxVehDist}/2$ without violating the distance constraint. The route scheduling phase is designed to allocate routes to vehicles, not to fragment those routes into smaller routes.

After the larger routes have been assigned to vehicles, the smaller routes are assigned to these existing vehicles in such a manner as to maintain equitable vehicle distances. This phase determines the fleet size.

During the allocation of nodes to vehicles, where an individual demand exceeds the remaining vehicle capacity and where split deliveries are permitted, the over-large demand must be split in such a way that the largest part of the demand that can be carried on the current vehicle is assigned to that vehicle.

If split deliveries are not permitted, a node is allocated to a vehicle only if the entire demand can be accommodated by the remaining free space on that vehicle. If the demand cannot be accommodated on a current vehicle, a new vehicle is assigned to service the node.

4.7.4 Incorporation of coordination

After the initial routing and scheduling phases, the resultant system is investigated for the potential for coordination, which may involve re-routing and re-scheduling. Individual types of coordination are discussed later after which an analytical scheme is presented by which systems can be considered for those forms of coordination. The order in which coordination types are investigated is discussed in §4.11 Problem Design.

Route scheduling

```

longRoutes := {all routes with distance  $\geq \text{maxVehCap} \times 0.5$ }
unassignedRoutes := [1..numRoutes]
fleet V := {vehicles}
FOR route r  $\in$  longRoutes
  assign route r to its own vehicle  $v_r$ 
  append r to assignedRoutes
  remove r from unassignedRoutes
  vehDist( $v_r$ ) := distance(route r)
NEXT
WHILE unassignedRoutes
  choose the route  $r \in$  unassignedRoutes with the largest distance
  IF  $\exists$  routes  $\{r\}$  s.t.  $\text{vehDist}(v \in V) + \text{distance}(\text{route } r) < \text{maxVehDist}$ 
     $r := r$  with maximum value of  $(\text{vehDist}(v \in V) + \text{distance}(\text{route } r))$ 
    assign route r to vehicle v
    vehDist(v) := vehDist(v) + distance(route r)
    remove r from unassignedRoutes
  ELSE
    assign route r to its own vehicle  $v_r$ 
    vehDist( $v_r$ ) := distance(route r)
    increase fleet size by 1
    remove r from unassignedRoutes
  ENDIF
END WHILE

```

Procedure 4.6: Route scheduling.

4.8 Size of near radius

The *near radius* is that distance within which sites are investigated for coordination. It enables investigation in the plane to be limited to a much smaller area. By limiting the search for suitable sites to be within some arbitrary distance of nodes and existing paths, the exploration takes place in the most likely places and ignores unrealistic locations.




Figure 4.12: The near radius is indicated by dotted lines.

In Figure 4.12, with the near radius indicated by dotted lines, it is clear that node C lies outside the near radius of the intersection between AC and BD and that only node D lies close to (i.e. within the near radius of) a route other than its own. Consequently, when investigating nodes near intersections, all these nodes can be ignored; when investigating nodes near edges, only node D need be considered. Similarly, if a node is outside the near radius of an intersection, then the intersection is outside the near radius of that node. Thus, node-based and intersection-based investigations return the same results. Increasing the size of the near radius in Figure 4.12 eventually would allow nodes to fall within the near radius of the intersection.

A small value of the near radius limits the number of coordination opportunities, perhaps unrealistically by ignoring the possibility of coordination involving a node that would prove suitable if investigated using a larger value. The results of employing too small a near radius may therefore be incomplete. A large value greatly increases the computation time required and may involve the

investigation of coordination at nodes that could not be effective in reducing vehicle distance or fleet size.

The effect on problem solutions of the size of the near radius is investigated for 10 problems involving 50 nodes in a Euclidean plane of dimensions 100×100 units, altering only the size of the near radius for each problem.


Size of near radius	Distance reduction (%)	No. candidates considered	Computation time (relative)
2	0	0	0.03
5	0.4	1	0.31
10	2.3	5	1.00
15	2.3	12	2.19
20	3.7	18	3.54
25	3.7	23	3.74
30	3.7	32	5.67

Table 4.3: Effect of varying the size of the near radius.

Increasing the near radius allows more interactions to be investigated, thus increasing the computation time but not necessarily increasing the value of the result. Having a small value for the near radius provides a faster result but risks missing an interaction that would have proved fruitful if the value of the radius had been larger. Thus, the value of the distance reduction increases with increasing radius up to a point beyond which no more valuable interactions will be uncovered.

Consider the 50 node system shown in Figure 4.13 (load plans are given in Appendix C) in which nodes near intersections are investigated for their potential coordination site suitability. With the near radius set to 20, 18 nodes are investigated. As the radius is decreased from 15 to 10, nodes 1, 5, 29, 32, 46, 49 and 50 are no longer investigated, without change in the final result because those nodes are not involved in coordination. Yet when then radius is increased to 20, six extra nodes are investigated resulting in an improved solution. Since no further improvement is found, there is no benefit in increasing the value of the near radius above 20 for this problem.








Figure 4.13: Changing the value of the near radius.

4.9 Orientation and route order

The interactions between nodes and routes depends not only on the spatial characteristic of the location of nodes but also on the temporal characteristic of location of vehicles. This is particularly important in load swapping where vehicles must at some stage be co-located, and in load transfer where the order of departure from a coordination site is critical. Each multi-noded route may be traversed in two directions, denoted +1 and -1 depending on the relevant size of the initial and final node identifiers, and multi-routed vehicles may service each route in any order. Although the total distance travelled is unaffected, potential interactions, and consequently potential reductions caused by coordination, depend on the most favourable combination of route order and orientation being adopted.

Consider the system shown in Figure 4.14 in which four vehicles service ten nodes on five routes. Because there are four multi-noded routes, there are 2^4 combinations of orientation; the one multi-routed vehicle has two possible route orders. Thus, there are 32 feasible combinations of orientation and route order, each of which must be investigated for interactions which may provide a fertile environment for the application of coordination in an attempt to reduce the total distance travelled. The number of close (i.e. within the near radius) approaches between vehicles is investigated by approximating the path of each vehicle by a discrete series of time steps. For the (time step, near radius) value of (0.01, 20) any combination of orientation and route order yields between 30 and 38 close approaches, stored as a *Close List* of the form (Vehicle A, Vehicle B, time step #,

distance apart). An example of a Close List is given in Table 4.5 and a snapshot from this combination is given in Figure 4.15.




Figure 4.14: Overall view of feasible routes. Colours indicate vehicles. The plum-coloured vehicle is multi-routed.

Vehicle A	Vehicle B	time step #	approach	comment
1	2	2	5.56	(b)
1	3	2	0.75	(a)
1	4	2	16.58	
2	3	2	6.27	(b)
2	4	2	19.04	(c)
3	4	2	16.15	
1	2	3	11.12	
1	3	3	1.50	(a)
2	3	3	12.55	(b)
1	2	4	16.68	
1	3	4	2.25	(a)
2	3	4	18.82	(b)
1	3	5	3.02	(a)
1	3	6	11.50	(a)
2	3	12	13.58	(b)
3	4	12	19.30	
2	3	13	3.58	(b)
2	4	13	11.76	(c)
3	4	13	9.29	

2	3	14	0.00	(b)
2	4	14	0.70	(c)
3	4	14	0.70	
2	3	15	0.00	(d)
2	4	15	10.70	(c)
3	4	15	10.70	
2	3	16	0.00	(d)
1	2	17	16.59	
1	3	17	16.59	(a)
1	4	17	19.09	
2	3	17	0.00	(d)
2	4	17	19.45	(c)
3	4	17	19.45	
1	2	18	6.59	
1	3	18	6.59	(a)
1	4	18	8.77	
2	3	18	0.00	(d)
2	4	18	9.45	(c)
3	4	18	9.45	

Table 4.4: Close List for orientation = $(-1 \ 1 \ -1 \ -1 \ 0)$ and route order = $(4 \ 5)$. (a): Vehicles 1 and 3 have many early close approaches because they travel in almost the same direction. (b): Vehicles 2 and 3 have close approaches because they arrive back at the depot at similar times. (c): Vehicles 2 and 4 have close approaches because vehicle 2 is back at the depot when vehicle 4 moves from the depot towards node 5. (d): Depot effect. Time step 1 has been omitted since all vehicles are at the depot.

Since all vehicles are obliged to travel through the depot, a large number of interactions can be expected to occur there. The number of close approaches increases near intersections where the intersections occur in temporal proximity. At the time of the snapshot in Figure 4.15, vehicles 2 and 3 had returned to the depot, vehicle 4 had reached node 5 and vehicle 1 was returning to the depot.




Figure 4.15: Snapshot of system in Figure 4.14 and Table 4.5.

4.10 Number of nodes

Changing the number of nodes in a system alters the environment by affecting the node density and distribution. It can influence the physical length of routes and the population size of routes. The influence on the effectiveness of coordination of increasing the number of nodes depends in part on the construction method employed. As the density of nodes increases, clustering methods tend to reduce the number of intersections but increase the number of close approaches; non-clustering methods tend to increase the number of intersections.

When there are more nodes, there are more routes, and most of those will not change. Because only a few routes are usually involved in coordination, there is generally a greater percentage improvement in total distance travelled when there are fewer routes in the load plan. As routes become longer, however, the sections involved in coordination assume a smaller influence over the whole route. Similarly, the presence of fully laden single-noded vehicles reduces the overall reduction in distance because these vehicles are usually resistant to the effects of coordination. The presence of nodes that are geographically isolated from others in their routes provides a rich environment for load splitting but, as the number of nodes increases, such isolated nodes appear less frequently because geographical isolation tends to decrease in systems with more nodes.

Because all routes are obliged to travel through the depot, greater reduction in total distance often can be achieved when coordination occurs at sites distant from the depot. When the number of nodes increases, more nodes exist at locations further from the depot thus increasing the probability of having longer edges

between depot and initial (or ultimate) node. The longer exposed edge provides greater opportunity for the interactions with other nodes and with other routes. When the form of coordination being applied relies on the presence of intersections, those intersections that are further from the depot often provide a more fertile environment for the application of coordination.

Although most nodes are not actually involved in forming coordination sites, many must be implicated in concomitant considerations. The number of nodes that can be investigated by coordination is limited by the computation method employed. This study has investigated systems of 250 nodes. The use of some functions in the MATLAB language has limited the number of vehicles to fifteen, which is acceptable, because the concept of coordination and its application can be demonstrated well within this limit.

The extent to which the environment influences the type of coordination effective in reducing distance travelled and/or fleet size is discussed below.

4.11 Problem design

As noted earlier, there are four main types of coordination found in the Euclidean plane: load splitting, load swapping (and load transfer), deferral (or obviation) of return to depot, and the creation of metastable states, each of which is discussed in more detail below. There is a strong environmental influence on the determination of the form of coordination that is most successful in reducing total distance travelled and/or fleet size. Most problems are designed using random data, but some have been constructed with particular environments and attributes to test the behaviour of specific forms of coordination.

Coordination type	Required environment	Resultant environment
Load splitting en route	• Close approach of routes	• Possibly fewer intersections
Load swapping	• Intersections • Concurrent approach of vehicles	• Intersections
Obviation of return to depot	• Intersections • Multi-routed vehicle	• Fewer routes
Formation of metastable state	• Intersections	• Possibly fewer intersections

Table 4.5: Relationship between environment and coordination type in the Euclidean plane.

Interactions between the types of coordination applied to a particular problem can both create and destroy opportunities for further use of coordination. More coordination opportunities occur where the resultant environment caused by the application of one type is a fertile environment for another type of coordination. Similarly, the application of coordination that reduces the number of intersections can prevent another type of coordination being used which requires the existence of intersections. Load swapping is used to reduce the fleet size and to allow feasible solutions where they might not otherwise exist, and occurs when the total load can be shared by fewer vehicles. Obviation of return to depot affects only multi-routed vehicles. The use of coordination as an interim step between two feasible solutions applies only when intersections are present in the initial solution. Whereas load swapping creates intersections, both load splitting and formation of the interim step tend to reduce the number of intersections. The application of load swapping can create an environment in which load splitting and interim step generation can be used effectively to reduce the total distance and/or the fleet size. The effect of varying the order of applying coordination type is shown in Table 4.6. Twenty-five problems were created in which any of the four types of coordination could have been applied; solution techniques differed only in the order in which coordination types were applied.

An example of the effectiveness of the order of application of coordination type is shown in the system of Figure 4.16 in which 25 nodes are serviced by 6 vehicles. Since the fleet is already at the minimum size, coordination is applied in order to attempt to reduce the total distance travelled. Load plans are given in Appendix C.

Order	% Reduction in total distance	Most effective sections
ABCD	8.2	A, B
ABDC	8.2	A, B
ACBD	7.1	A, C
ACDB	5.9	A, C
ADBC	4.8	A, D
ADCB	4.9	A, D
BACD	5.7	A, B
BADC	4.7	B
BCAD	5.7	A, B
BCDA	5.7	A, B
BDAC	4.7	B
BDCA	4.7	B
CABD	5.9	A, C
CADB	5.9	A, C
CBAD	7.0	B, C
CBDA	7.0	B, C
CDAB	6.0	C
CDBA	6.0	C
DABC	5.7	A
DACB	5.7	A
DBAC	4.5	D
DBCA	4.5	D
DCAB	4.8	D
DCBA	4.8	D

Table 4.6: Effect of order of application of coordination type. Percentage reduction refers to the average reduction of value relative to non-coordinated plans for 25 problems. A = load splitting, B = load swapping, C = obviation of return to depot, D = interim step generation.

The most effective reduction is delivered when load splitting (A) is applied first followed by load swapping (B) unless either would produce infeasible results. Subsequent applications of methods C and D did not significantly improve the load plan, regardless of the order of applying C and D.

There was no noticeable benefit in re-applying a technique once its initial application had been followed by another type. Consequently, improvement techniques such as ABACBD have not been reported.

In some instances, conditions resulting from the application of methods A and B are inappropriate for the application of a third or fourth type of coordination. An example is provided in Figure 4.16 in which load splitting removes all intersections, other than one involving a single-noded route with demand equal to maxVehCap , thus depriving the system of a suitable environment in which an interim step could be formed. When the initial feasible solution is improved by means of interim step formation, load splitting has a reduced influence on the result.

Unless stated otherwise, the following default values apply.

- There is one depot.
- Maximum vehicle capacity, maxVehCap , is 200.
- Maximum vehicle distance, maxVehDist , is 220.
- Demands are integers in the range $[0, \text{maxVehCap}]$.
- All demands are known in advance.
- Fleet size is determined by the initial scheduling phase and is not increased unless doing so would result in a large decrease in the total distance travelled.
- Vehicles are homogeneous with respect to their operating characteristics.
- The number of nodes is fixed at the start of an investigation unless the system is stochastic.
- Near radius is 20.
- Time step is 0.01.
- Space step is 0.01.
- Load splitting is permitted.
- Coordination occurs at existing nodes or on existing routes.




Figure 4.16: Effect of order of application of load splitting and creation of interim step.

4.12 Load Swapping

Load swapping is employed in order to reduce the number of routes, thereby reducing the total distance travelled and/or reducing the fleet size. When

$$\text{sum}(\text{Total Load}) \leq \text{Max Load} \times \text{Fleet Size},$$

the total load could possibly be shared between a smaller number of routes without overloading any vehicle, if the actual demands permit it, without splitting any load. If load splitting is allowed, demands can be splintered subject to size constraints created by the splitting criteria. The spare capacity on the receiving vehicle and the load (or load portion) to be swapped must be compatible. This form of coordination can be used for both single-routed and multi-routed vehicles.

A consequence of applying this method to multi-routed vehicles is that the inter-route return to the depot may be removed, even though the motivation for the coordination is a reduction in number of routes and/or number of vehicles rather than removing the return to depot. The employment of coordination in order to remove an inter-route return to depot is discussed later.

The potential for load swapping increases where a node lies close to another route. This allows that other route to detour through the node without greatly increasing its distance. The coordination site must be able to accommodate all vehicles involved in the swapping procedure.

Since it is unlikely that both vehicles involved in load swapping will arrive at the coordination site simultaneously, one will be forced to wait for the other. The waiting time induced must be acceptable to the vehicle forced to wait.

4.12.1 Reduction in total distance

Where a node is excluded from a vehicle by load constraints, demand delivery at another node creates a reduction in the total load that needs to be carried and allows the previously excluded load to be accommodated by load swapping.

Consider the example shown in Figure 4.17. The system consists of a single depot and ten nodes with demand units arbitrarily integrally divisible. That is, a load of 95 can be considered as 95 separate units. Before coordination, the nodes are serviced by three routes and a fleet of two vehicles. The total load of 397 units is just within the 2×200 limit (where maxVehCap has value 200) so there is a possibility that it could be shared between two vehicles servicing two routes rather

than the non-coordinated three routes. The results of scheduling are given in Table 4.7.

Route #	Distance	Load	Route
1	103.350	108	0 2 7 6 0
2	173.987	194	0 4 10 8 1 9 0
3	88.157	95	0 5 3 0
	365.494	397	

Vehicle #	Distance	Route(s)
1	174	2
2	192	1 3

Table 4.7: Without coordination. maxVehCap = 200 units.




Figure 4.17: Before application of load swapping. Arrows indicate direction of travel.

Using the node-based approach of Procedure 3.1, each node is investigated for its potential as a coordination site, with nodes that are within the near radius of routes (other than those that serve them) selected as candidates for coordination. The effects of coordination are compared among these candidates and the most effective is selected as the actual coordination site. In the case under consideration, Nodes 7 and 9 are candidates for coordination, and each must be investigated with regard to each route orientation resulting in nine possibilities:

1. Coordination at node 7, orientation -1,-1
2. Coordination at node 7, orientation -1,+1
3. Coordination at node 7, orientation +1,-1
4. Coordination at node 7, orientation +1,+1
5. Coordination at node 9, orientation -1,-1
6. Coordination at node 9, orientation -1,+1
7. Coordination at node 9, orientation +1,-1
8. Coordination at node 9, orientation +1,+1
9. No coordination.

The coordination resulting from the most favourable possibility (1-4) is adopted if the non-coordinated plan provides a less favourable result. Possibility 7 is discussed in detail:

The route taken by Vehicle One is shown in red, the route for Vehicle Two is blue. When node 9 is considered as a candidate for coordination point, the route chains between depot and coordination candidate are

- Vehicle One (red): 0 9 ... ,
- Vehicle Two (blue): 0 2 7 9 ... ,

so that by the time Vehicle Two enters the coordination site, it has already off-loaded demand(2) and demand(7). The remaining nodes can now be allocated to the vehicles by a process of progressive lading. Since the demands for nodes that are encountered before the coordination site must travel on their original vehicles, demands for nodes 2 and 7 must be carried by Vehicle Two. The demand for the candidate could be carried by either vehicle. The demands for the other nodes can now be distributed between the two vehicles without regard to their original schedules since both vehicles will travel through the candidate node. Allocate loads to vehicles by putting the largest unallocated demand onto the vehicle with the greatest free space as in Table 4.8, allocating the route segments to Vehicles One and Two, followed by improvement by node exchange and k-opt. Thus

- the route chain for Vehicle One (red) stays at 0 9 1 8 10 4 0,
- the route chain for Vehicle Two (blue) becomes 0 2 7 9 6 3 5 0.

Vehicle A	sub-total	Vehicle B	sub-total
demand(8)=107	107	demand(6)=100	100
demand(4)=70	177	demand(5)=75	175
demand(1)=13	190	demand(3)=20	195
demand(10)=1	191		

Table 4.8: Progressive lading at node 9.

Vehicles One and Two are now loaded. From the depot, Vehicle One carries the demands for nodes 3, 4, 6, 8 and 9, and Vehicle Two carries demands for nodes 1, 2, 5, 7 and 10. At node 9, the loads are re-distributed with vehicle one giving up the demands for nodes 3 and 6 and collecting the demands for nodes 1 and 10. Vehicle Two carries a load of 197 up to the candidate and 195 away from the candidate; it is never overloaded. The resulting routes are shown in Figure 4.18. Vehicle One travels the same route as it would have in the non-coordinated system in Figure 4.17, although the payloads on leaving the depot are different. Because Vehicle Two no longer returns to the depot between routes, its distance travelled is reduced by 13.7%.




Figure 4.18: Load swapping at node 9.

One vehicle has been forced to wait for the other. Assuming that one time unit is equivalent to one distance unit, the waiting time is calculated as

$$\begin{aligned} Dist(Depot, Node\ 2) + Dist(Node\ 2, Node\ 7) + Dist(Node\ 7, Node\ 9) &= 47.358 \\ Dist(Depot, Node\ 9) &= 29.158 \\ \text{difference in distance} &= 18.200\ \text{units} \\ &\approx 18.2\ \text{time units} \end{aligned}$$

which may be acceptable.



The other possibilities are investigated in the same manner and the results compared:

Possibility	Reduction in distance	
1	0%	Original routes
2	7.9%	0 2 7 9 1 8 10 4 0 0 7 6 3 5 0
3	5.7%	0 6 7 5 0 0 4 10 8 1 9 7 2 0 0 3 0
4	0.7%	0 7 9 1 8 10 4 0 0 6 7 2 5 3 0
5	9.4% (Vehicle Two), 4.9% (overall)	0 4 10 8 1 9 0 0 2 7 9 6 3 5 0
6	1.0% (Vehicle Two) 0.8% (overall)	0 4 10 8 1 9 0 0 6 9 7 2 3 5 0
7	13.7% (Vehicle Two) 7.2% (overall)	0 9 1 8 10 4 0 0 2 7 9 6 3 5 0
8	13.7% (Vehicle Two) 7.2% (overall)	0 4 10 8 1 9 0 0 2 7 9
9	0%	Original routes

Table 4.9: Results of coordination to reduce distance travelled.

The greatest reduction is achieved by Possibility 2, which is accepted as the final route plan.




Figure 4.19: The preferred exchange plan.

Distance reduction is more likely where the original routing structure allows at least one of the vehicles involved in coordination to be near empty when the load swapping occurs.

Distance reduction is more likely where coordination occurs early in the route rather than later because a greater part of the route is affected by the exchange. Nodes serviced before the swapping occurs are not involved in improvement by swapping.

A reduction in total distance does not necessarily imply a reduction in the number of routes.

Distance reduction is improved where nodes that are close together are serviced by the same vehicle.

The existence of a multi-routed vehicle is not a necessary condition for distance reduction by load swapping as the following example demonstrates.

Consider the scheme shown in Figure 4.20 and Table 4. of a single depot and nodes with demand units arbitrarily integrally divisible. Only those nodes affected by coordination are shown.




Figure 4.20: Single-routed fleet before reduction in total distance by load swapping.

Route	Distance	Load	Route
1	152.28	190	0 2 3 4 5 0
2	195.54	136	0 1 6 7 0
347.82			

Table 4.10: Single-routed fleet before reduction in total distance by load swapping.

Using the node-based approach of Procedure 3.1, node 5 is identified as a possible coordination site. Using the techniques described above, coordination by load swapping at node 5 results in the plan shown in Figure 4.21 and Table 4.11 with a distance saving of 4.7%.




Figure 4.21: Reduction in distance by load swapping at node 5. Arrows indicate direction of travel.

Route	Distance	Load to node 5	Load away from node 5	Route
1	162.30	190	200	0 5 4 3 2 1 0
2	169.14	136	118	0 5 6 7 0
	331.44			

Table 4.11: Distance reduction by load swapping at node 5.

At node 5, demand(5) is delivered by the red vehicle, creating enough free space to accommodate demand(1) which is then transferred from the blue vehicle. Both vehicles continue on their adjusted routes. Neither vehicle has experienced any waiting time other than that required to deliver demand(5) and to transfer demand(1) from the blue vehicle to the red vehicle. Node 1, which had been excluded from the red vehicle because of capacity constraints, can be carried without such violation since delivery of demand(5) increases the available capacity. ■

4.12.2 Reduction in fleet size

A reduction in fleet size without a concomitant increase in distance can be achieved through load swapping by improving the distribution of loads over the fleet. Even routes that are not involved in the exchange of payloads are affected by the coordination as the vehicles that have been allocated to them are re-assigned and the fleet size reduced. Consider the route structure shown in Figure 4.22 - nodes are labelled a to j ; individual demands are shown next to the nodes; colours represent vehicles; the maximum vehicle capacity is 200 mass units and the maximum vehicle distance is 220 distance units.




Figure 4.22: Before coordination.

Without coordination, three vehicles are required to service the four routes since no combination of the four individual route distances allows a smaller fleet to service all routes. Table 4.12 gives pre-coordination scheduling details.

Route #	Distance	Load	Route
1	176.23	116	0 i j a f e 0
2	134.21	186	0 b c 0
3	36.76	200	0 h 0
4	87.71	191	0 g d 0
434.91			
Vehicle	Distance	Route(s)	
1	176.23	1	
2	134.21	2	
3	124.47	3 4	

Table 4.12: Before coordination.

Using the node-based Procedure 3.1 and a relatively small near-radius, the only candidate for coordination point is node e. However, using this node produces no improvements in fleet size or in distance travelled without violating distance and capacity constraints, so node e must be rejected as a coordination site. Increasing the near-radius allows node b to be considered as a coordination site. In this

example, no other node is realistically close to any second route. Since the sum of the loads carried by the blue and green vehicles does not exceed the value of $2 \times \text{maxVehCap}$, both vehicles can participate in load swapping subject to distance constraints.

When considering node b as a candidate for coordination site, both orientations of routes 1 and 2 need to be considered, resulting in five possibilities.

Possibility	Orientation	Reduction	
		Fleet	Distance
1	-1, -1	1	2.1%
2	-1, +1	No swapping	
3	+1, -1	1	2.1%
4	+1, +1	No swapping	
5	No swapping	0	0%

Table 4.13: Possible combinations for investigation of load swapping at node b .

Possibilities 1 and 3 produce suitable results: orientation -1,-1 produces a result by load swapping; orientation +1,-1 produces the same effect by load splitting without swapping. The load plans are shown in Table 4.14 and Figure 4.21.

		Vehicle	Distance	Load towards node <i>b</i>	Load away from node <i>b</i>	Route
Possibility 1 Orientation -1,-1	Demand(<i>b</i>) is delivered intact	1	168.81	200	95	0 <i>b f a j i</i> 0 <i>h</i> 0
			+ 36.76			
		2	132.43	132	126	0 <i>b e c</i> 0 <i>g d</i> 0
			+ 87.71			
			425.71			
Possibility 3 Orientation +1,-1	Demand(<i>b</i>) is split into 105 + 6	1	168.81	200	0	0 <i>i j a f b</i> 0 <i>h</i> 0
			+ 36.76			
		2	132.43	132	95	0 <i>b e c</i> 0 <i>g d</i> 0
			+ 87.71			
			425.71			

Table 4.14: After coordination at node *b*. Fleet reduced by one, distance reduced by 2.1% overall, 3.0% in affected routes.




Figure 4.23: Load swapping reducing the fleet size. Arrows indicate direction of travel.

Only one combination (possibility 1) of route order and orientation allows two vehicles to pass close to or through node b at about the same time, enabling node b to act as a coordination point without imposing severe waiting time on the other vehicle. In the resulting load plan, delivery of demand(b) frees enough carrying capacity to enable demand(e) to be delivered by the green vehicle. The demands for all nodes serviced by these two vehicles can be distributed between the vehicles at the depot in any manner subject to capacity constraints and re-distributed at the coordination site. Consequently, on leaving the depot the blue vehicle carries demands for nodes i , j , a and b , a total of 200 units; the green vehicle carries demands for nodes c , a and f , totalling 132 units. Demand for node b is delivered and the vehicles swap demands for nodes f and e . The green vehicle now has ample capacity to carry the demand for node e . Since removal of this node would reduce the blue vehicle's distance, node e is adopted into a new green route. The reduction of 3.0% in distance travelled on routes 1 and 2 is enough to permit two vehicles (rather than the original three) to combine routes 1-4 without violating distance constraints. The new route 1 is shorter than the original route 1, the new route 2 is slightly longer than the original route 2, but the reduction in route 1 outweighs the increase in route 2 and allows route 1 to be combined with route 3

so that the blue vehicle can service both routes without exceeding the maximum distance.

In possibility 3, the load for the coordination site, node b is split between the two vehicles subject to divisibility constraints. Up to 105 units can be carried by the blue vehicle, leaving at least 6 units to be carried by the green vehicle. If the load does not have to be re-combined before delivery to the customer, the green vehicle will not have to wait for the other vehicle before delivery.

In each case, the fleet size is reduced by one without increasing the total distance travelled.



Usually, load swapping requires that the coordination site not be at the end of a route. An exception to this occurs where the vehicle involved is able, via coordination, to reduce the number of routes that it services by combining them and thus removing the requirement for an inter-route return to depot.

A structural characteristic common to many scenarios that prove receptive to improvement by load swapping is the slight overlap of routes in the original non-coordinated scheme. The creation of a coordination site and the subsequent improvement by node exchange and k-opt reduces the overlap.

4.12.3 Transfer

Transfer of loads can be considered as a special case of load swapping in that it is a one-way exchange of payload from the donor vehicle to the recipient vehicle. In addition to the capacity and distance constraints applied above, there is a time constraint: the recipient vehicle cannot leave the coordination site until the donor vehicle has transferred the goods.

A variation on this transfer theme is the concept of a roving vehicle that is not constrained to begin and end each work cycle at the depot. Rather, it can move freely among the customer nodes and surrounding territory, collecting payload from vehicles and transferring to other vehicles or to nodes. This can prove particularly useful where a route containing an isolated node distant from the other route-members rearranges on coordination to a smaller route that does not contain that node; the isolated node is incorporated into another route. This is demonstrated in the example below, where the green route consists of a cluster of nodes as well as an isolated node (node 9). The routing scheme is shown in Table 4.15 and Figure 4.24.




Figure 4.24: Node 9 is distant from the other nodes serviced by the same vehicle.

The isolated node is unable to be accommodated into nearby routes because of capacity constraints (red route) and distance constraints (blue route).

The vehicle servicing the red route is able to travel further only after delivering its payload; the vehicle servicing the blue route is able to carry more but unable to travel further without violating distance constraints. The structural characteristic of:

- isolated node in an otherwise clustered route
- vehicle with spare carrying capacity but no spare distance capacity
- vehicle with no spare carrying capacity but spare distance capacity
- the above occurring within some "suitable" distance of each other

is a fertile area for successful improvement by coordination.

Route #	Distance	Load	Route
1	186.27	172	0 8 7 6 5 4 3 2 1 24 25 22 20 21 0
2	147.76	198	0 27 16 15 23 26 14 13 9 0
3	66.24	190	0 28 0
	400.27		

Table 4.15: Routing scheme prior to transfer.

Upon the creation of a coordination site at the node closest to both the departure point for the recipient vehicle and the recipient node, the vehicle servicing the node co-located at the coordination site (the donor vehicle) is able to deposit at least part of the demand for the recipient node at the coordination site. Having done this, the donor vehicle proceeds without further involvement in this coordination activity, although it is free to participate in any other coordination plan. The only constraint on this donor vehicle is that it deliver the goods to the coordination site at such a time that the recipient vehicle is able to process them without violating any timing constraint (for example, the recipient vehicle might have to be at another location, such as the depot, no later than a particular time).

Having arrived at the coordination site, the recipient vehicle either must wait for the donor vehicle to transfer the goods, or may collect the goods and depart for the recipient node, perhaps as part of another multi-noded route.

After delivery of the demand to the recipient node, the vehicle is able to participate in further coordination.

Where the requirement for the recipient vehicle to begin and end its journey at the depot is relaxed, the vehicle becomes a rover, able to roam the Euclidean plane collecting payloads from coordination sites and delivering them. This is a region of exploration suitable for further investigation.

The coordination scheme is shown in Table 4.16. Establishment of a coordination site at node 8 means that node 9 is no longer part of the green route. The red vehicle services node 28 thus freeing up carrying capacity. It travels to the coordination site, collects the remaining part of demand(9), combines it with the portion already carried, and delivers it to node 9. The splitting of node 9 was not a requirement for transfer to occur; if the total demand of the nodes serviced by the blue vehicle had been lower, demand(9) could have been carried intact. The total distance reduction is 2.7% overall, and 5.0% in the altered routes.






Figure 4.25: After servicing node 28, the red vehicle is able to continue to a coordination site at node 8, collect demand(9) and deliver it before returning to the depot.

Vehicle	Distance	Load towards node 8	Load away from node 8	Route
1	186.27	197	172	0 8 7 6 5 4 3 2 1 24 25 22 21 0
2	113.52	-	-	0 16 14 13 26 23 15 27 0
3	89.80	10	35	0 28 8 9 0
389.59				

Table 4.16: A coordination site is established at node 8.

In the example described here, there is a distance reduction created by the removal of the isolated node from an otherwise clustered route. The isolated node is serviced by a vehicle that, without coordination, could not have accommodated the demand but had distance capacity that enabled it to reach the node via a coordination site. The demand was delivered to the coordination site by another vehicle that had enough carrying capacity to accommodate the demand but had not the distance capability to reach the node.

4.13 Obviating return to depot

As the result of a routing schedule, one or more vehicles may service more than one route each. The length of one route is determined by

- the carrying capacity of the vehicle,
- the sum of the demands of the nodes allocated to it, and
- limitations on the length of a route or on the time of a route.

A vehicle is also limited by the maximum time it can be on the road or by the maximum distance it can travel without a break. It is sometimes possible for a particular vehicle to service more than one route without violating distance and timing restrictions. In a pure delivery scenario, the distance involved in returning to the depot to collect that second (or subsequent) load is normally unproductive time, since no customer is being serviced, but occasionally is used to delay the departure of a subsequent route in order to satisfy feasibility requirements subject to timing constraints. Coordination points can be used in an attempt to eliminate or to delay that inter-route return to the depot.

When coordination sites are created in order to reduce the need for multi-routed vehicles to return to the depot, split deliveries can be incorporated but are not essential. The load for at least part of the later route is accommodated on a second vehicle, which then drops these part-loads at a location that can be reached by the vehicle travelling the earlier route. There may be a need to return to the depot before the end of that second route thus deferring rather than precluding the return to the depot.

Both single-noded and multi-noded routes are receptive to improvement by obviation of return to depot.

4.13.1 Single-noded routes

In the case of single-noded routes, the demand for the solitary node may be shared among other routes, subject to the constraints of distance, capacity and load splitting.

Consider the situation shown in Figure 4.26. Colours represent vehicles; nodes are labelled a to i ; demands are shown beside each node; the vehicle capacity is 200 units. The red vehicle services both a multi-noded and a single-noded route. The arrowed node of the single-noded route will have its demand delivered intact but carried split.




Figure 4.26: Situation before coordination. The multi-routed red vehicle services a single-noded route.

Without the application of coordination, the routes would be those in Table 4.17 with the red vehicle servicing routes 1 and 2. None of the vehicles is filled to capacity at any time. Since load splitting is permitted, the entire load for route 1 can be divided among the other routes. All combinations of route order and orientation need to be examined when considering the creation of coordination sites. Whenever the single-noded route follows the multi-noded route in a vehicle’s route order, the multi-noded route can actively participate in coordination rather than merely receiving the benefit of it. In this case, if the red vehicle’s route order is 2 1 rather than 1 2, three vehicles can share the demand of the single-noded route among them compared with two vehicles otherwise.

Assuming a route order of 2 1, after the completion of route 2 the red vehicle would have travelled back to the depot, collected the load for node *a*, travelled to node *a* and returned to the depot.

Using load splitting and two coordination points (co-located with nodes *c* and *i*) the inter-route return to the depot is removed. The new routes are shown in Figure 4.27 where arrows indicate the direction of travel.

Route number	Load	Route
1	200	0 b c g i a 0
2	200	0 c d e f 0
3	186	0 h i 0

Table 4.18: After coordination. Coordination sites have been created at nodes *c* and *i*.

■

The multi-routed vehicle is not constrained to service a multi-noded route, as the following example demonstrates. Had the red vehicle in the previous example serviced the same three nodes, but each with large demands that required them to be on single-noded routes, it would have been multi-routed with three single-noded routes. The routing scheme is shown in Figure 4.28 and Table 4.19.




Figure 4.28: The multi-routed red vehicle services three single-noded routes.

Vehicle	Route
red	0 a 0 b 0 g 0
green	0 c d e f 0
purple	0 h i 0

Table 4.19: Multi-routed vehicle servicing only single-noded routes.

Route order for the multi-routed vehicle must be considered when investigating the potential of coordination. In this case there are six possible route orders for the red vehicle, one of which results in $0 \ b \ 0 \ g \ 0 \ a \ 0$, corresponding to the order in which those nodes were serviced in the example above. The following coordination plan enables the red vehicle to deliver demand(a) without returning to the depot after servicing node g :

All vehicles leave the depot with full loads. After servicing node b , the red vehicle continues to node c where, the green vehicle having waited, 18 units are transferred from the red vehicle to the green one. The red vehicle returns to the depot, collects demand(g) and travels to g where it delivers demand(g) and meets the green vehicle again. This time, the transfer is from the green vehicle to the red. The green vehicle returns to the depot while the red vehicle continues on to node h where it collects the balance of demand(a) before travelling to node a and delivering its demand. Finally, the red vehicle returns to the depot. The total amount that can be delivered to node a using this scheme is

$$4 \times \maxVehCap - \sum_{x=b}^i demand(x)$$

which has a value of 132 for this example.




Figure 4.29: One inter-route return to depot is removed for a vehicle servicing multiple single-noded routes. Arrows indicate direction of travel.




Figure 4.30: Removal of inter-route return to depot between two multi-noded routes. The structure on the left is re-designed by the creation of a coordination site at node 9.

Fertile structural characteristics for the introduction of coordination in an attempt to reduce distance travelled and/or fleet size by obviating the inter-route return to depot for multi-noded routes include:

- a vehicle (other than the multi-routed subject) with spare carrying capacity
- coordination site on the earlier route accessible after delivery of suitable payload
- appropriate combination of broken and made links
- total demand able to be carried by the existing vehicles.

The roading company of Problem 1 (see Chapter 1 and Figure 3.1) utilises transfer of bitumen from one vehicle to others at a coordination site in order to remove the requirement for each vehicle to return to the distance supply yard for replenishment of stock.

4.14 Coordination site as interim step

Coordination sites can be created which may violate load and distance constraints. These metastable coordination sites are created in an intersection-based approach and eliminated in later steps to produce routing structures with reduced fleet size or reduced total distance. This technique can be used to investigate the removal of intersections from a feasible solution to create another feasible solution.

The potential coordination site is created by the intersection of two routes, as shown in Figure 4.31. At the site, four part-routes impinge: the before- and after-encounter for each of the two routes.




Figure 4.31: Investigation of the formation of an interim step. An intersection close to a node (a) is considered as though it occurs at the node (b). Four part-routes are created, shown as coloured lines.

In order for the coordination site to be co-located with a node, the intersection is assumed to occur at a node close to it, thus creating the interim step. All nodes within the near radius of the intersection are considered as candidates for coordination site. The sensitivity of the technique to the size of the near radius is discussed later.

Five nodes make up the interim coordination site system: the coordination site node, the node immediately preceding it and the node immediately following it on one route, and the two nodes on the intersecting route that define the intersecting edge. When the coordination site is an interim step, those four part-routes rearrange in a dissolution process so that the coordination site is in only one

of the resulting routes, not in both. This differs from the use of a coordination site to swap loads, where the coordination site appears in both routes.






Figure 4.32: Coordination as an interim step. An intersection in a feasible solution has been removed to produce another feasible solution.


When the interim step splits apart, one of five possible arrangements results, shown in Figure 4.33 (in which the red node represents the potential coordination point). Some results seem at first glance to be subsets of others but each has been chosen because of the manner in which the "loose" nodes are re-assigned.




(i) 3,1,1 system




(ii) 2,1,1,1 system



(iii) 3,2 system



(iv) 4,1 system



(v) Depot 4,1 system

Figure 4.33: Structures resulting from dissolution of the interim step.

The structures are not unique; for example, there are two type (iii) structures, shown in Figure 4.34, and four type (iv) structures.



Figure 4.34: The two type (iii) dissolution structures. One will be the original; in the other, the node co-located with the coordination site has been transferred to the other route that forms the intersection.

The links shown in the dissolution of the interim step are regarded as permanent in that they can not be modified in the subsequent improvement procedure which attempts to reassign all other links. The type (v) dissolution in Figure 4.33 only exists when the depot is included in the original set of five nodes. The depot may appear more than once in the original set since it occurs twice in every route but is not considered to act as a coordination point since this is a trivial instance.

After a candidate node has been identified, the vehicle is located in an *a posteriori* approach. As in other forms of coordination, the creation and dissolution of an interim step may result in an environment which also is receptive to improvement by coordination. Such cascading of coordination potentials should be taken into account when determining the best options to adopt.

The potential use of coordination in the creation and dissolution of an interim step is indicated in routing structures by

- initial feasible solution
- node within the near radius of an intersection
- resulting feasible solution.

Consider the route structure shown in Figure 4.35 in which ten nodes are served by three routes which intersect in three places; nodes are labelled *a* to *j*; colours represent vehicles. The near radius is indicated by the dotted circles, each

of which is centred on an intersection. Four nodes lie within the near radii and each must be considered as candidates for coordination sites.




Figure 4.35: Before coordination. Circles represent extent of the near radius.

Route #	Distance	Load	Route
1	152.269	196	0 a h f b 0
2	195.548	162	0 d c e 0
3	<u>181.166</u>	186	0 j g i 0
	528.982		

Table 4.20: Route data for the scheme shown in Figure 4.35.

The proximity of each intersection with a node raises the possibility of the use of coordination as an interim step to reduced the total distance travelled and/or fleet size. For each of the three intersections, possible coordination sites are identified as in Table 4.21.

Intersection	Possible coordination site	Nodes involved
1	e	0, a, e, c
2	b	0, b, e, c
3	d	0, d, i, g
3	i	0, d, i, g

Table 4.21: Interim step investigation.

Investigation of the first intersection will be shown in detail to demonstrate the use of the interim step. Although five nodes are required to form the metastable state, these nodes need not be unique if the depot is involved since the depot appears twice in every route. The metastable state for intersection 1 is shown in Figure 4.36.




Figure 4.36: The metastable state for intersection 1 in Table 4.21.

Dissolution type	Fixed links
(i)	<i>0-e-a 0-e-c 0-e-0 a-e-c</i>
(ii)	<i>e-0 e-c e-a</i>
(iii)	<i>{0-e-a, 0-c} {0-a, 0-e-a}</i>
(iv)	<i>0-e-a-0 0-e-c-0 0-e-c-a</i>
(v)	<i>0-c-a-0</i>

Table 4.22: Fixed links from the dissolution of the metastable state shown in Figure 4.36.

Several types of dissolution for intersection 1 yield the same result, shown in Table 4.23. The minimum total distance represents a reduction of 13.5% overall and 18.9% in the affected routes (since route 1 is unchanged). No intersections are present in the final routing which shows a petal structure. Although the loads are

not as evenly distributed, the route distances are more equitable. Similar investigations of intersections 2 and 3 produce the same final result.

The size of the near radius is set to embargo calculations involving links that clearly will not result in an improved solution. In Figure 4.35, it can be seen that the near radius excludes node c , which is eventually involved in a changed link. This is not a problem because there is a high redundancy in the calculations and because nodes that (in the initial feasible solution) are routed with the five nodes involved in forming the metastable state, are also involved in forming interim steps when their linkages have been broken. Thus, node c , which is linked with node d , is investigated along with node d .

Type	Routes	Distance	Comments
(i)	0 e a 0	144.38	Also from:
	0 h f b 0	91.47	type(ii),
	0 c d i 0	155.26	both type (iii),
	0 g j 0	<u>86.78</u>	type(iv).
		477.89	
(i)	0 e c d 0	528.98	Original routes
(i)	0 e c 0	174.85	
	0 a h f b 0	152.28	Also from:
	0 i d 0	139.46	type (ii),
	0 g j 0	<u>86.78</u>	type (iv)
		553.37	
(i)	0 e 0	67.08	
	0 a h f b 0	152.28	
	0 c d i 0	155.26	
	0 g j 0	<u>86.78</u>	
		461.40	
(ii)	0 e j g 0	150.19	
	0 c d i 0	155.26	
	0 a h f b 0	<u>152.28</u>	
		457.73	

continued on next page ...

... from previous page			
(iv)	0 e c a 0	237.10	> maxVehDist
	0 c a 0	196.07	
	0 e 0	67.08	
	0 h f b 0	91.47	
(v)	0 d i 0	139.46	
	0 g j 0	<u>86.78</u>	
		580.86	

Table 4.23: Routes resulting from the dissolution of intersection 1 in Figure 4.35. The minimum total distance of 457.73 represents a reduction of 13.5% overall.

The initial feasible solution, created using a modified parallel savings heuristic, could not have produced the feasible solution generated here by using improvement by coordination as an interim step. The final solution has the appearance of a structure generated using a sweep algorithm.

Route #	Distance	Load	Route
1	152.28	196	0 a h f b 0
2	155.26	200	0 c d i 0
3	<u>150.19</u>	148	0 e j g 0
	457.73		

Vehicle	Distance	Route(s)
green	152	1
red	155	2
blue	150	3

Table 4.24: Using coordination point as interim step.




Figure 4.37: After the application of interim step.

■

Every time an interim step is used, the links connecting the nodes may have changed, possibly creating new intersections, each of which needs to be investigated for possible coordination.

The final route structure may be the result of the successive application of many interim steps, some of which have been induced by earlier coordination steps. The next section demonstrates the effect of these cascading interim steps.

4.14.1 Cascading interim steps

The product of an application of coordination as an interim step may itself contain intersections close enough to nodes to warrant further investigation. Consider the routing scheme shown in Figure 4.38 and described in Table 4.25. The initial feasible solution contains eight intersections, each of which will be investigated for improvement by coordination as an interim step. In this example, to highlight the effect of the interim step, split deliveries will not be used. The near radius is set at such a level as to preclude investigation of combinations that could not be expected to produce favourable results.




Figure 4.38: During the investigation of coordination as an interim step, eight intersections will be investigated.

Route #	Distance	Load	Route
1	153.199	187	0 22 19 17 0
2	95.650	142	0 5 1 0
3	86.914	178	0 18 25 7 0
4	71.573	173	0 23 8 0
5	129.514	199	0 20 12 0
6	109.875	192	0 6 15 9 14 11 0
7	104.941	172	0 21 13 0
8	35.777	125	0 2 0
9	79.398	159	0 3 0
10	12.166	144	0 4 0
11	106.019	129	0 10 0
12	114.630	160	0 16 0
13	128.250	200	0 24 0
1227.904			

Vehicle	Distance	Route(s)
1	153	1
2	201	4 5
3	220	7 12
4	220	9 10 13
5	216	6 11
6	218	2 3 8

Table 4.25: Route structure before coordination for the routes in Figure 4.38. maxVehCap=200, maxVehDist=220, near radius=10.

Because of the high values of some demands, six of the thirteen routes are single-noded and three of these single-noded routes are involved in forming intersections. If the demands of potential coordination sites are high enough that no other node may be serviced by the same vehicle without violating loading constraints, the corresponding intersection should be removed from consideration. Also removed from consideration are those intersections which do not have a nearby node other than those defining the intersection (except where the depot is involved). Having excluded any possible coordination sites that rely on too few nodes or on too heavy demands, it can be seen from Table 4.26 that this scheme

has eight possible candidates for coordination site, although only one involves five unique nodes.

Intersection 6 is a classic demonstration that the near node does not have to be involved with either of the routes forming the intersection.

Inter-section	Links involved	Near node	Demand of near node	Prohibition	Five nodes?	OK?
1	20-12, 0-24	-	n/a	n/a	yes	yes
2	20-12, 0-10	20	80	n/a	yes	yes
3	20-12, 0-16	-	n/a	n/a	yes	yes
4	20-12, 0-13	12	119	n/a	yes	yes
4	20-12, 0-13	13	120	n/a	yes	yes
5	20-12, 21-13	12	119	n/a	no	no
5	20-12, 21-13	13	120	n/a	no	no
6	0-12, 13-21	23	92	12, 13	yes	yes
7	0-23, 13-21	2	125	23, 13	yes	yes
8	0-8, 13-21	-	n/a	n/a	yes	yes

Table 4.26: Options. "Prohibition" refers to nodes in column 2 that cannot be trucked with the node at the potential coordination site without violating capacity constraints. The depot (node 0) can be involved more than once.

While most of the options in Table 4.26 yield no benefit and create no further intersections, several combinations are fruitful:

- (a) 20-12, 0-13, 13
- (b) 0-12, 13-21, 23
- (c) 0-23, 13-21, 12.

The combination (a) of the intersection defined by links 20-12 and 0-13 with near node 13 provides only a small improvement in the total distance travelled of about 6 units (~0.5%) which might be considered minor except that further intersections are formed which might have greater benefit. The cascading interim step is potentially worthwhile and must be investigated before the full effect of accepting the first interim step can be determined. The results of accepting the coordination site at node 13 as an interim step are shown in Figure 4.39. Only three intersections survive the re-design by the first interim step, and two new ones are

created. The new intersections are investigated in the same manner as the first intersections (see Table 4.26).




Figure 4.39: After a first application of coordination as an interim step in combination (a), three intersections have been eliminated and two have been changed.

Because the same five nodes are involved in both secondary investigations, some of the calculations will be duplicated. When consideration is given to the capacity constraints, there are only two possible results for the combinations of Table 4.26, shown in Figure 4.40.

Inter-section	Links involved	Near node	Demand of near node	Prohibition	Five nodes?	OK?
1	12-21, 0-23	8	81	n/a	yes	yes
2	12-21, 0-8	23	92	12	yes	yes

Table 4.27: Interim step investigation of the results of combination (a).




Figure 4.40: A second interim step formation. The structure on the left is the result of the first interim step. All other options were disadvantageous.

In Figure 4.40, the original structure is returned by a type (iii) dissolution using node 8 as the coordination site, and by a type (ii) dissolution using node 23 as the coordination site. The successful structure on the right of Figure 4.40 is returned both by a type (ii) dissolution using node 8 as the coordination site and by a type (iii) dissolution using node 23 as the coordination site. All other dissolutions were rejected either because they violated capacity constraints or because they produced longer travel distances than the original structure. Since no new intersections were created, the investigation chain for combination (a) ends.

The final results of accepting a coordination site at node 13 with the intersection defined by links 20-12 and 0-13, i.e. combination (a), are shown in Figure 4.41. The investigation route was not unique, since the same result was found by more than one possibility.

When investigating combination (b) (0-12,13-21,23), two type (ii) dissolutions each yield the routes 0-21-23-0, 0-8-12-0, and 0-13-20-0 shown in Figure 4.41 with a total distance of 1177.8. One new intersection has been created, defined by 21-23 and 0-8, but no application of an interim step to this provides any

improvement. No other dissolutions of combination (b) provided any benefit or generated intersections that provided any benefit.

The combination (c) of the intersection defined by links 0-23 and 13-21 with near node 12 provides no improvement, resulting, via a type (i) dissolution, in an increase in total distance and the routes 0-2-21-0, 0-8-23-0, 0-13-20-0 and 0-12-0. It also creates four new intersections which are investigated in the same manner as those above, resulting in the same routing structure found by combinations (a) and (b). The result of investigating combination (c) is the structure shown in Figure 4.41.




Figure 4.41: The result of cascading investigations of using coordination sites as interim nodes.

When the results of combinations (a) – (c) are compared, they are found to have identified the same feasible solution, shown in Figure 4.41. The reduction in total distance is ~1% overall and 2.5% in the affected routes.



4.14.2 Stability of intersections

The node that is close enough to an intersection to become a coordination point may still be close to an intersection even after coordination. Indeed, the intersection whose existence initiated the use of coordination is not always destroyed in the improved routing scheme. The following example demonstrates this.

Consider the routing structure shown in Figure 4.42. The dissolution is type (i) in Figure 4.33, resulting in a 7% reduction in distance. The original intersection disappears and is replaced by another in almost the same location involving the same near node. Even though the intersection is defined by only three of the four nodes that defined the original intersection, it is co-located with the original and has the same near nodes.

■




Figure 4.42: Stability of intersections. An intersection (indicated by the red arrow) involving the node at the coordination site still exists in improved routes.

Consider the routing structure shown in Figure 4.43. The dissolution is type (iii) and the resulting routes are improved by a 2-opt procedure to become those shown. There is a 7% improvement in distance although the number of intersections is unchanged.

■




Figure 4.43: The number of intersections is unchanged.

4.14.3 Limitations and advantages of coordination as an interim step

Coordination can be used in improving feasible solutions by forming a metastable state in an interim step and dissolving that structure in a variety of ways to produce a feasible solution by re-assignment of the route fragments. If the produced solution is more advantageous than the original, the new solution is accepted as the improved solution. The original solution can always be formed as one of the possible results of reassigning the dissolved structure.

The use of interim steps to improve a feasible solution produces other feasible solutions that may be found by a combination of other methods, including inter-route node exchange. It is of use primarily in the improvement of poor initial solutions that include intersections. Therefore, construction heuristics that customarily inhibit the formation of intersections, such as the sweep method, are not receptive to this particular use of coordination as an improvement technique.

The cascading effect of interim solutions must be considered before accepting or rejecting alternatives.

Although the technique attempts to reduce total distance travelled by reducing the number of intersections, the final solution might contain the same number of intersections as the original. In addition, an intersection that is almost co-located with the original intersection may be present in the final feasible solution.

Although the theory requires five nodes to form the metastable state, these nodes need not be unique. The depot may be considered more than once since it is obliged to appear twice in every route. An exception to this is the "roaming vehicle" scenario, mentioned briefly above, in which a vehicle moves from node to node without necessarily being domiciled at the depot.

Nodes other than the five involved in forming the metastable state are involved in creating a feasible result, since "loose" nodes may be combined into other routes. In addition, nodes that are in the same routes as the five nodes forming the metastable state are also involved in the investigation since their linkages are broken and rearranged. As a result, investigations involving adjacent nodes (that are being considered as coordination sites) produce many of the same results from different dissolution types.

Because of the variety of dissolution types, more than one investigation path may lead to the best result.

This use of coordination can not easily be used in conjunction with other types of coordination that require the presence of intersections, since it seeks to

reduce the influence and existence of intersections. Since load swapping and load transfer both depend on the presence of intersections, interim step formation is not recommended where payload exchange is preferred.

4.15 Appraisal of Euclidean coordination

This study has identified several types of coordination as effective methods for improving feasible solutions to the capacitated VRP in the Euclidean plane. They include

- split deliveries
- load swapping
- load transfer
- obviation or deferral of inter-route return to depot
- metastable state creation as an interim step.

Because each form of coordination occurs under different conditions, it is important when applying the improvement to use the form most suitable for the prevailing climate. Empirical evidence suggests that the order in which the types of coordination have been identified in the above list yields the greatest improvement.

The use of split deliveries [DT89] as an improvement method is well established although it has not been discussed in the literature explicitly in the context of coordination. Where split deliveries are permitted, they are used effectively in conjunction with other coordination methods as integral to those other methods rather than as separate steps. As an example, in situations where load splitting is permitted, the splitting is most effectively practised at the swapping location, so that loads that are carried in split mode from the depot to the coordination site are re-combined at the coordination site by being swapped with other loads.

Load swapping allows initial feasible solutions to be improved by reducing both distance and fleet size. Because swapping is a two-way process, the vehicles involved in the exchange must be present at the coordination site at the same time, a requirement which may require one of the participating vehicles to wait for the other. This interdependence of temporal and spatial characteristics means that changes to the orientation and route order may significantly alter the potential for acceptable coordination. Distance reduction is more effective where coordination occurs early in the route, rather than later, because a greater part of the route is affected by the exchange; nodes serviced before the swapping occurs are not

involved in improvement by swapping. If the original routing structure allows at least one of the vehicles involved in coordination to be near-empty when the load swapping occurs, there is more freedom in re-assignment of payloads. Distance reduction is improved where nodes that are close together are serviced by the same vehicle. A reduction in total distance does not necessarily imply a reduction in the number of routes.

A reduction in fleet size without a concomitant increase in distance can be achieved through load swapping by improving the distribution of loads over the fleet. This involves routes that are not active participants in the exchange process, since the vehicles that were initially assigned to them are re-assigned with the reduction in fleet size. Changes in vehicle assignment may result in alterations to route order and therefore the time at which nodes are serviced.

A structural characteristic common to many scenarios that prove receptive to improvement by load swapping is the slight overlap of routes in the original non-coordinated scheme. Exchange of loads and the subsequent improvement by node exchange and k-opt reduces the overlap and the total distance travelled.

Load transfer is a special case of load swapping since it may be thought of as a one-way exchange of payload from a donor vehicle to a recipient vehicle. It differs from two-way exchange by relaxing the constraint involving the concurrent presence of donor and recipient vehicle at the coordination site. For load transfer to occur, it is necessary that the recipient vehicle not leave the site before the arrival of (and transfer by) the donor vehicle.

Load transfer is particularly suitable as an improvement method by reducing total distance by removing an isolated node from an otherwise clustered route in an environment that includes

- isolated node in an otherwise clustered route
- vehicle with spare carrying capacity but no spare distance capacity
- vehicle with no spare carrying capacity but spare distance capacity
- the above occurring within some "suitable" distance of each other.

Under coordination by load transfer, the isolated node is serviced by a recipient vehicle that, without coordination, could not have accommodated the demand but had distance capacity that enabled it to reach the node via a coordination site. The demand is delivered to the coordination site by a donor vehicle that has enough free payload to accommodate the demand but does not have the distance capability to reach its intended node. Under these circumstances, the coordination site is created at the node closest to both the isolated node and

the node preceding it, subject to feasibility and constrained by the timing requirement that the donor vehicle not leave before the transfer is completed. The vehicle servicing the node at the coordination site is the donor vehicle, and having deposited the demand (or part of it) of the isolated node its involvement in this coordination activity is complete; it is then free to participate in any other coordination that may be necessary. The recipient vehicle is similarly constrained not to leave the coordination site until the transfer is complete; it, too, is then free to participate in any other coordination that may be necessary, subject to feasibility.

A further relaxation to the problem occurs when a vehicle is no longer constrained to return to the depot at the end of a route but is able to roam the Euclidean plane engaging in coordination. This is a region of exploration suitable for further investigation. The potential for coordination may be enhanced by removing the requirement for vehicles to return to the depot because they are able to travel further from the depot and increase the area of operations without violating distance constraints.

Where multi-routed vehicles are employed, there is the possibility of reducing the total distance travelled by removing or deferring the inter-route return to depot. As in load swapping, split deliveries may be incorporated. For the multi-routed vehicle, the load for at least part of the later route is accommodated on a second vehicle, which then drops these part-loads at a location that can be reached by the vehicle travelling the earlier route. If only part of the payload of the later route can be accommodated while servicing the nodes on the earlier route, the inter-route return to depot is postponed to a more convenient time rather than removed entirely. The requirement that the vehicle be multi-routed does not limit the number of nodes serviced by the routes beyond the normal capacity and distance constraints; both single-noded and multi-noded routes are receptive to improvement by obviation of return to depot. Where single-noded routes are involved, the demand of the single node is shared among the other routes. In the case of multi-noded routes, the routes are redesigned rather than re-distributed. Such cases require an environment that includes

- a vehicle (other than the multi-routed subject) with spare carrying capacity
- coordination site on the earlier route accessible after delivery of suitable payload
- appropriate combination of broken and made links
- total demand able to be carried by the existing vehicles.

In contrast to the other forms of coordination which are employed subject to the constraints of capacity and distance (among others), the establishment of a metastable state as an interim step is unconstrained by such requirements. However, the subsequent and immediate dissolution of the metastable state and formation of a feasible situation is as constrained as the non-coordinated system. Creation of the metastable state requires only that there be a node and an intersection sharing the same neighbourhood. The ensuing formation of a feasible solution offers as one of the alternatives the re-establishment of the initial feasible solution. Because the formation of the resultant feasible solution involves the insertion or re-assignment of only a few nodes, the method used to create the initial solution has no effect on the final result beyond creating the initial environment. Consequently, situations in which intersections abound close to nodes may be fertile climates for using metastable state coordination; similarly, systems which present few intersections and which have few nodes near intersections, are not suitable for distance reduction by this form of coordination.

The effect of coordination as a method for reducing the total distance travelled and/or the fleet size of initial feasible solutions to the VRP in the Euclidean plane is determined in part by the environment created by the initial solution and in part by the type of coordination employed, in addition to the usual constraints of distance, capacity, location and problem structure. The application of coordination changes the environment and affects the climate for subsequent applications of coordination.

Generally, where the environment permits, the most effective order of application of coordination in the deterministic Euclidean scenario is

1. load swapping
2. load transfer
3. removal of inter-route return to depot
4. interim step formation.

If load splitting is permitted and attracts no penalty, it should be employed as an integral part of the other forms of coordination, allowing part-loads to be swapped or transferred. Other than the creation of the metastable state in interim step formation, all coordination is carried out subject to the constraints of capacity, distance and location. Coordination is accepted only if it enables a feasible solution to be formed that has either a lower total distance or a smaller fleet than the non-coordinated form. This may be a limitation since the application of one type of

coordination may result in an infeasible solution that nevertheless provides a suitable environment for another type of coordination whose subsequent application can result in an overall reduction. This has not been investigated in this study and is an area that deserves further research.

This study has investigated coordination only at sites co-located with existing nodes in a purely deterministic setting, concentrating only on the beginning and end of the part-routes involved in coordination. The next chapter deals with coordination in the rectilinear grid, in which the actual paths available to vehicles assume a greater importance.

CHAPTER 5

COORDINATION IN THE RECTILINEAR GRID NETWORK

The aim of this chapter is to investigate the effectiveness of coordination in a non-Euclidean setting. To achieve this, there is an initial brief consideration of the differences between Euclidean and rectilinear scenarios and the reasons for studying coordination in the rectilinear grid. Then the grid itself is investigated and advantage taken of its rectilinear nature to store information about it and to provide a means for performing computations on the grid. In addition, grid reduction is described as a method of removing some non-essential edges to greatly simplify computation. Having defined the grid, route construction and route scheduling methods are considered. Once routes have been created, scheduled and improved, they may be suitable for further improvement by coordination. As user-defined values which influence the potential for coordination, the maximum distance constraint and maximum vehicle capacity constraint are discussed. Having considered the underlying grid and the methods used to create and improve routes, the types of coordination are investigated in detail in isolation – load swapping, deferral of return to depot, and split delivery. Techniques for selecting suitable coordination sites are discussed and alternative coordination sites are evaluated. The influence of the physical environment on the potential for coordination is investigated as the effect of grid density. This includes a discussion on the effect of depot location on the potential for coordination. The order of application of the various types of coordination is then discussed. Finally deterministic rectilinear coordination is evaluated. The following chapter considers dynamic coordination in the rectilinear grid.

5.1 Introduction

A rectilinear grid network has been chosen because it closely simulates the road layout in a city and because the potential paths are clearly circumscribed. Unlike

the Euclidean plane, in which paths are defined by the nodes through which they pass, the rectilinear grid forces paths to be defined by their component edges.

Where in the Euclidean plane the shortest path between two nodes is the unique straight line joining the nodes, in the rectilinear grid there may be many shortest paths. Figure 5.1 shows four paths of shortest length between two nodes. The multiplicity of shortest paths increases the number of interactions that potentially can occur between routes.




Figure 5.1: Paths of equal length between nodes.

Unlike the Euclidean scenario where routes intersect only at points, neighbouring paths in the rectilinear grid may share edges. In Figure 5.2, the pink route shares two edges with the red and black paths, one edge with the blue path and no edges with the green path. Although the red, black, blue and green paths represent the shortest distance between nodes A and B, their interactions with the route joining nodes C and B are quite different. If a coordination site were to be established between the two routes, its placement would depend to some extent on which path route AB took.




Figure 5.2: Different routes sharing common edges.

5.2 Grid construction and definition

In this study, rectilinear grid networks [B93] are generated by forming a complete Manhattan graph of size 21 by 21 and then deleting a varying number of edges to produce three levels of grid density: sparse, medium and dense, examples of which are given in Appendix A. This grid-size has been chosen because it allows resultant routes to contain a reasonable number of nodes and to permit enough potential interactions to investigate coordination. Three levels of grid density have been used in order to vary the structural characteristics of the grid that can influence the

behaviour of vehicles, such as cul-de-sacs and forbidden regions (containing no edges).

All edges are two-way, enabling vehicles to traverse routes in both directions. In practice, routes rarely can be travelled with the same ease in each direction because of the influence of terrain, one-way streets, and difficult intersections. In addition, traffic regulations (e.g. [NZRC]) usually treat left- and right-turning traffic differently, regardless of the side of the road on which it travels. While recognising that exclusively two-way systems are more limited in their practical applications than mixed two-way/one-way/restricted-turning systems, it is expected that they are able to provide a clear indication of the capabilities of coordination and that the results obtained from their study can be applied to other systems.

In this study, each grid is defined by two 0/1 matrices: EH representing the horizontal edges and EV representing the vertical edges.

$$EH(i,j) = \begin{cases} 1 & \text{if there is an edge to the right of vertex } (j,i) \\ 0 & \text{otherwise} \end{cases}$$

and

$$EV(i,j) = \begin{cases} 1 & \text{if there is an edge above vertex } (i,j) \\ 0 & \text{otherwise} \end{cases}$$

All changes to a particular grid are accompanied by minor changes to the contents of EH and EV without affecting the size of the matrices. The upper limit on the size of problem that can be tackled using this method of grid definition depends on the storage capacity of the computing device. For a 21×21 grid, EH has dimensions 20×21 and EV has dimensions 21×20.

5.3 Node location on the grid

Nodes are located randomly at the vertices of the grid, rather than along edges. This simplifies the computations without loss of generality since a node that lies on an edge can be considered to lie a fraction of the length of the edge from either of the two vertices defining that edge. Edge location of nodes would introduce another set of variables into the computations made up of one of the end vertices defining the edge and the fractional displacement in a particular direction from the vertex, and may yield improved solutions particularly where the ratio of distance

travelled: distance between vertices is small. The underlying concepts do not change since vertex location can be considered as a special case of edge location in which the displacement along the edge is always zero.

Restricting node locations to vertices has several consequences:

1. paths consist of complete edges
2. point intersections occur at vertices
3. common sections of paths are multiples of complete edges
4. routes have integer-valued lengths.

As a result of consequences 2 and 3, coordination sites involving intersections occur at vertices. As a result of consequences 1 and 2, coordination sites involving detours also occur at vertices. Accordingly, all coordination sites occur at the vertices of the rectilinear grid.

5.4 Grid network distances

A path from point $A=(x_A, y_A)$ to point $B=(x_B, y_B)$ is made up of points (x, y) defined by a continuous vector-valued one-to-one function $\alpha=(u, v)$ on the interval $[0, 1]$ with

$$\begin{aligned}\alpha(0) &= (u(0), v(0)) = (x_A, y_A) \\ \alpha(t) &= (u(t), v(t)) = (x, y), \quad t \in (0, 1) \\ \alpha(1) &= (u(1), v(1)) = (x_B, y_B).\end{aligned}$$

Let $t_1, t_2, \dots, t_{n-1}$ be points in $(0, 1)$ with $t_{i-1} < t_i$ for $i = 2, \dots, n-1$. The points $(x_i, y_i) = \alpha(t_i)$, $i = 1, \dots, n-1$ are intermediate points of a path $P(A, B)$ with length

$$\text{len}(P(A, B)) = \lim_{n \rightarrow \infty} \sum_{i=1}^n (u(t_i) - u(t_{i-1}) + v(t_i) - v(t_{i-1})).$$

A rectilinear path, $P(A, B)$, from A to B is defined (e.g. in [DS01]) in terms of a finite set $t_i \in [0, 1]$ for $i = 1, \dots, n-1$, with $t_{i-1} < t_i$, such that on every other interval (t_{i-1}, t_i) , $u(t)$ is constant and $v(t)$ is either increasing or decreasing, and on alternate intervals, (t_i, t_{i+1}) , $v(t)$ is constant and $u(t)$ is either increasing or decreasing. A staircase path is defined as a nondecreasing or nonincreasing rectilinear path, and is a shortest path from A to B . A rectilinear path proceeds from one point to another in a sequence of connected straight line segments alternating between vertical and horizontal; in the staircase path, all horizontal segments have the same

direction and all vertical segments have the same direction. Consequently, all staircase paths between the same two points have the same length, but this is not necessarily the case with rectilinear paths.

In this study, the shortest distance between two adjacent vertices on the grid is taken to be 1 unit. Since the length of a rectilinear path from A to B is equivalent to the sum of the lengths of the horizontal and vertical segments of the path, this is equivalent to 1 + the number of vertices between, but excluding, A and B . Other researchers [e.g. F93] have used larger and varying values but this work differs in that the nodes are located only at vertices and not part way along edges. Maintaining a constant distance of one unit between adjacent vertices simplifies calculations without loss of generality. The 21 by 21 grid has dimensions 20×20 .




Figure 5.3: Rectilinear and staircase paths.

Where barriers exist, a staircase path between two points may not exist and a shortest path may be represented not by a staircase path but by a more circuitous rectilinear path. In this study, barriers are caused by the physical layout of the grid (non-existent edges) rather than by constraints on traffic across edges. Each edge is assumed to have infinite carrying capacity, enabling any number of vehicles to traverse it at any one time in either direction. In practice, all roads have a maximum traffic density capability which, when exceeded, results in gridlock. For the relatively small problems investigated here with no more than 15 vehicles (a limit again imposed by the MATLAB programming language), the assumption of infinite carrying capacity of edges is plausible.

It will be shown in §5.6 and subsequent sections that coordination in the rectilinear grid may require the involvement of every intersection in the region known as the coordination hull, and in §5.15 that each vertex has the potential to be used as a coordination site. Accordingly, to accommodate these requirements, procedures are developed that identify potentially useful intersections and vertices and ignore intersections that will not be used in creating useful solutions. This

reduces the computations needed to identify shortest paths and stores those paths in a format that can be used later in the coordination process. Although it is based on the same principles as Dijkstra's Algorithm, it has the advantage that it generates a data structure that can be used readily for later processing and from which suitable other paths can be generated. The procedure itself identifies suitable vertices using only the matrices that define the grid structure, and by not involving all vertices provides a more efficient method of identifying shortest paths and their lengths in the rectilinear grid.

The length of the shortest path between each pair of nodes, and between each node and the depot, is calculated using Procedure 5.1. Beginning at the source node, a sequence of paths is traversed until either an arbitrary distance limit is met or the destination is reached. At each intersection, the rectilinear nature of the grid is exploited by recognising that no more than four branches may form the intersection. The branches are followed in the arbitrary order of preference east, north, west and south, where north is taken as the direction of the y-axis of the grid, where those edges exist. The distance travelled to each intersection is recorded so that the traversal can be backtracked to the latest intersection should the current path prove fruitless. Procedure 5.1 refers to two functions, CreatePath and FAR.

CreatePath(source,destination,EH,EV,TL,X) creates a path from source to destination using the network defined by EH and EV but which will not travel through point X and will not exceed length TL. The path is stored in a 9-column matrix, NetPath, with structure:

Column								
1,2	Location on active path							
3,4	Location on 1st alternate path							
5,6	Location on 2nd alternate path							
7	Distance from destination							
8	Distance travelled so far							
9	Boolean: is sum of columns 7,8 > target length? (1/0=Y/N)							

For example, part of the path from 6,12 to 7,11 in Figure 5.4 would be represented as

6	12	0	0	0	0	2	0	1
7	12	5	12	6	11	1	1	1
7	11	0	0	0	0	0	2	1

indicating that at the location (6,12), alternative paths existed leading from vertices at (7,12), (5,12) and (6,11).




Figure 5.4: Grid section.

The function $\text{FAR}(\text{oldPath}, \text{EH}, \text{EV}, \text{TL}, \text{X})$ finds routes other than oldPath from the source (identified as the start of oldPath) to the destination (identified as the end of oldPath) through the network defined by EH and EV , without travelling through X and without exceeding a distance of TL units. It is very like CreatePath except that it starts by backtracking through oldPath to the latest intersection and then grows a new path from there. Consequently, it finds all shortest paths between source and destination.

```

/* Length of shortest paths
/* Initialise
/*TL = upper limit of path length, NumNodes = number of nodes
GS := size(EH,1) /*Grid Size
NetDist := repmat(0,NumNodes,NumNodes+1)
FOR Sid := 1:NumNodes          /*Source ID
  /*Inter-nodal distances from Source to Destination
  FOR Did := Sid+1:NumNodes    /*Destination ID
    TL := 2*GS
    ThisPath := CreatePath(Sid,Did,EH,EV,TL,[0,0])
    TL := ThisPath(end,8)-1
    WHILE size(ThisPath,1)>1
      TL := ThisPath(end,8)-1
      oldThisPath := ThisPath
      ThisPath := FAR(ThisPath,EH,EV,TL,[0,0])
      newThisPath := ThisPath
    END WHILE
    NetDist(Sid,Did) := oldThisPath(end,8)
  END FOR
/*Node-Depot distances
TL := 2*GS
ThisPath := CreatePath(Sid,[Depot],EH,EV,TL,[0,0])
TL := ThisPath(end,8)-1
WHILE size(ThisPath,1)>1
  TL := ThisPath(end,8)-1
  oldThisPath := ThisPath
  ThisPath := FAR(ThisPath,EH,EV,TL,[0,0])
  newThisPath := ThisPath
END WHILE
NetDist(Sid,Did) := oldThisPath(end,8)
END FOR

```

Procedure 5.1: Finding the length of shortest paths between nodes.

```

/* CreatePath
/* Initialise
CarryOn := 1; options := [0 0 0 0]; NetPath := [0 0 0 0 0 0 0 0]
NetPath(1,1:2) := S; CurrentPosition := NetPath(end,1:2)
NetPath(1,7) := |distance between CurrentPosition and D|
IF NetPath(1,7) ≤ TL
    NetPath(1,9) := 1
END IF
WHILE CarryOn
    LineToAdd := [0 0 0 0 0 0 0 0]; BackTrack := 0
    IF ∃ an east of CurrentPosition
        cat1 := 1
    END IF
    IF ∃ an edge north of CurrentPosition
        cat2 := 2
    END IF
    IF ∃ an edge west of CurrentPosition
        cat3 := 3
    END IF
    IF ∃ an edge south of CurrentPosition
        cat4 := 4
    END IF
    IF size(NetPath) > 1
        OldPosition := NetPath(end-1,1:2)
        Compare CurrentPosition with OldPosition and adjust cat values
        so that backwards travel does not occur.
    END IF
    options := [cat1,cat2,cat3,cat4]; Remove zeros from options
    NP := [0 0 0 0 0 0 0 0]
    WHILE any(options)
        WayToGo := min(options)
        adjust NP to reflect WayToGo
        options(WayToGo) := 0
    END WHILE
    LineToAdd(1:8) := NP(1:8)
    IF LineToAdd(1:2) == X
        LineToAdd(1:6) := LineToAdd(3:8)

```

```

ELSEIF LineToAdd(3:4) == X
  LineToAdd(3:6) := LineToAdd(5:8)
ELSEIF LineToAdd(5:6) == X
  LineToAdd(5:6) := LineToAdd(7:8)
END
adjust LineToAdd(7)
LineToAdd(8) := NetPath(end,8)+1
adjust LineToAdd(9)
IF ~LineToAdd(9) | ~LineToAdd(1)
  BackTrack := 1
END IF
IF LineToAdd(9) & LineToAdd(1)
  append LineToAdd to NetPath
  OldPosition := CurrentPosition
  CurrentPosition := NetPath(end,1:2)
  BackTrack := 0
END IF
IF D has been reached, i.e. ~LineToAdd(7)
  CarryOn := 0; BackTrack := 0
END IF
WHILE BackTrack
  BackTrack := 0
  FirstLineToRemove := max(find(NetPath(:,3)))
  Shuffle := 1
  WHILE Shuffle
    Shuffle := 0; BackTrack := 0
    FirstLineToRemove := max(find(NetPath(:,3)))
    IF FirstLineToRemove > 1
      truncate NetPath from FirstLineToRemove to end
    ELSE
      LineToAdd(1) := 0
      CarryOn := 0; Shuffle := 0; Looping := 0; BackTrack := 0
      NetPath := [[S] 0 0 0 0 0 0 0]
    END IF
    IF LineToAdd(1)
      Looping := 0
      LoopHere := find(NetPath(1:2) == LineToAdd(1:2))
      IF LoopHere

```

```

    Looping := 1; Shuffle := 1; BackTrack := 0;
    FirstLineToRemove := max(find(NetPath(:,3)))
  END IF
ELSEIF first two elements don't exist
  BackTrack := 1; Shuffle := 0; Looping := 0
  ELSE
    CarryOn := 0; Shuffle := 0; Looping := 0; BackTrack := 0
    NetPath := [[S] 0 0 0 0 0 0 0]
  END IF
END IF
IF ~Looping & CarryOn
  update LineToAdd(7)
  LineToAdd(8) := FirstLineToRemove - 1
  update LineToAdd(9)
  IF ~ LineToAdd(9)
    Shuffle := 1
  ELSE
    OldPosition := NetPath(FirstLineToRemove-1,1:2)
    CurrentPosition := LineToAdd(1:2)
    truncate NetPath from FirstLineToRemove to end
    append LineToAdd to NetPath
    Shuffle := 0; BackTrack := 0
    IF ~ LineToAdd(7)
      CarryOn := 0
    END IF
  END IF
END IF
END WHILE
END WHILE
END WHILE

```

Procedure 5.2: Creating a path between two points.

Consider the example in Figure 5.5 in which the shortest distance between nodes 1 and 2 is sought. Starting from node 1, the no-choice-string is followed to the first intersection (A) where the eastern branch is taken. Similar action is taken at each intersection, with the northern branch taken if an eastern branch does not exist. At E, the path leads to a cul de sac (node 10), forcing the path to back-track to the previous intersection (E) and assume the next branch (north). When the

path, shown in green, eventually arrives at its destination, node 2, a record is made of the distance and a shorter path sought by reducing the limit to (distance found – 1) and resuming the search at the previous intersection (H). When node 2 is not reached before the path reaches a length equal to the limit, this new path is rejected and the search resumes from the next intersection (G) on the successful path. This process is repeated until a shorter path (red) is found, in this case using a successful branching from the intersection at B. The new best distance is recorded and the process repeated with a reduced limit using the intersection prior to B (A). Since this does not lead to a reduction in shortest path, the new shortest path is taken to include the red route that branches from B and the process is repeated along this path in a fruitless effort to identify a shorter one. Although other paths of equal length are identified, none is found to be shorter so the distance between nodes 1 and 2 is recorded in the matrix, NetDist, and the paths themselves can be extracted from oldPath.




Figure 5.5: Finding the shortest distance between two nodes.

5.5 Depot location

Although the depot can be located at any vertex in the grid, a realistic location would not be on the boundary. Once positioned, it affects the system in several ways:

- Because no node can be co-located with the depot in this study, nodes cannot be situated at that location; this may affect node density by creating a forbidden vertex.
- Grid reduction is influenced (see §5.7).
- Depot location relative to nodes influences route construction.
- The potential for coordination is changed by affecting:
 - the surplus distance availability for detours between coordination site and routes;
 - the number, location and type of coordination sites (see later);
 - the interactions between routes.

Although grid reduction is influenced by the location of the depot, the reduced grid is not limited in its effectiveness since any edges that have been removed in reducing the grid are no more useful than the remaining edges. Differences occur where the total distance constraint is invoked. When the difference between the total distance which a vehicle may travel and the route distance allotted to that vehicle is large, the vehicle in question can establish a coordination site further from the depot and it can detour further to collect payload from a coordination site. Thus, the outcome of coordination is sensitive to the location of the depot. Consider the 15-node systems of Table 5.1 and Figure 5.6, which differ only in the location of the depot (slightly off-centre at [10,8] or [12,12]). Detailed analysis of these schemes is given in Appendix B. The pre-coordination systems show only a small effect of depot location whereas the effect on the coordinated system is more pronounced; while one system shows a distance increase of $\sim 1\%$, the other has a reduction in distance of $\sim 12\%$.

The effect of re-locating the depot from a relatively central position to one on the edge of the grid is considered in the discussion of medium grid density later in this chapter.

For much of the following investigation, the depot is located slightly off-centre at (60%,60%), equivalent to [12,12]. The static location reduces the effect of depot location on the outcome of the final solution, enabling the consequences of coordination to dominate the results.




Figure 5.6: Slight changes in depot location can produce different outcomes. Upper: distance increase $\sim 1\%$, 1 coordination site. Lower: distance decrease $\sim 12\%$, 2 coordination sites.

Depot location	Without coordination		With coordination		
	Total dist	Fleet size	Total dist	Fleet size	No. coordn sites
[12,12]	226	6	228	5	1
[10,8]	222	6	198	5	2

Table 5.1: Slight differences in depot location can have large differences in outcomes.

5.6 Connectedness of the grid

- An *outcrop* is defined as a small string of edges that has few neighbouring edges other than the string that permits connection to the rest of the grid.

A modified coordination hull (outlined in red in Figure 5.7) is formed by the nodes and depot(s) with regard to the edges connecting them. When a node is located on an outcrop of edges, the edges connecting that outcrop to the main body of edges must be included within the coordination hull. Only that part of the grid that falls inside the boundary of such a modified coordination hull need be connected. As in the Euclidean case (see Chapter 4), coordination can occur only within the coordination hull. All edges that are outside the hull are extraneous and can be eliminated from consideration by a process of grid reduction.




Figure 5.7: Sparse grid network showing a modified coordination hull (red outline).

5.7 Grid reduction

- An *essential point* is defined as a node or depot.

Grid reduction is the process of removing from a rectilinear grid edges that are not vital to the application of coordination. Edge removal is predicated upon the assumption that an edge has value only if it touches an essential point, or if it contributes to a shortest path between two essential points, or if the vertices at either end provide a site for coordination. As a result of such a reduction, all edges outside the coordination hull will have been removed from consideration, thereby removing edges that will not contribute to coordination and that will not be part of any shortest path between nodes. The removal of edges within the coordination hull reduces the number of calculations in determining the shortest paths and reduces the number of edges that need to be considered during route construction.

Maze routing procedures are used in very large scale integrated chip design and robot planning to reduce the total length of paths where the underlying frame of reference is a mesh grid. In order to solve the maze-routing problem on a reconfigurable mesh architecture, Ercal and Lee (see [E97]) propose a tree pruning

algorithm to eliminate edges in the maze not on the shortest path between two points. It is similar in principle to the method of rectilinear grid reduction (see [R98]) used in this study, where edges not on a shortest path between two intersections are considered for removal. Hertz and Mittaz (see [HM01] and [H00]) employ a similar concept in replacing a chain, C , of traversed but unserviced edges by a shortest chain linking the endpoints of C .

5.7.1 Vertex classification

The vertices of a grid can be classified according to the number of edges impinging upon them, as in Table 5.2 and Figure 5.8.

- A *cul-de-sac* is defined as a string of edges that does not lead to more than one intersection.
- A *no-choice-string* is a string of edges that includes no intersections, other than those at either end of the string.

Once a no-choice-string is encountered, a vehicle is compelled to continue along it in the prescribed direction of travel until the no-choice-string has been traversed. The only intersections associated with a no-choice-string are those that define the ends of it. In some cases, a no-choice-string may be replaced by an alternative route from one end of the no-choice-string to the other.

Number of Edges	Vertex Classification
0	Unused vertex
1	End of a cul-de-sac
2	Member of no-choice-string
3	3-way intersection
4	4-way intersection

Table 5.2: Classification of vertices with respect to the number of edges attached to them.




Figure 5.8: Vertex classification.

A cul-de-sac that does not include a node or depot will never be used as an alternative route or as a site for coordination since it does not lead to an essential point; it is superfluous to any consideration of vehicle routing and can be removed from the grid with no concomitant loss of grid effectiveness.

A no-choice-string that does not include a node or depot and that is no longer than an alternative route between its ends can always be replaced by another route of the same length or less; it is also superfluous to the overall routing scheme.

As cul-de-sacs and no-choice-strings are removed from the grid, the degree status of vertices incident on removed edges changes. These vertices need to be re-classified and re-assessed until they do not exhibit changing status. As an example, consider the grid of Figure 5.9 in which removal of the cul-de-sac converts two edges into a no-choice-string, which subsequently might be removed. Although the final grid shows two possible arrangements, they are equivalent and all shortest paths between the nodes will be considered during the investigation of coordination. Had the cul-de-sac not existed, as shown in the lower part of Figure 5.5, the two edges would have been classified as no-choice-strings from the outset and been removed.






Figure 5.9: Classification changes with edge removal.

5.7.2 Reduction procedures

Initially, only those edges not involving essential points are examined for removal (shown schematically in Procedure 5.3). When all such edges have been dealt with, the edges involving essential points are considered (see Procedure 5.4). The process is repeated until all extraneous edges have been removed.



Procedure 5.3: Grid reduction by removal of extraneous edges.



Procedure 5.4: Removal of edges involving essential points.

Every no-choice-string that remains in the reduced grid can be considered as a representative of the phase space of shortest paths between the two nodes defining the ends of the no-choice-string. For example, any one of the four strings shown in Figure 5.1 could remain in the reduced grid, unless they were all subsequently eliminated in the reduction process.

The no-choice-string (red) in Figure 5.10, between points A and B, is the same length as the path (black) between A and B that passes through C. The presence of intersections along the black path provides the opportunity for interactions with other paths and allows more flexibility in routing since sections of the black path can be used by vehicles other than those travelling between A and B, whereas the red path is used only by vehicles travelling between A and B. The red path can be replaced by the black path but not vice versa. Consequently, any no-choice-string is replaced by a path of equal length that includes an intersection.



Figure 5.10: A no-choice-string is removed if it is no shorter than another path between the same endpoints. Numbers indicate length of path between A and B.

Where two no-choice-strings have the same end-points, as in Figure 5.11, either can be chosen as the representative path in the reduced grid because the distance between A and B is the same regardless of path taken and because each path has identical characteristics. Since each is a staircase path, neither has an advantage over the other, regardless of the number of corners.




Figure 5.11: Of two no-choice-strings with the same endpoints, no more than one is retained in the reduced grid.

When considering no-choice-strings that include essential points within the string, the no-choice-string is divided into two pieces at the essential point, and each piece is considered separately. The presence of the node or depot alters the

priority given to the paths by ensuring that a section of no-choice-string defined by the end of the no-choice-string and the essential point is retained if there is no shorter path from the essential point to the appropriate end of the no-choice-string. The upper grid in Figure 5.12, shows a node at point D that is 1 unit from one end (B) of a no-choice-string and 2 units from the other end (A). The string is divided at the node and each piece considered separately:

Red: $\text{length}(DB) < \text{length}(DACB) \Rightarrow \text{edge DB is retained}$

Blue: $\text{length}(DA) < \text{length}(DBCA) \Rightarrow \text{path DA is retained}$

Both the 1-unit and 2-unit sections are retained because the alternative paths are longer in each case.

The lower grid in the same Figure shows a slightly different situation where the node is 1 unit from one end and 4 units from the other end of a no-choice-string.

Red: $\text{length}(DB) < \text{length}(DACB) \Rightarrow \text{edge DB is retained}$

Blue: $\text{length}(DA) \geq \text{length}(DBCA) \Rightarrow \text{path DA is not retained}$

Using the same criteria as before, the 1-unit section is retained because the alternative path is longer, but this time the 4-unit section of no-choice-string is removed because it is no shorter than the alternative path. As a result, the no-choice-string has been converted into a cul-de-sac. Had the node not existed at D, the entire no-choice-string would have been eliminated in the reduction process. ■




Figure 5.12: No-choice-strings that include essential points.

This technique of dividing no-choice-strings containing essential points at the point ensures that grid reduction does not result in isolated nodes and depots.

When the essential point is at the end of a no-choice-string, similar considerations determine which edges, if any, to remove. As above, a path can be removed if it is no shorter than another path between the same endpoints, and as long as the point is not left isolated. The grid in Figure 5.13 shows two situations involving a node at the end of no-choice-strings. In the upper grid, the node lies at the intersection (B) defining one end of a no-choice-string and is 2 units away from the other end of the string (A) regardless of which path is taken. Consequently, the no-choice-string is removed without losing the shortest path between node and point A. If the degree of B had originally been 3, the node would now be part of another no-choice-string; if the degree of B had been 4, the node would now lie on a three-way intersection, possibly part of other no-choice-strings depending on the degree of its neighbours.

The lower grid in the same Figure shows another case where the node lies at the end of a no-choice-string, except that this time the alternative path is longer than the string, resulting in the retention of the no-choice-string in the reduced grid.




Figure 5.13: Node or depot at end of no-choice-string.

Where more than one essential point is contained in a no-choice-string, the constraints are amended to include:

- a string or part-string will not be removed if its retention was confirmed when considering another essential point, and
- a string or part-string will not be removed if it forms the shortest path between two essential points.

This is achieved without calculating the shortest paths between nodes, by dividing the no-choice-string into $(n+1)$ sections (where n is the number of nodes on the no-choice-string) depending on which node is being investigated, as shown in Figure 5.14. For each section, the distances involved in travelling from the node (red in Figure 5.14) to the other end of the section using both paths are compared and the shorter path accepted as above; in the case of a tie, the path including the intersection is preferred. The lower part of Figure 5.14 shows the comparisons that are involved when determining which part of the no-choice-string to retain with respect to the red node.




Figure 5.14: Several nodes on a no-choice-string. The arrows indicate the divisions of the no-choice-string with respect to the relevant node (red).

An example is given in Figure 5.15 which shows two nodes (at B and D, respectively) where B is 1 unit from an intersection at C, and D is 1 unit from the intersection at A. When considering the node at B, edges BD and AD would be removed if not for the presence of the node at D. Had they been removed, the node at D would have been isolated. Since the node at D lies inside a no-choice-string, consideration of the string surrounding it involves division of that string into two parts at D. Because $\text{length}(AD) < \text{length}(ACBD)$, edge AD is retained in the reduction, regardless of the result from consideration of the node at B. Similarly, because $\text{length}(BD) < \text{length}(BCAD)$, edge BD is retained regardless of any other consideration. The retention of edges when considering a particular node is absolute and outweighs any argument for their removal by any other means.






Figure 5.15: Two nodes on a no-choice-string.

In Figure 5.15, edge BD would have been retained anyway since it is the shortest path between the two nodes. This retention of the shortest path is demonstrated in Figure 5.16 where there is no other excuse for retaining edge BD. In this case, $\text{length}(BC) < \text{length}(BDAC)$, ensuring the retention of the edges making up BC; $\text{length}(DA) < \text{length}(DBCA)$, ensuring the retention of the edges making up DA. Edge BD is retained solely because it represents the shortest path between the nodes at B and D. Consequently, the grid-section in Figure 5.16 is fully reduced unless the degree status of points A and C changes.






Figure 5.16: Edge between nodes is retained.

The criteria for edge removal depend on the replacement capability of edges (so that no two equivalent paths travel between contiguous intersections) and on the presence of cul-de-sacs, thus facilitating the determination of shortest paths. As a result, not every remaining edge will be involved in a shortest path between essential points, as demonstrated in Figure 5.17. Although the grid in the Figure can not be reduced further, it contains several edges (shown in red) that do not form part of any shortest paths between essential points.




Figure 5.17: In this reduced grid, the red edges are not part of any shortest paths between nodes.

The reduced grid is constructed of a combination of cul-de-sacs and no-choice-strings in which every cul-de-sac is terminated by a node or depot and no-choice-strings are terminated by intersections and/or essential points. Every no-choice-string is unique in that it cannot be replaced by a shorter path between the ends.

Once the grid has been reduced, routes can be constructed more easily, since there are fewer edges to consider. Typical reductions are shown in Appendix A.




Figure 5.18: Grid before (upper) and after (lower) reduction.

The number of edges removed during grid reduction depends on the physical nature of the initial grid and the location of essential points. As the number of edges in the initial grid increases, more vertices are employed, and the boundary of the coordination hull tends to adopt a more regular shape. The effect of grid reduction was investigated for grids of each of the three densities constructed with ten problem sets for each node density varying from 10 nodes to 300 nodes with the results shown in Figure 5.19. For the particular grids chosen (shown in Appendix A), the sparse grid contained 364 vertices, the medium grid contained 439 vertices, and the dense grid contained 441 vertices, any of which was capable of being a node location.




Figure 5.19: Effect of number of nodes on grid reduction. Colours indicate grid density – red: sparse; green: medium; blue: dense. Dotted lines indicate no. edges in original grid before reduction.

As the number of nodes increased to about 300, the number of edges remaining in the reduced grid approached about 90% of the number of edges in the original grid, and the range of values decreased. The reduction in range of values with increasing node density is the result of several factors. Increasing the number of essential points reduces the number of edges to be removed during the reduction process, since edges that include essential points are treated differently from those without: cul-de-sacs with essential points remain, whereas those without are

removed; no-choice-strings containing essential points can result in some or all pieces of the string being retained, as in Figures 5.12, 5.15 and 5.16.

Increasing the edge density reduces the number of cul-de-sacs thereby reducing one of the major causes of edge reduction.

Another effect of increasing edge density and increasing node density is the lowering of the number of exposed no-choice-strings (that is, no-choice-strings that include a corner). Reduction of the grid in Figure 5.20(a) is initiated by the absence of several edges at the arrowed location which is an exposed no-choice-string (shown in blue) resulting in the grid shown to the right. In this case, it is the same length as the no-choice-strings to the right of the upper left node and below the upper right node, but since these two no-choice-strings are not exposed, they do not present a vulnerable point at which grid reduction can be initiated. This reduction would have been prevented by

- increasing the edge density to remove the exposed no-choice-string, as in (b), or by
- converting the no-choice-string to include a node, as in (c).

Thus, increasing the number of nodes or increasing the number of edges reduces the effect of grid reduction by removing exposed no-choice-strings, as shown in Figure 5.19.




Figure 5.20: Grid reduction is affected by node location and edge density.

and dashed strings as any point on the string that is replaced, thereby preserving path proximity. ■

Whenever a string is replaced by a shorter string, the distance between vertices on two paths is no greater than the distance between the source vertex and the destination vertex on the replaced path.

5.7.4 Uniqueness of reduction

Since the number of edges attached to a vertex cannot increase during grid reduction, intersections are neither created nor increased in complexity, only left untouched, simplified or eliminated. String removal occurs only between adjacent intersections (for no-choice-strings) or between intersection and string end (for cul-de-sacs). No sequence of edges can be removed if there is an intersection inside it.

Results of reduction may not look the same but they are functionally comparable. For example, whether the dotted line or dashed line is retained in Figure 5.21, there will still be five vertices between the white vertices on the blue path. ■

Consider also the reduction shown in Figure 5.22 in which choice of the initial no-choice-string to remove results in different subsequent steps but the same outcome. Initial removal of the BC no-choice-string gives a more rapid result than initial removal of the AC no-choice-string, which involves the subsequent removal of a cul-de-sac. Regardless of the choice of initial no-choice-string for removal, the result is unchanged because intersections are not created. ■




Figure 5.22: Piecewise reduction of a grid section.

5.7.5 Effect of node location

Edges that involve nodes are treated differently from other edges in order to maintain access to the node, thus the layout of the nodes influences the reduced grid by affecting the shape of the coordination hull and the regions surrounding each point. Grid reduction effectively “frays” the areas of a grid in much the same way as a moth frays the weak spots of a woollen sweater, by using irregular areas as starting points for erosion (reduction).

Figure 5.23 shows the influence of an essential point on the neighbouring reduction while maintaining edge density in a complete grid. Four nodes at the corners of the grid create a coordination hull, within which there is no opportunity for reduction to be initiated. When one of those nodes is moved, the effect on the hull and subsequent reduction is dramatic, with some edges being eliminated. When all four nodes are moved to adjacent positions, the coordination hull is minimised and most edges are eliminated.




Figure 5.23: Effect of node location on a complete grid.

When a complete grid has some edges removed to form a grid of lesser density, the absence of edges can lead to the creation of exposed no-choice-strings (see Figure 5.20) and to other structural characteristics reminiscent of geographical features such as ‘bays’, ‘streams’, ‘lakes’ and ‘peninsulas’, as indicated in Figure 5.24. As in geography, a definition of such features is at best vague:

- *lake* may be thought of as an area that contains no edges but which is bounded by a path
- *peninsula* is a region of grid bounded by a path and connected to another (usually larger) area of grid by a no-choice-string
- *bay* is a region of grid with an absence of edges that has, as at least one boundary, the edge of the grid, and has a somewhat regular boundary otherwise
- *stream* is like a bay but which may contain many turns, usually longer than it is wide.

The relative location of nodes to these structural features influences the reduction outcome, as shown below.




Figure 5.24: Geographical features may be recognised in the structure of a rectilinear grid.

The presence of essential points on peninsulas can have a dramatic effect on the shape of the coordination hull; the location of even one node on a peninsula ensures that the part of the feature connecting the node and the main body of the grid will be retained, even if that connection is serpentine. Had there been no

essential point on the peninsula, the entire feature would have been eliminated during the reduction process.

The effect of nodes near lakes depends on the presence or absence of exposed no-choice-strings and cul-de-sacs forming part of the feature's outline. In Figure 5.20, which shows nodes close to an area now identified as a lake, the presence of a node on the exposed no-choice-string halts any reduction of that and subsequent no-choice-strings.

Streams, where they exist in the reduced grid, increase the perimeter of the coordination hull. The presence of nodes near, and on opposite sides of, such features changes the shape of the final grid because the shortest path between the nodes is forced to pass around the stream. Figure 5.25 shows two nodes 1 unit apart on opposite sides of a stream, with a shortest path between them shown in pink. Even if no other essential points lie in the grid shown, the pink path will be retained in the final reduced grid; had the nodes been on the same side of the stream, the stream feature would have been eliminated entirely in the reduction process (in the absence of other nodes).




Figure 5.25: The effect of nodes on opposite sides of a stream is to retain some semblance of the stream in the reduced grid.

The effect of node location on the reduced grid can be seen in Figures 5.26 and 5.27 below (in which two grids differ only in the location of essential points), each of which involves ten nodes and one depot on the same initial grid with the same depot location. Although they differ only in node location, one result is compact while the other is spread out. The number of edges removed differs markedly depending on the hull that describes the essential points, and increasing the number of nodes within that hull does not necessarily directly influence the reduction process. Consequently, the complexity of the grid reduction procedure is difficult to determine.




Figure 5.26: Before and after grid reduction. A compact result.





Figure 5.27: Before and after grid reduction. A spread result.

5.8 Coordination opportunities

Coordination sites are limited by the nature of a grid to vertices and edges and can not occur at the spaces in between and, for the purposes of this study, coordination is further restricted to occur only at vertices. There is therefore only a finite number of possible coordination sites, all of which are known in advance of the route construction phase. This has the advantage that it allows emphasis to be placed on the paths rather than on the nodes, in contrast to the situation in the Euclidean plane, where all paths were limited to straight lines between nodes.

In the rectilinear grid, route deviation can occur at any vertex on the existing path, including a vertex co-located with a node. There are therefore a finite number of deviation points. The types of coordination possible include load splitting, load swapping and changing the time at which a multi-routed vehicle returns to the depot.

5.9 Solution construction

5.9.1 Initial data

A fixed number of nodes are created at random vertices, each with known delivery demand and zero pickup demand. The fleet of identical vehicles, each with a maximum capacity, maxVehCap , and maximum total distance allowance, maxVehDist , begins the work cycle from the single depot at the same time and all routes end at the depot. All edges of the rectilinear grid are two-way and have infinite carrying capacity in each direction. Load splitting is permitted, and the demand of each customer must be satisfied during the current work cycle.

5.9.2 Route construction

Routes are constructed using a two-phase process such as that suggested by [FJ81]:

- phase one involves choosing a seed node and clustering other nodes about it;
- phase two constructs routes within the clusters.

The vehicle capacity constraint limits total payload carried by any vehicle at any particular time to a value not exceeding maxVehCap . Although one vehicle may service more than one route, no route may be served by more than one vehicle. Thus, maxVehCap limits the availability of nodes for inclusion in a cluster by not

allowing any node to be accepted into a cluster if doing so would mean that the total demand of nodes in that cluster exceeded maxVehCap . A first-come-first-served principle applies in that once a node has been accepted into a route, it will not be rejected in favour of another even if the demands can be carried more efficiently.

The vehicle distance constraint is accommodated by limiting the size of a cluster. Any modification to the cluster can be made in an improvement step if the vehicle distance constraint is violated.

Several clustering methods [B93] were investigated and evaluated in conjunction with the routing methods.

- Smallest neighbourhood

The sum of the distances from the three closest unassigned nodes to each node is determined. The seed node is the node that minimises the sum of the distances to its three closest unassigned neighbours. Other neighbourhood sizes were investigated but smaller neighbourhoods produced results that required more improvement to produce realistic results, while larger neighbourhoods produced results that were no better than the size 3 neighbourhood.

- Largest neighbourhood

As for the smallest neighbourhood, but the seed node is that node with the greatest sum of distances to its three closest unassigned neighbours.

- Maximum distance from depot

The seed node is that node which is farthest from the depot.

- Isolated node

Choose as the seed node, an unassigned node that is the most distant from the depot or from other unassigned nodes.

- Path deviation

The seed node is that node that is most distant from the depot. Nodes are assigned to the seed based on how much they deviate from the shortest path between the seed node and the depot.

Once nodes have been clustered, they are routed within each cluster either by a farthest insertion method or by a nearest insertion method. The initial node in each emerging route is the depot; the next node to be included is that unassigned node that is farthest/nearest to the last node assigned. Inter-nodal distances are determined by the smallest number of edges in a path leading from one node to the other. The insertion method is further adjusted by k-opt versions of inter-route

node transfer / node exchange / edge exchange and intra-route edge exchange. Feasible routes are created using Procedure 5.5.

Step One	Create routes based on clustering algorithm and nearest/farthest insertion with regard to capacity constraint.
Step Two	Improve these routes using edge exchange and k-opt procedures.
Step Three	With regard to distance constraint, improve these routes using edge exchange and k-opt procedures.

Procedure 5.5: Route construction without coordination.

Generally, routes generated using a smallest neighbourhood/nearest insertion method in Step One require fewer adjustments in Steps Two and Three, and result in shorter distances than other combinations.

The differences are illustrated in Table 5.3 in which maxVehCap has value 200, maxVehDist has value 50 and no splitting of demands is permitted. On improvement, nearest insertion and farthest insertion produced routes that differed in orientation.

LNNI and LNFI

0 6 13 11 0	dist 48	load 195
0 12 5 0	dist 24	load 72
0 7 3 0	dist 38	load 130
0 9 1 8 14 0	dist 32	load 178
0 15 10 4 0	dist 48	load 195
0 2 0	dist 28	load 166
Total distance	218	

SNNI and SNFI

0 11 7 10 4 0	dist 38	load 181
0 12 3 5 0	dist 34	load 198
0 8 2 0	dist 40	load 169
0 15 1 0	dist 26	load 193
0 13 6 0	dist 40	load 177
0 14 9 0	dist 32	load 132
Total distance	210	

Table 5.3: Route generation using different clustering methods.



Smallest neighbourhood heuristics have been reported [B93 p25] as resulting in clustering nodes that are close to each other and consigning isolated nodes to later routes of greater distance. In this study, the presence of demands approaching the

value of maxVehCap and the application of vehicle capacity constraints limited the number of nodes on a route and prevented such clustering problems from occurring.

5.9.2.1 Inter-route node transfer

Inter-route node transfer consists of removing a node from the route to which it was assigned during route generation, and inserting it into another existing route or creating a route for which it is the only node. This improvement method, therefore, can create and delete routes.

When a node is transferred from one route to another, three arcs are destroyed and three created. Each arc may consist of one or more edges, as in Figure 5.28. Transfer is deemed worthwhile when

$$[d + e + f] - [a + b + c] < 0$$

subject also to the vehicle capacity constraint,

$$\text{total demand on route created with transferred node} \leq \text{maxVehCap}.$$

The vehicle capacity constraint is not applied to the route from which the node has been transferred since that node is being removed from a route which was already constrained by maxVehCap , making it impossible for the total demand on that route to be increased.




Figure 5.28: Inter-route node transfer. The arrowed node is transferred from one route to another.

Positioning of the transferred node in the new route is carried out in a manner that minimises the sum of the lengths of the created arcs. Where the node being transferred lies initially on a single-noded route, that route ceases to exist after transfer. Where the node is transferred to a single-noded route, a new route is

created. Changing the number of single-noded routes may have consequences in the type of coordination that can be applied (see §5.14).

Transferring a node from one route to another has temporal consequences as well as spatial ones. All nodes on the enlarged route subsequent to, and possibly including, the transferred node are visited later than they would have been had the node not been transferred. Similarly, nodes on the route from which the transferred node originated may be visited earlier than they would have been had the node not been transferred. Such changes may affect the manner in which coordination is applied (see §5.13).

5.9.2.2 Inter-route node exchange

When two nodes are exchanged between routes, four arcs are destroyed and four created, as shown in Figure 5.29. The exchange is considered worthwhile if

$$[e + f + g + h] - [a + b + c + d] < 0$$

subject to the vehicle capacity constraint.

Since two nodes are being exchanged, they usually occupy the position in the new route that is being vacated by the node that they are replacing, but this is not always the case, as in Figure 5.30. When a simple exchange is not adequate, a modified k-opt improvement is used to find the best position for each node. Because placement of the nodes is independent of the positions from whence each came, they can be considered separately.




Figure 5.29: Inter-route node exchange. The arrowed nodes are exchanged between routes.

Such an improvement of one of the final routes is shown in Figure 5.30 in which the lengths $(g' + h') < (g + h)$, resulting in one of the transferred nodes not replacing the other transferred node in the sequence. Consequently, each inter-route node exchange is accompanied by a k-opt improvement to determine the insertion positions that result in the greatest reduction in total distance. Because

rectilinear grid coordination is sensitive to the closeness between routes for detour purposes (see e.g. §5.13), failing to improve the routes created by inter-route node exchange could result in an environment either that is not suitable for coordination, or that appears falsely suitable for coordination but which ought to have been improved by a k-opt process.




Figure 5.30: Inter-route node exchange requiring a different placement of a transferred node (arrowed).

5.9.2.3 Intra-route edge exchange

During intra-route edge exchange, the sequence of nodes is altered by swapping two pairs of adjacent nodes within a node sequence and causing at least one of the edges to be traversed in the opposite direction. In Figure 5.31, the sequence of edges (*a b c d e f*) becomes (*g e h i b j*); the node sequence (1 2 3 4 5 6 7) becomes (1 5 6 4 2 3 7).




Figure 5.31: Intra-route edge exchange. Edges b and e are swapped within a route.

The exchanged edges may contain a common node as in Figure 5.32 in which edges b and c , both of which contain node 3, are exchanged. The edge sequence $(a\ b\ c\ d)$ becomes $(e\ c\ b\ f)$ with both exchanged edges being traversed in the opposite direction, and the node sequence $(1\ 2\ 3\ 4\ 5)$ becomes $(1\ 4\ 3\ 2\ 5)$.



Figure 5.32: Intra-route edge exchange in which the exchanged edges share a common node.

When edges are investigated for exchange, it is the effect on neighbouring edges that determines whether the exchange proceeds. For edges containing a common node, as in Figure 5.32, the exchange is worthwhile when

$$e + f < a + b$$

subject to vehicle capacity constraints, and when the edges do not share a common node, as in Figure 5.31, the exchange is worthwhile when

$$g + h + i + j < a + c + d + f.$$

That is, when the sum of the lengths of both the surrounding edges and the intervening edges is reduced by the exchange, the exchange can proceed, subject to vehicle capacity constraints.

In the context of rectilinear grid coordination, the effect of altering a node sequence or of reversing the orientation of a sequence can be extreme. As discussed in Chapter 3, the order in which nodes are visited and the time delays

between visits to a particular region by multiple vehicles help determine the suitability of sites for load transfer and load swapping, and also assist in determining the suitability of vehicles for delivering to, or collecting from, coordination sites. This is discussed in more detail in §5.13.

5.9.2.4 Inter-route edge exchange

As in the intra-route version of exchange, inter-route edge exchange depends for its viability on the effect on neighbouring edges. During exchange, shown in Figure 5.33, four edges are destroyed and four edges are created, with the exchange considered worthwhile, subject to vehicle capacity constraints, when

$$h + i + j + g < a + c + d + f$$




Figure 5.33: Inter route edge exchange. Edges b and e are swapped between routes.

Unlike the intra-route version, neither of the exchanged edges is required to be traversed in the opposite direction, so the order in which nodes that define the ends of the transferred edges are visited is unchanged. The times at which they are visited may be altered, which may affect the value of vertices along those edges as potential coordination sites. In addition, the creation of edges g , h , i and j may alter the potential of the region for coordination (see §5.13).

5.9.3 Route scheduling

In this study, the upper limit on the distance a vehicle can travel during a work cycle is maxVehDist . Because a maximum vehicle capacity exists, a route may be allocated nodes whose demands don't exceed maxVehCap but whose distance falls well within the maximum distance limit, thus allowing the vehicle servicing that route to service other routes before reaching the maximum distance constraint. The

consequences of allowing multi-routed vehicles are discussed in §5.14. The effect of varying the value of `maxVehDist` is discussed §5.10.

Once the routes have been constructed, they are allocated to vehicles in a scheduling phase in a manner that seeks to minimise the number of vehicles used without violating the distance constraint. At this pre-coordination stage, routes are treated as intact entities without being split, so, initially, any route that is longer than half the maximum vehicle distance is assigned to a new vehicle; no two such routes can share any vehicle without violating the distance constraint.

Several objectives for scheduling could be applied:

1. equitable times (including service times) away from the depot
2. equitable customer lists (i.e. each vehicle services a similar number of nodes)
3. maximum vehicle capacity
4. equitable travel distances

Objective 1 was not chosen for several reasons. The problem includes a maximum vehicle distance and not a maximum vehicle time, which allows the coordination method to be applied that seeks to obviate or delay the inter-route return to depot without consideration of the time taken. While such an added constraint could have been accommodated in this study, its inclusion does not add to the understanding of the essential processes involved in coordination and is therefore considered a prospect for future study. As a result, the only temporal consideration in this study is concerned with the location of vehicles and the demands of nodes. The stochastic demands of nodes is investigated in the next chapter.

Objective 2 was rejected because it, too, was not considered to advance the understanding of coordination processes. In practice, equitable customer lists are important where the service is customer-intensive rather than payload oriented.

Equitable travel distances are not always possible because some routes are created with distances equal to or close to the value of `maxVehDist`. Although the remaining routes could be compiled to form combinations of approximately equal length, there is no advantage in having most vehicles with spare capacity for a pure delivery scenario. Had the problem required pickup, as well as delivery, such spare capacity would have been desirable, but since there is no advantage in having spare capacity in each vehicle, it has not been targeted. Were equitable routes a favourable objective, the inclusion of a detour for the purposes of coordination (affecting one or more routes) might require the re-scheduling of all routes.

For these reasons a combination of Objectives 3 and 4 was adopted which results in most vehicles having similar distances to travel and sometimes allows one vehicle to travel a much shorter distance. After the initial allocation of routes with greater than half the maximum vehicle distance, remaining routes are allocated to vehicles so that the longest unallocated route is assigned to the vehicle with the smallest remaining free distance that can accommodate that route, thus producing vehicle distances that are reasonably equitable with respect to total distance travelled.

As an example, consider the scheduling of vehicles with routes of distances 50, 44, 32, 26, 24, 18, 16, 8, and 6 in a scheme where maxVehCap has value 50. The results of using Objective 3 and a combination of Objectives 3 and 4 are shown in Figure 5.34, where the application of combined Objectives 3 and 4 fills all but one of the vehicles and where the application of Objective 3 alone is unable to create equitable vehicle distances. The combined objectives can respond rapidly to increases in the distance of a route by utilising the spare capacity of vehicle 5 whereas Objective 3 requires juggling of all distances.




Figure 5.34: Scheduling vehicles using two different objectives. The dashed line indicates the value of maxVehCap.

A default value of maxVehDist of $(2.5 \times \text{the length of a grid side})$ entitles even a small ten-node system to have enough vehicles to allow the effects of coordination to be evident with little computation. The factor has been chosen to allow a vehicle to service any node on the grid, regardless of edge density, without violating the distance constraint. The effect of varying the value of maxVehDist is discussed later.

5.9.4 Incorporation of coordination

Once initial routes have been created, scheduled, and improved, situations can be investigated for their coordination potential. Rich environments include

- ❑ intersecting routes
- ❑ overlapping routes
- ❑ routes close to each other
- ❑ very short routes
- ❑ nodes close together but distant from other nodes, serviced by independent routes
- ❑ multi-routed vehicles
- ❑ routes that are slightly too long to be combined with others without violating distance constraints.

These situations may be suitable for coordination in the form of split deliveries, load swapping or deferral (or even removal) of inter-route return to depot, or some combination of these methods.

Three types of coordination site can be identified:

Type 1:

Site is visited by a vehicle before it starts its regular route, to deposit goods for later uplift by another vehicle or by the same vehicle. This type of site typically is established by a vehicle that has spare distance capacity but not enough carrying capacity to participate effectively in coordination. The site is created close enough to the depot so that the vehicle can reach it and return to the depot without violating the distance constraint before servicing the route that was allocated to it in the pre-coordination phase. Since the vehicle has no location requiring servicing other than the coordination site, it is able to carry a full payload to the site and deposit it there before returning to the depot preparatory to servicing its own route. The quantity of goods that can be deposited by a vehicle i at a type 1 site, $\text{deposit}_{i,1}$, is

$$\max(\text{deposit}_{i,1}) = \text{maxVehCap}$$

Since there is no limit on the quantity of goods that a site can accommodate, up to maxVehCap can be deposited each time such sites are visited in this manner. The same location can be used as a type 1 site by more than one vehicle. In Figure 5.35, the vehicle servicing the black route

establishes a type 1 site (arrowed) before returning to the depot, collecting the payload for its allocated route and servicing the nodes on that route. The deposit of a trailer by a stock truck in Problem 2 (see Chapter 1) is an example of the creation of a type 1 site in a pickup scenario.

Type 2:

Site is visited by a vehicle by detour (i.e. it travels to a location which would not be visited if the coordination site did not exist) in order to deposit or uplift goods. The deposit of goods enables another vehicle (or the same vehicle later) to collect and deliver them to another location; the collection of goods occurs after the site has been established by another vehicle (or the same vehicle previously) and enables the collection vehicle to service an area that it would not have had the distance capability of reaching or that it would not have had the capacity to carry the goods to service. In Figure 5.35, the vehicle servicing the green route detours to the coordination site while servicing the green route, and deposits or uplifts goods. The quantity of goods that can be deposited by a vehicle i at a type 2 site, $\text{deposit}_{i,2}$, is limited to the spare payload that the vehicle has at the beginning of its route when the demands of the $n(i)$ nodes that it services have been loaded, i.e.

$$\max(\text{deposit}_{i,2}) = \text{maxVehCap} - \sum_j \text{demand}(j) \quad j = 1, \dots, n(i)$$

and the quantity of goods that can be collected from a type 2 site by a vehicle that has serviced $m(i)$ of its $n(i)$ nodes, $\text{collect}_{i,2}$, is

$$\max(\text{collect}_{i,2}) = \text{maxVehCap} + \sum_{j=1}^{m(i)} \text{demand}(j) - \sum_{k=m+1}^{n(i)} \text{demand}(k)$$

$$j = 1, \dots, m(i)$$

$$k = 1, \dots, n(i)$$

where the nodes are serviced in the order $1, \dots, m, m+1, \dots, n$.

Figure 5.35 represents several scenarios involving a type 2 site at the arrowed location:

- the green vehicle establishes a type 2 site which later is visited by the red vehicle, enabling the red route to include a node whose demand it would not otherwise have had the capacity to accommodate;
- the red vehicle establishes a type 3 site (see below) depositing enough payload for the green vehicle to collect and deliver to a node that it would not have been able to service because of capacity constraints;

- the black vehicle establishes a type 1 site, delivering enough payload to fully service the green vehicle's second route; the green vehicle services the single node on its first route and, instead of returning to the depot to collect payload for the second of its routes, travels to the coordination site, collects all the payload previously deposited there, and uses it to service the nodes on its second route;
- some combination of the above scenarios.

Type 3:

Site is visited by a vehicle while on its regular route to deposit or collect payload; the vehicle is not involved in travelling any extra distance to use this site than it would if the site had not existed. The quantity of goods that can be deposited by a vehicle i at a type 3 site, $\text{deposit}_{i,3}$, is limited to the spare payload that the vehicle has when the demands of all its $n(i)$ nodes have been loaded, i.e.

$$\max(\text{deposit}_{i,3}) = \text{maxVehCap} - \sum_j \text{demand}(j) \quad j = 1, \dots, n(i)$$

and the quantity of goods that can be collected from a type 3 site by a vehicle that has serviced $m(i)$ of its $n(i)$ nodes, $\text{collect}_{i,3}$, is

$$\max(\text{collect}_{i,3}) = \text{maxVehCap} + \sum_{j=1}^{m(i)} \text{demand}(j) - \sum_{k=m(i)+1}^{n(i)} \text{demand}(k)$$

$$j = 1, \dots, m(i)$$

$$k = 1, \dots, n(i)$$

where the nodes are serviced in the order $1, \dots, m, m+1, \dots, n$.

In Figure 5.35, the red route uses the arrowed location as a type 3 coordination site in one of several scenarios:

- the black vehicle establishes a type 1 site at the location, depositing goods that the red vehicle will later uplift, enabling the red vehicle to service one or more of the later nodes on its route, thereby extending the range of the red route;
- the red vehicle establishes a type 3 site at the location depositing goods that the green vehicle will uplift.

A location can be used as more than one type of site, as shown in the discussion above in which the one site has two designations; it can be used in each

of the three types in the one scenario. For example, in Figure 5.35, the black vehicle establishes a type 1 site before attending to its own routes; the green vehicle detours to that location, using it as a type 2 site, to collect some of the load deposited there by the black vehicle; the red vehicle passes through the same location, collecting some of the load deposited there by the black vehicle, thereby using it as a type 3 site.




Figure 5.35: A site can have more than one classification depending on how it is used.

During the investigation of the potential of coordination, type 1 sites should be investigated before either of the other two types of site since type 1 sites can be created with maximum loads earlier than other types of site; the earlier a site is established, the greater the potential use of it as a type 2 or 3 site for collection purposes. It is preferable to create a type 1 site on the actual route of a vehicle to minimise extra distance by the collecting vehicle; such a location would then be used as a type 3 site by the collecting vehicle without incurring extra travelling distance. Similarly, type 2 sites should be created close to regular routes to minimise the addition of extra distance. Type 1 sites can be located using a surplus table which lists, for each vehicle, distance travelled, load, surplus distance and surplus load; type 1 sites can be established within half the surplus distance from the depot to enable the vehicle to return to the depot and service its own route without violating distance constraints. (Examples of this are discussed below.)

Vehicle #	Distance	Load	Surplus distance	Surplus load
1	$dist_1$	$load_1$	$\maxVehDist - dist_1$	$\maxVehCap - load_1$
2	$dist_2$	$load_2$	$\maxVehDist - dist_2$	$\maxVehCap - load_2$
...	...	...	...	...
p	$dist_p$	$load_p$	$\maxVehDist - dist_p$	$\maxVehCap - load_p$

Table 5.4: Surplus table for a system with p vehicles.

5.10 Maximum distance constraint

For much of this study, the maximum distance, maxVehDist , for the fleet of identical vehicles has been set arbitrarily at $(2.5 \times (\text{length of grid side}))$. This value was chosen because it is slightly greater than the sum of the horizontal and vertical boundaries of the rectilinear grid and it was expected that any single-noded route would not exceed this value. This allowed the creation of feasible single-noded routes.

When the maximum distance is constrained to be a lesser value, there is a danger that a node that is a long way from the depot will not be serviceable within one work cycle. That is, no vehicle will be able to deliver goods to the node and return to the depot without violating either the maximum distance constraint or the requirements to start from, and return to, the depot at the beginning and conclusion respectively of each work cycle. Thus, there is a minimum value of the maximum distance constraint, equal to twice the maximum distance between the depot and any node, $\text{Node}(i)$, ($i=[1, \text{NumNodes}]$).

$$\min(\text{maxVehDist}) = 2 \times \max_i(\text{Dist}(\text{Depot}, \text{Node}(i)))$$

Because of the grid structure and node location, the maximum distance between depot and nodes varies between problems, thereby altering the minimum value that maxVehDist can assume in order to create feasible solutions. Two possibilities for determining the value of maxVehDist for any problem present themselves:

1. Set MaxVehDist at such a high level that any problem will have a feasible solution.
2. Determine $\min(\text{MaxVehDist})$ for each problem after calculating the node-depot distances and adjust the value of maxVehDist to be no less than this. How much above this minimum to set the value will need to be determined for each problem.

This study uses possibility 1 since it involves fewer calculations and is expected to provide more general results.

As the maximum distance value is increased, more nodes can be serviced by any one route during the routing process and more routes can be combined for any vehicle during the scheduling process. Because the scheduling process works in discrete values, the relationship between maxVehDist and fleet size is not smooth

but varies in a stepwise manner, as shown in Table 5.5 and Figure 5.36 (data in Appendix D).




Figure 5.36: Effect of maxVehDist on fleet size.

maxVehDist	Fleet Size
30	6
40	6
50	5
60	5
70	4
80	4
90	3
100	3
110	3
120	3
130	2
...	...
240	2
250	1

Table 5.5: Effect of maximum distance constraint on fleet size.

Fleet size is affected by both the number of nodes that can be combined onto any route and the number of routes that can be serviced by any vehicle. As an individual route gets longer, its distance approaches the value of maxVehDist; long routes tend to be serviced by single-routed vehicles since there is not enough surplus vehicle distance to accommodate a second route. Thus, small values of

maxVehDist result in large fleets of single-routed vehicles and large values of maxVehDist result in small fleets of single-routed vehicles.

The interaction between maxVehDist and maxVehCap is discussed in §5.11.

5.11 Maximum vehicle capacity constraint

5.11.1 Coordination by methods other than split delivery

When splitting of demands is not permitted, maxVehCap determines the maximum value of an individual demand that can be dealt with. The maximum vehicle capacity influences the potential for coordination by affecting the existence and location of intersections, the surplus distance capacity of vehicles, and the surplus carrying capacity.

The value of maxVehCap is a factor in determining the number of nodes serviced by a vehicle and thus the number of routes, and thereby influences the number of interactions between routes, which in turn affects the potential for coordination.

As the value of maxVehCap increases, more nodes can be serviced by the same vehicle; individual route lengths tend to increase, and intersections tend to occur away from the depot. Since an increase in the distance of individual routes to values greater than $(\text{maxVehDist} \times 0.5)$ increases the fleet size (because these longer routes are unable to be combined without violating vehicle distance constraints) fleet reduction may be achieved by reducing the length of routes to values less than the value of $(\text{maxVehDist} \times 0.5)$, even if this results in an increase in total distance travelled. In this study, the values of maxVehDist and maxVehCap are user-defined in keeping with the real-world situation [e.g. NZTAA95].

The effect on the surplus carrying capacity influences the establishment of type 1 coordination sites and thereby influences the form of coordination that may occur.

Eventually the value of either maxVehCap or maxVehDist becomes a limiting factor in the route structure. Table 5.6 shows the results for a system in which only the value of maxVehCap varies (data in Appendix D); increasing it beyond 550 results in no reduction in distance or fleet size because the value of maxVehDist limits the route lengths.

maxVehCap	total distance	fleet size
200	318	8
220	288	7
240	267	7
245	268	7
250	256	6
300	240	6
350	228	6
400	232	6
450	222	5
500	230	6
550	214	5
600	214	5

Table 5.6: Typical effect of maxVehCap on total distance and fleet size.

The load plans for values of maxVehCap of 240 and 245 are shown in Table 5.7. As the value of maxVehCap is increased from 240 to 245, the number of multi-routed vehicles drops from 3 to 2, thus reducing the potential for distance reduction created by obviating an inter-route return to depot.

MaxVehCap=240; MaxVehDist=50				
Without coordination			With coordination site at Y	
Vehicle	Distance	Route	Distance	Route
1	48	0 6 7 0 10 25 24 22 0	28	0 6 7 10 25 24 22 0
2	48	0 3 8 2 0 15 18 0	50	0 3 8 Y 2 0 15 18 0
3	50	0 4 14 0 1 5 0	50	0 4 14 0 1 5 0
4	38	0 11 13 12 17 0	38	0 11 13 12 17 0
5	25	0 16 9 0	25	0 16 9 Y 0
6	38	0 19 20 0	38	0 19 20 0
7	38	0 21 23 0	38	0 21 23 0
	285		267	

MaxVehCap=245; MaxVehDist=50				
Without coordination				
Vehicle	Distance	Route		
1	50	0 6 7 8 0 4 11 5 13 0		
2	40	0 1 2 3 0 15 18 0		
3	28	0 14 17 0		
4	38	0 10 25 24 22 0		
5	38	0 9 16 12 0		
6	38	0 19 20 0		
7	38	0 21 23 0		
	268			

Table 5.7: Effect of changing maxVehCap from 240 to 245.

Routes altered by coordination are shown in bold type.

When maxVehCap is increased to 245, the vehicle servicing nodes 6 and 7 is able to accommodate the demand for node 8 and the distance increases from 10 to 16, thus preventing that vehicle from also servicing route 0 10 25 24 22 0, which has a distance of 38, without violating the vehicle distance constraint. Any viable combination of route 0 6 7 8 0 with another route passes very close to the depot and there is no benefit in removing the inter-route return to depot. Therefore, the removal of the comparable return to depot in the system with maxVehCap=240, which resulted in a substantial decrease in total distance travelled, is not present when the vehicle capacity is increased.

The other inter-route return to depot is created by a vehicle servicing route 0 1 2 3 0 and any other route. Because no node on route 0 1 2 3 0 is more than two distance units from the depot, the benefit in removing the inter-route return to depot can be no more than 1% and is not considered worthwhile.

It should be noted that an increase in the number of routes and in the total distance is not always accompanied by an increase in the fleet size. Consider, for example, the illustrative structures in Table 5.8 where fleet size decreases in Structure B even though both total distance and number of routes increase. In Structure A, four routes have lengths greater than $(0.5 \times \text{maxVehDist})$, requiring a fleet size of at least four; the fifth route is short enough to be served by one of the four allocated vehicles without requiring the fleet size to be increased. In Structure B, three routes have lengths greater than $(0.5 \times \text{maxVehDist})$, so at least three vehicles are required. The remaining routes in Structure B are able to be serviced by those three vehicles without violating the distance constraint, so no extra vehicles are required. Thus, although the routes in Structure B have a greater total distance, the individual route lengths can be combined more efficiently, and a smaller fleet size is required.

	Structure A		Structure B	
	route 1	48	route 1	40
	2	32	2	32
	3	30	3	23
	4	27	4	27
	5	10	5	10
			6	18
total distance		147		150
fleet size		4		3

Table 5.8: Effect of individual route distances on total distance and fleet size ($\text{maxVehDist}=50$).

5.11.2 Coordination by split delivery

When the only form of coordination available is the splitting of demands, the number of routes created may be a minimum, where maxVehCap determines the minimum number of routes created by

$$\text{Minimum no. routes} = \text{ceil}(\text{total demand}/\text{maxVehCap}).$$

The fleet size is determined by how efficiently the routes can be combined to produce vehicle distances without violating the vehicle distance constraint. Consequently, altering the value of maxVehCap can influence the fleet size by changing the number of routes and hence the way routes are combined for vehicles. In the following example, maxVehDist has value 50 and only the value of maxVehCap is varied; results are displayed in Tables 5.9 and 5.10, and in Figure 5.37. When maxVehCap is increased from 240 to 250, there is no change in the minimum number of routes that can be created and no effective change in the resulting route structure; only the value of the split sizes is affected. The presence of large demands permits such buffering to occur because the variety of ways in which larger demands can be split makes systems that include them less sensitive.

When the value of maxVehCap is reduced from 240 to 235, the minimum number of routes changes from 8 to 9 and the subsequent change in the way the distances of these routes combine produces a larger fleet.

MaxVehCap=240, 250; MaxVehDist=50		
Vehicle	Dist	Route
blue	44	0 6 7 3 0 2' 10 25 22' 0
red	34	0 5 11 13 4 12' 0
pink	47	0 15' 14 0 2" 8 9 16 1 0
green	34	0 12" 17 18 15" 0
gold	40	0 15'" 19 20 21' 0
peach	40	0 21" 23 24 22" 0
	239	
MaxVehCap=235; MaxVehDist=50		
Vehicle	Dist	Route
red	34	0 4 5 11 13 12' 0
green	34	0 12" 17 18 15' 0
pink	48	0 15" 14 0 24" 25 22 1 0
gold	40	0 15'" 19 20 21' 0
peach	40	0 21" 23 24' 0
blue	38	0 10 9 8 3 2 0
black	28	0 16 6 0 7 0
	262	

Table 5.9: Effect of changing maxVehCap from 235 to 250.

Primes indicate split deliveries.

maxVehCap	total dist	fleet
235	262	7
240	239	6
250	239	6

Table 5.10: Effect of maxVehCap on total distance and fleet size for the standard problem in Appendix D.




Figure 5.37: Different values of maxVehCap may produce different schemes. These two scenarios differ only in the value of maxVehCap. Splitting occurs at nodes 2, 12, 15, 21 and 22 (upper), and at nodes 12, 15, 21 and 24 (lower). Colours indicate vehicles.

5.11.3 Mixed systems with common demand

It has already been shown (§5.10.2) that large demands buffer the coordination process during splitting. In an attempt to determine the relationship between the value of maxVehCap and the value of non-buffering demands, topologically distant nodes are assigned large demands while all other nodes are assigned the same or almost the same smaller demand (denoted 'CD' for 'common demand').

Since coordination flourishes in an atmosphere of route interaction and surplus capacities, the potential for coordination is enhanced when the value of CD is small enough to allow the creation of many routes and yet large enough so that integer multiples of CD do not reach the value of maxVehCap. The latter condition allows surplus carrying capacity to exist and thus produces one of the environmental requirements for coordination.

During coordination, the surplus carrying capacity of one vehicle is used to carry part of the demand of a node that would not be on the route currently serviced by that vehicle if coordination were not applied. That demand may be shared among several other vehicles, all of which must approach a coordination site without excessively increasing the total distance travelled. In order for this to happen, several vehicles need to have enough surplus carrying capacity to accommodate the demand. Consequently, there is a relationship between the values of maxVehCap and CD, and this relationship is determined partly by the network structure.

Let

$$R = \text{ceil}(CD/(\text{maxVehCap mod } CD)).$$

The value of R is an indication of the number of routes required to carry the split delivery of one node of demand CD. For coordination to be effective, R+1 routes need to interact: R for coordination and one to be improved. The minimum number of routes in a feasible solution, min Routes, is determined as above.

Typical results are shown in Figure 5.38 for a system with maxVehCap=235. When CD has value 70 or 100, three routes are required to accommodate one demand of value CD; for other values of CD in the range 50-100, two vehicles interact to carry a demand of value CD among their combined surplus capacities. When CD has value 90, no one vehicle is able to accommodate a larger demand at the same time as a demand of CD without applying coordination. Consequently, the application of coordination has the potential for considerable improvement. For a system with maxVehCap=235, the greatest improvements by coordination can be expected for CD values of approximately 80-90, in the expectation that with at least 10 or 11 routes, two will be close enough to interact and be able to

accommodate the demand of one node of a third route. When CD has value 110, 8 routes are required to interact to improve a ninth route; with a minimum of 13 routes, this is unlikely, hence in this study CD values of 110 are not investigated.




Figure 5.38: Effect of CD on R and minimum number of routes for a system with maxVehCap=235.

Investigations at different grid densities (see later) confirm this pattern for all types of coordination regardless of depot placement and grid topology.

5.12 Problem design

In the studies of the fundamental types of coordination in the sections immediately following (§5.13 - §5.15), problems are generated with the following default values, unless stated otherwise:

- single depot pure delivery system (as in §5.9.1)
- maxVehCap=200
- maxVehDist=50
- objective is to reduce the total distance travelled and the fleet size.

Default values have been chosen to allow feasible, single-noded routes and to generate systems with enough vehicles and routes to demonstrate the effect under

consideration but not so many as to cause obfuscation. Customer demand is generated from a normal random distribution to produce an integer in the range $[0, \text{maxVehCap}]$ with all demands known in advance; the next chapter deals with dynamic demands and node locations. Coordination sites are limited to vertices; the number of nodes is determined arbitrarily in advance. Grids are generated at each of the sparse, medium and dense levels. The examples shown below have been chosen to demonstrate the type of coordination under discussion. Where splitting is allowed, the minimum acceptable customer demand is assumed to have value 1.

5.13 Load swapping

The exchange of payloads is predicated on the ability of the vehicles involved to accommodate the total payload and their concurrent presence at a site that is accessible to each vehicle without violation of distance and capacity constraints. It can be an effective method for dealing with a scenario in which a vehicle has surplus capacity but limited distance, and another has surplus distance but limited capacity.

Figure 5.39 shows three vehicles meeting at a coordination site (arrowed); activity at the coordination site is given in Table 5.11. In this example, the red and blue vehicles are unable to service all the nodes they can reach without violating capacity constraints; the blue vehicle is unable to service all the nodes whose demands it can accommodate without violating distance constraints. The situation is resolved by the creation of a coordination site, x , at which the three vehicles meet and swap loads, as shown in Table 5.11. The rural bus company of Problem 4 (see Chapter 1) utilises a similar scheme in which passengers are taken to a cordination site (one of the schools) at which they rearrange their carriage to their final destinations.




Figure 5.39: Three-vehicle swapping.

Vehicle	Route	At x
red	0 a x b c 0	transfer: 35 to blue, 15 to green collect: 60 from blue; 10 from green
blue	0 f g x h 0	transfer: 60 to red collect: 35 from red
green	0 d x e 0	transfer: 10 to red collect: 15 from red

Table 5.11: Activity at the coordination site in three-way swapping.

The concept of load swapping in the rectilinear grid network can be investigated by considering the swapping between only two vehicles, since swapping between more than two vehicles is a natural extension of this investigation but does not add substantially to the underlying principle. Each of the two affected routes can have many paths before and after coordination since a rectilinear grid can offer a multitude of paths between any two vertices. Each vertex that is not part of a no-choice-string represents an intersection which might be part of an equal length path between two nodes. Even the shortest route between two nodes can be made up of many equally valid paths, as in Figure 5.1.

If the routes intersect, neither vehicle might have to deviate from its uncoordinated path, as with the green and red routes in Figure 5.39. Intersection in the rectilinear grid can occur at a vertex or at one or more strings of edges, as in Figure 5.40. Where temporal flexibility permits, the intersection region (as in Figure 5.40) is a natural coordination site. Each intersection site must be checked for coordination.




Figure 5.40: Intersection at points and regions.

If the routes do not intersect, one or both vehicle(s) must deviate from the uncoordinated route(s) to effect coordination. When selecting sites for swapping to

occur, the possible deviation points need to be determined. Where only one vehicle deviates, the points are those vertices that lie on one of its shortest paths and are closest to the second route. Where both vehicles deviate, every vertex within a coordination hull consisting of the two routes, in all their possible configurations, becomes a possible coordination site.

The two examples below show load swapping without detour and with one vehicle detouring. Although there is a subsequent reduction in path duplication, the coordination is driven by the imperative to swap loads rather than to reduce path overlap.

Consider the routes shown in Figure 5.41 (colours represent vehicles; nodes are labelled with letters; demand is shown next to each node; unaffected nodes have been removed). It would be advantageous to have nodes *b* and *j* serviced by the same vehicle since they are close together and distant from other nodes. Unfortunately, this would mean combining nodes *g* and *h*, which is not feasible since the magnitude of their demands would violate the vehicle capacity constraint.




Figure 5.41: Routes prior to swapping.

Vehicle	Route	Load	Distance
blue	0 h j 0	35	42
magenta	0 g b 0	87	44

Table 5.12: Prior to swapping.

The situation is resolved by introducing a coordination point, *x*, as shown in Figure 5.42, and using it as a venue to swap loads between the two vehicles. Both

vehicles travel to x , the magenta vehicle delivering the load for node g en route. At x , the magenta vehicle has free space of $(200-68=132)$ units and the blue vehicle still has the demands for nodes h (120) and j (15). The demand for j is transferred to the magenta vehicle which then travels to nodes j and b while the blue vehicle travels to node h ; both then return to the depot. The total distance of 74 units represents a reduction of 14%. Load swapping does not result in a reduction in fleet size in this case.

■




Figure 5.42: Routes with load swapping.

Vehicle	Route	Load in to x	Load out of x	Distance
blue	$0\ x\ h\ 0$	135	120	30
magenta	$0\ g\ x\ j\ b\ 0$	187	83	44

Table 5.13: Load swapping at coordination site, x .

Consider the situation in Figure 5.43 in which nine nodes on two routes are serviced by two vehicles. Since the routes overlap slightly and approach each other for a short stretch, there is a possibility that coordination in the form of load swapping could occur. This is advantageous only if there is a resulting distance reduction, since the number of routes cannot change when loads are swapped.

Vehicle	Route	Load	Distance
red	$0\ a\ b\ j\ c\ 0$	200	36
green	$0\ f\ g\ h\ i\ t\ 0$	195	48

Table 5.14: Prior to detour and swapping.




Figure 5.43: Routes prior to detour and swapping.

The demand for node t is unable to be accommodated by the red vehicle since this would violate the capacity constraint, so the green vehicle traverses a short path to service t . Node j would be served better by the green vehicle since the green route passes directly through j on its path to node i , but nodes t , j and i cannot be carried by the same vehicle since this would exceed the maximum carrying capacity of any vehicle.

The situation is resolved by establishing a coordination point, x , at a location on or near to each route at a point before node t so that j and i can be serviced by one vehicle. One solution is shown in Figure 5.44 where the red vehicle detours to x which lies on the green route; details are given in Table 5.15. On leaving the depot, the vehicles carry the demands for the nodes allocated to them in Table 5.14, along the routes indicated in Table 5.15 until they reach the coordination point, x , where the demands for nodes t and j are swapped. The red vehicle no longer has to travel to j and is able to take a short path between t and c and thence to the depot; the green vehicle no longer makes the short journey to t . The total distance is reduced from 84 units to 76 units, a reduction of 14%. The elimination of path overlap is a consequence of the load swapping.



Load swapping seeks to reduce the total distance travelled rather than the fleet size. It relies on the abilities of multiple vehicles to access a coordination site at the same time or within an acceptable time span. Usually, suitable coordination sites exist at or near regions of path overlap which, as a result of load swapping, are reduced or eliminated.




Figure 5.44: Routes with detour and load swapping.

Vehicle	Route	Load into x	Load out of x	Distance
red	$0\ a\ b\ x\ t\ c\ 0$	160	185	30
green	$0\ f\ x\ g\ h\ j\ i\ 0$	188	163	46

Table 5.15: Detour and load swapping.

5.14 Deferral of return to depot

Multi-routed vehicles return to the depot between routes. It is sometimes possible to collect some of or all of a load for a later route while still servicing the earlier route. This involves a second vehicle carrying some of the later route's load and depositing it at a location that the multi-routed vehicle can reach. In some cases, it may be possible to remove the multi-routed vehicle's requirement to return to the depot.

Consider the system shown in Table 5.16 and Figure 5.45 (colours represent vehicles; nodes are identified by letters; numbers represent demands; nodes not involved in subsequent coordination have been removed from the figures and tables) in which two vehicles traverse three routes. The red vehicle is multi-routed: it services the nodes on route 1 before returning to the depot to collect the loads for the nodes on route 2. The distance constraint prevents the green vehicle from servicing route 2, whereas the loading constraint prevents the red vehicle from

combining the nodes on route 2 with the nodes on route 1 into one non-coordinated route.




Figure 5.45: Multi-routed vehicle (red) prior to coordination.

Route #	Route	Load	Distance	Vehicle	Route(s)	Distance
1	$0\ t\ k\ r\ 0$	98	26	red	1 2	50
2	$0\ b\ a\ 0$	93	24	green	3	41
3	$0\ p\ m\ f\ g\ 0$	106	41			

Table 5.16: Prior to coordination.

Results of coordination are shown in Figure 5.46 and Table 5.17. The green vehicle travels close enough to route 2 to detour towards it and drop off demand for nodes a and b at coordination point, x . The increase in distance for the green vehicle is more than compensated for by the reduction in distance for the red vehicle which no longer is compelled to return to the depot to collect demand for nodes a and b . Distance has been reduced by 11%; fleet size is unchanged. ■

The coordination site, x , used in this example is not the only suitable vertex; three other vertices produce the same distance reduction. A discussion of the alternative coordination sites appears in §5.14.




Figure 5.46: No inter-route return to depot for red vehicle.

Vehicle	Route number	Route	Load in to x	Load out of x	Distance
red	1	0 t k r x b a 0	0	93	38
green	2	0 x g e f m p 0	199	106	43

Table 5.17: No inter-route return to depot for red vehicle.

When using type 1 coordination sites to obviate a return to depot for a multi-routed vehicle, the coordination sites are most effective when the depot falls outside a hull formed by nodes served by the vehicle and connecting edges. Two multi-routed vehicles are shown in Figure 5.47 and described in Table 5.18. In Figure 5.47, green lines indicate the hull formed by nodes served by the green vehicle and include edges on the shortest paths between those nodes. Similarly, black lines indicate the hull for the black vehicle. Since the depot lies outside the green hull, the establishment of a type 1 coordination site on the hull will reduce the total distance travelled by the green vehicle. There will be no benefit in establishing such a site for the black vehicle since the depot already lies on the hull.

■

Vehicle	Route
green	0 a b 0 f 0
black	0 c d 0 g 0

Table 5.18: Positioning a type 1 site.




Figure 5.47: Positioning a type 1 site.

Deferring or obviating an inter-route return to depot for a multi-routed vehicle has the potential to reduce both total distance and fleet size. Where path overlap is not present before coordination is applied, the establishment of type 2 and 3 sites may facilitate coordination by providing a site to which one of the vehicles can sensibly detour. If this is inappropriate because of the increase in distance imposed by the detour, the establishment of a type 1 site may be possible, where one vehicle adds extra distance to its roster in order that another (multi-routed) vehicle is able to reduce its total distance by travelling only to the coordination site and not to the depot for replenishment of payload. In such a case, the type 1 site should be located as far from the depot as practicable.

5.15 Split delivery

5.15.1 Reducing path duplication

In the rectilinear grid, vehicles are restricted to travelling along pre-defined edges; they are not free to roam as in the Euclidean plane. As a consequence, several paths may overlap, creating possibly degenerate stretches. Reducing the incidence of these stretches, where possible, can result in a reduction of the total distance travelled.

Consider the situation shown in Figure 5.48 in which ten nodes are serviced by four vehicles. There are three main stretches of duplicated paths: depot to node *a*, depot to node *e*, and depot to a point near node *i*. Smaller stretches involve the unavoidable two-way traversal of cul-de-sacs.




Figure 5.48: Before coordination. Several stretches of overlap.

Vehicle	Route	Load	Distance
red	0 a h c g 0	187	34
blue	0 b e j i 0	177	44
magenta	0 d 0	200	34
green	0 f 0	11	30

Table 5.19: Before coordination.

The original scheme is dominated by the demand of maxVehCap for node d , which is delivered intact and appropriates a vehicle which cannot then be used for any other node. When coordination is applied, this large load can be spread among more than one vehicle. In seeking to remove areas of overlap, node i is grouped with nodes g , c and h ; node f is grouped with node a . The load for node d is delivered by two vehicles, resulting in a distance reduction of 19% and a fleet reduction of 25%. Stretches of superposition still exist but they are not as extensive. The longest such stretch (from the depot to node d) is covered three times, with the blue vehicle travelling in one direction and the magenta vehicle travelling in both directions. Before coordination, this stretch had been traversed four times, twice each by the green and blue vehicles. These results are shown in Figure 5.49 and Table 5.20.

■

Vehicle	Route	Load	Distance
red	0 i g c h 0	178	37
blue	0 a f d 0	200	40
magenta	0 j b e d 0	197	30

Table 5.20: Reduced overlap.

The mitigation of path overlap is a popular theme [RBL96] in vehicle routing. Split delivery can reduce path overlap by the more efficient use of vehicles, such as in the example above where overlap in one region is removed by splitting the demand of a node that is located in a quite different region. The paths available in the rectilinear grid network are more limited in number than their Euclidean counterparts, so it is not surprising that there are regions of path overlap remaining after the application of coordination by split delivery. The presence of such stretches of overlap does not diminish the value of the application of coordination.




Figure 5.49: Reduced overlap.

5.15.2 Fleet reduction

Fleet size is dependent upon the manner in which individual routes are combined so as not to violate distance constraints. Coordination points can exist in order to reduce slightly the distances of a number of routes so that they can be combined more efficiently using fewer vehicles. The concomitant decrease in total distance may be quite small. The prospect of an improvement by coordination is signalled by the total distance being only slightly longer, by ε , than the value required for a reduction in the fleet size, and the number of routes, NumRoutes, being greater than the minimum required:

$$\frac{\text{total distance}}{\text{maxVehDist}} = \min(\text{FleetSize}) - 1 + \varepsilon, \quad 0 < \varepsilon < 1$$

$$\text{and} \quad \text{NumRoutes} > \min(\text{NumRoutes})$$

$$\text{where} \quad \min(\text{FleetSize}) = \text{ceil}\left(\frac{\text{total distance}}{\text{maxVehDist}}\right)$$

$$\text{and} \quad \min(\text{NumRoutes}) = \text{ceil}\left(\frac{\text{total load}}{\text{maxVehCap}}\right)$$

That is, ε is the non-integer part of $\frac{\text{total distance}}{\text{maxVehDist}}$. The magnitude of ε that enables this form of coordination to exist may be unique to each grid network structure. Investigation of the critical value for ε is an area for future research.

Consider the situation shown in Figure 5.50 in which four vehicles service ten nodes on six routes (colours represent vehicles, nodes are labelled a to i , demands are shown beside each node). The vehicle capacity is 200 units with maximum distance 50 units. Without coordination, the routes would be those in Table 5.21 with a fleet size of 4. Nodes c and d could not be carried by the same vehicle because the distance created by returning to the depot would have violated the distance constraint. A return to the depot is necessary because the sum of the demands for nodes c and d violate the loading constraint, making the inclusion of nodes c and d on the same route an infeasible solution.




Figure 5.50: Four vehicles service the routes before coordination.

The total distance of 152 units must be carried by at least four vehicles, which is the fleet size here, but is only two units over the limit for a fleet size of three. If routes 3 and 4 were slightly shorter, they could be combined onto the

same vehicle. Since the total load of 901 could be carried by five routes instead of the six found here, there is the possibility of using coordination to shorten at least one of routes 3 and 4 or to combine routes 3 and 4 into one route.

Route #	Load	Route	Dist	Vehicle	Route(s)	Dist
1	82	<i>0 f i 0</i>	30	red	1 2	50
2	85	<i>0 a e b 0</i>	20	green	5 6	50
3	62	<i>0 c 0</i>	28	magenta	3	28
4	9	<i>0 d 0</i>	24	black	4	24
5	59	<i>0 g 0</i>	30			
6	44	<i>0 h 0</i>	20			

Table 5.21: Without coordination.

A coordination point is created at x along the overlap in paths from the depot to each of nodes c and d ; the updated route structure is shown in Table 5.22. The actual placement of x along the region of overlap occurs so as to minimise the distance ($c-x-d$) since each of the other routes utilises x as a type 3 site without incurring additional distance.

Route 1 carries an extra 32 units from the depot; it delivers 162 to node c and collects 37 from x to deliver 69 to node d . Route 2 drops 18 at x before continuing on to nodes f and i ; route 3 drops the remaining 19 required for node d at x before continuing on to node g . The load for node d is carried split but recombined at x before delivery, and is therefore delivered intact.

Route #	Route	Load	Dist	Vehicle	Route(s)	Dist
1	<i>0 c x d 0</i>	200	46	black	1	46
2	<i>0 x f i 0</i>	200	30	red	2 4	50
3	<i>0 x g 0</i>	178	30	green	3 5	50
4	<i>0 a e b 0</i>	185	20			
5	<i>0 h 0</i>	144	20			

Table 5.22: Fleet reduction after coordination.

Nodes c and d are now served by the same vehicle on the same route. The other routes involving x are not affected since x lies on the original paths and no detour has been created. There is only a 5% reduction in the distance of the affected routes, but a 25% reduction in fleet size. ■




Figure 5.51: Routes after coordination, showing fleet reduction.

The value of ε is an indication of the amount of adjustment that a routing scheme needs in order to achieve an improvement in the fleet size; the suitability of any routing scheme for fleet reduction by coordination depends also on the grid network structure. The example above shows the use of split deliveries to a type 3 coordination site in order that an inter-route return to depot is rendered unnecessary which, in turn, results in enough of a reduction in total distance that fewer vehicles are required. A possible advantage here is that the demand for each node is delivered intact, reducing the administration required at the customer's delivery point.

5.16 Alternative coordination sites

Sometimes, more than one coordination site is possible. Reasons for this include:

- vehicles have spare distance capacity and are able to reach equally good coordination sites
- equivalent paths may exist because of the nature of the reduced grid
- changing the orientation of one or both routes might allow loads to be dropped and collected in a sequence suitable for coordination.

Each potential coordination site must be evaluated and compared with all other coordination sites so that the site is chosen whose use results in the greatest reduction in distance and/or fleet size. This includes considering each orientation of a route, since both temporal and spatial characteristics are important.

Consider the example in Figure 5.46 in which two vehicles meet so as to avoid an inter-route return to depot for one of them. Although one vehicle slightly increases its travel distance, the overall distance is reduced because the other vehicle no longer is required to go back to the depot to collect loads for nodes on its second route. Since both vehicles did not travel the maximum allowable distance, alternative coordination points may exist. These are shown by the coloured circles in the figures below; dark rims indicate nodes. For each vehicle, any vertex within $\frac{1}{2} \times (\text{spare distance units})$ of its route (or of an equivalent route) could be reached without violating the distance constraint. Since the red vehicle had 12 spare distance units and the green vehicle had 7 spare distance units, any vertex lying within 6 units of the red route (Figure 5.52) and within 3.5 units of the green route (Figure 5.54) is a potential coordination site. Any such vertex whose use did not decrease the total distance travelled would be unacceptable.




Figure 5.52: Vertices acceptable to the red route as coordination sites. Red circles indicate vertices within 6 units of the red route.




Figure 5.53: Vertices acceptable to the green route as coordination sites. Green circles lie within 3.5 units of the green route.

Several nodes are coloured both red and green, shown in Figure 5.54 as yellow circles. Each will be considered as a potential coordination site and rejected if its use does not result in a reduction in total distance.




Figure 5.54: Potential coordination sites.

For each potential site, the new total distance is calculated and the site rejected if the distance is not less than the original (without coordination) distance of 91 units. Successful sites are shown in the figure below with circles colour-coded to indicate their distance savings.




Figure 5.55: Distance reduction values at potential coordination sites.

From Figure 5.55, four vertices produce the maximum distance reduction of 10 units. Any of them would have been acceptable as the coordination site. One of the outcomes is shown in Figure 5.46, the other resulting routes are shown in the figures below. A comparison of the routes is given in Table 5.23. When the

coordination site is moved to either of the best two left-most sites, the green route changes to an equivalent path without affecting the order in which the nodes of route 2 are visited. In Figure 5.45 and Figure 5.55, the green vehicle could have continued along the red route past node *b* before approaching its original route. Recall that each result is one of possibly many equivalent results, depending on the nature of the reduced grid.




Figure 5.56: Green vehicle detours to upper left coordination site (arrowed).




Figure 5.57: Red vehicle detours to lower left coordination site (arrowed).

Often, coordination methods do not exist in isolation – it is possible that the application of one type of coordination creates an environment that is conducive to the application of another type (discussed in more detail in §5.18). Where alternative coordination sites exist, the acceptance of any particular site may have consequences for cascading coordination in that, by accepting the site, the environment for other types of coordination may be prevented or created. Whether this is advantageous or not depends on the individual problem and the resources which are able to be dedicated to it. In this study, it is recognised that where multiple potential coordination sites exist, the site is accepted that results in the greatest immediate benefit to the objective function.




Figure 5.59: Three equivalent paths for the green vehicle.

5.17 Grid density

In considering the interaction between grid density and coordination, three standard layouts are considered, one of each grid density: medium, sparse and dense. Since the medium density grid more closely resembles real-world applications and because it can be expected to provide a middle-of-the-road insight, it is investigated in more detail than the sparse and dense grids, and the results compared.

For each problem, the following defaults are assumed:

- single depot pure delivery system of 25 nodes with all demands known in advance
- maxVehCap=235
- maxVehDist=50
- objective is to reduce the total distance travelled and the fleet size
- no limit on the fleet size in the unimproved system
- fleet increase is acceptable only if the accompanying distance reduction is at least 5%.

Because it will be shown (in §5.18) that splitting tends to dominate other forms of coordination when they are applied together, three scenarios are investigated for each problem – no coordination, splitting only, and coordination other than splitting. In each case, some nodes have high demands of 150, strategically located to make the problems more interesting and located so that they are unlikely to be serviced by the same vehicle, and all other nodes have a common demand (CD).

When CD has value 150, all nodes have the same demand and no buffering exists; results are shown in Table 5.24 and discussed later.

	Medium							
	Sparse		D1				Dense	
	dist	fleet	dist	fleet	dist	fleet	dist	fleet
Unimproved	(558)	(12)	(480)	(10)	(654)	(16)	(462)	(10)
Splitting	13%	0	18%	1	21%	4	20%	2
Other	19%	2	20%	1	23%	4	16%	2

Table 5.24: Reductions in distance and fleet size using splitting and other forms of coordination in grids of different density when all nodes have the same demand (CD=150). Figures in brackets refer to actual values.

5.17.1 Medium grid

Two forms of a 186 edge medium grid are examined, each with the same layout of nodes and differing only in the position of the depot. Depot location D1 is reasonably central on the 'shore' of a lake, while depot location D2 lies on the outer edge of the grid. Several nodes lie on no-choice-strings on the edge of the grid, and the grid is also characterised by an abundance of no-choice-strings within the grid. With 65 intersections most shortest paths between depot and node exhibit multiple redundancy, although some of the edges can be expected to be more heavily travelled than others, particularly the edge of the lake.




Figure 5.60: Medium density grid. Alternative depot locations are labelled D1 and D2.

As well as the lake, another dominant topological feature is the peninsula containing nodes 5 and 11 that juts out into an inlet bordered by nodes 3, 4, 6 and 7. Because of the large distance from nodes 5 and 11 to nodes 6 and 7 on the other side of the inlet, only nodes on the peninsula side of node 11 are available for inclusion with nodes 5 and 11 in any feasible routing scheme.

For each depot the effect of location of the high demand nodes is investigated by considering four different scenarios: all nodes of high demand, two of high demand diametrically opposite, and one with four nodes of high demand at corners of the grid. For all but the first case, the high demand nodes are located in such a manner that they are unlikely to be serviced by the same vehicle. Contiguous locations would be equivalent to a single demand of much larger value; locating the high demand sites only slightly apart would be equivalent to creating a

split delivery situation. For the standard layout, the following nodes were allocated high demands:

- all nodes
- nodes 13 and 23
- nodes 16 and 19
- nodes 13, 16, 18, and 23.

Detailed results are given in Appendix E, and general results are given in Table 5.25.

Vehicles using the depot at D1 and servicing nodes on the right hand side of the medium density grid may be forced to travel around the lake edge, which has the effect of partitioning the nodes into smaller groups depending on where they lie in relation to the lake and on the relative locations of other nodes serviced by the same vehicle. Nodes 22-25 tend to be grouped together because vehicles servicing them tend to traverse the lake in the same direction; nodes 2, 3, 6-9, and 16 tend to be grouped together because they lie in a mass on the off-lake side of the depot. Node 10 lies some distance above the depot and can be combined with either group. Reasons for the groupings can be seen by considering the arrows in Figure 5.61 that indicate the direction of shortest paths to depot D1; any vehicle one unit to either side of a double-headed arrow is forced to travel in one direction if it is solely depot bound.




Figure 5.61: Arrows indicate the direction of shortest paths to the depot D1.

				Nodes 13 and 23				Nodes 16 and 19				Corner nodes			
	CD	min Routes	R	splitting		other		splitting		other		splitting		other	
				dist	fleet	dist	fleet	dist	fleet	dist	fleet	dist	fleet	dist	fleet
Depot location D1	50	7	2	0	0	4	0	3	0	0	0	6	1	9	1
	60	8	2	8	0	6	0	11	1	12	1	1	0	7	0
	70	9	3	0	0	0	0	4	0	6	0	0	0	0	0
	80	10	2	13	1	13	2	15	1	13	1	11	1	11	1
	90	11	2	11	0	13	0	16	0	17	0	5	0	9	0
	100	12	3	9	0	9	0	0	0	7	0	8	-1	8	-1
Depot location D2	50	7	2	2	1	4	1	6	0	5	0	2	0	5	0
	60	8	2	7	1	7	1	7	0	7	0	3	0	3	0
	70	9	3	6	1	6	1	1	0	4	0	1	0	2	0
	80	10	2	10	0	12	0	8	1	11	1	12	2	13	2
	90	11	2	13	0	14	0	12	1	15	2	12	1	13	2
	100	12	3	11	1	10	1	4	0	4	0	9	1	9	1

Table 5.25: Reductions in distance (%) and fleet size using splitting and other forms of coordination in the medium grid. Slight increases in fleet size are tolerated where they accompany a distance decrease > 5%.

min Routes = $\text{ceil}(\text{total demand} / \text{maxVehCap})$,
 $R = \text{ceil}(\text{CD} / (\text{maxVehCap} \bmod \text{CD}))$ (see §5.11.3).

In contrast, when the depot is located at D2 the nodes are not so clearly partitioned. With most nodes a greater distance from the depot, there are more possible shortest paths and the effect of the lake is less pronounced. This is consistent with the real-world situation in which distant geographic features have diminished influence on local decisions. Nodes 16 and 9 are no longer grouped with nodes 2, 3 and 6-8 since the depot is now located within their midst, and nodes 8 and 25 are no longer exclusive to any particular group.




Figure 5.62: Arrows indicate the direction of shortest paths to the depot D2.

Because of the constraints of capacity and distance, the minimum number of routes may not always be achieved. For example, the D2 system of ($CD=90$, min routes=11; high demand nodes: 13, 16, 18, 23) has 12 routes when improved both by split deliveries and by other means even though these result in distance reductions of 12% and 13% respectively. This result is probably caused by the vehicle capacity constraint that prevents demand(16) and either demand(9) or demand(7) being carried by the same vehicle, causing node 16 to be serviced by a single-noded vehicle.

When the depot is located at an extremity, node-depot distances involving nodes at opposite extremes will be much greater than when the depot is more centrally located. As such distances approach and exceed the value of $\text{maxVehDist}/2$, the fleet size may be increased since each route of distance greater than or equal to $\text{maxVehDist}/2$ is serviced by a separate vehicle. This is evidenced

by the greater fleet size and larger distances when the depot is located at D2 compared with those with depot at D1, as shown in Table 5.24.

When the depot is located near an edge of the grid, route density may be increased near the depot because of the greater limitations on access to it; similarly, if a depot is removed from an area near a constraining geographical feature, route density may be decreased in that area. When the depot is shifted from D1 (on the lake edge) to D2 (on the grid edge), route density increases in the region near D2 and drops in the region near D1.




Figure 5.63: Limited access to a depot can increase the route density near it.

As expected, coordination both by split delivery and by other means (shown in Table 5.25 and Figure 5.64) produces the greatest reduction in total distance at values of CD around 80-90 and is less effective for values of CD approximately 50, 70 and 100, with the trends unaffected by depot location.




Figure 5.64: Variation of distance reduction with CD and depot location. For each depot location (D1 and D2), CD values range from 50 to 100.

Caution must be exercised when applying sequential coordination to ensure that deliveries and pickups (denoted by 'del' and 'col' respectively below) are made in the correct order (a transferred load is delivered to the coordination site before it can be collected from there) and that the final route can be serviced in the allotted time. Both types of coordination can be found in the same routing scheme, as shown in Table 5.26 in which letters refer to coordination sites. There does not appear to be any pattern to the occurrence of sequential and simultaneous coordination.




Figure 5.65: Sequential and simultaneous coordination. Intervening deliveries to other nodes have been omitted.

CD=80 (Mostly sequential)	CD=70 (Simultaneous)	CD=60 (Mostly simultaneous)
0 5 11 T ^{del} 0	0 X ^{del} 5 11 4 0	0 4 5 11 X ^{del} 0
0 13 12 T _{col} 4 0	0 13 12 0	0 13 X _{col} 12 17 S ^{del} 0
0 17 14 W ^{del} 0	0 Y ^{del} 15 14 17 0	0 15 14 S _{col} 18 19 0
0 15 W _{col} 18 19 X ^{del} 0	0 18 19 Y _{col} X _{col} 1 0	0 20 21 V ^{del} 22 0
0 1 X _{col} 20 21 Y ^{del} 0	0 20 21 22 0	0 23 24 V _{col} 25 0
0 23 24 Y ^{del} 0	0 23 24 0	0 10 9 8 W ^{del} 0
0 22 Y _{col} 25 10 Z ^{del} 0	0 22 25 10 0	0 3 W _{col} 6 7 16 0
0 8 Z _{col} 9 16 V ^{del} 0	0 8 9 16 0	0 2 0
0 3 V _{col} 6 7 0	0 7 6 3 0	
0 2 0	0 2 0	

Table 5.26: Routing schemes for the system with depot at D1 and high demand nodes 13 and 23, differing only in the value of CD. Routes indicated may be traversed in the reverse order to accommodate delivery and collection timetables.

The following observations are made:

1. When the depot is at an edge of the grid, the influence of the value of maxVehDist is more pronounced.
2. When the depot is central, the value of maxVehCap is a greater influence than the value of maxVehDist.
3. Nodes that are geographically close yet far from the depot tend to lie on the same route.
4. The benefits from splitting the deliveries of nodes close to the depot are not as great as those far from the depot, so where there is choice, remove distant nodes from routes via coordination.
5. Splitting can be emulated in non-split environments by creating coordination sites co-located with, or adjacent to, nodes that would have received split deliveries if the rules had permitted.
6. Some coordination schemes are sequential, some simultaneous, some a mix.
7. Coordination is more effective where there are fewer routes required to interact than when there are more routes present. For example, the system (CD=60; min routes=8; R=2) gives greater improvement than (CD=70; min routes=9; R=3).
8. The minimum number of routes need not be achieved in order to produce distance and fleet reductions.
9. Coordination can increase fleet size if there is a concomitant acceptable decrease in distance.

5.17.2 Sparse grid

The sparse grid shown in Figure 5.66 has 143 edges and 39 intersections with the single depot located off-centre. The layout is characterised topologically by long arms of no-choice-strings containing few nodes that spider out from a small central mass. Because the servicing of most nodes requires vehicles to traverse no-choice-strings, the sparse layout contains heavily travelled edges some of which converge in the region around node 12. Consequently this region is a potentially rich area for establishing coordination sites. Because of the expense of travelling between the tendrils, nodes tend to be grouped by their sharing the same no-choice-string, and the number of nodes per route is lower than that for grids of higher density. Tendrils also cause overall distances to be greater than more compact structures with centrally located depots, as shown in Table 5.24. Notwithstanding these comments, observations made in the medium grid are applicable in the sparse grid.

The grid is investigated with high demands at all nodes, and with high demands at nodes 8 and 21. Because the grid has relatively few edges, the effect of the lake bordered by the depot is small. General results are given in Table 5.27; specific results are given in Appendix E. Tendrils tend to limit the effectiveness of splitting as an improvement method, while the more dense regions of the grid lend themselves to use as potential sites by other forms of coordination. The effect of the sparse structure on the potential for coordination is discussed below in a comparison between splitting and other forms of coordination.




Figure 5.66: Sparse grid.

5.17.3 Dense grid

The dense grid shown in Figure 5.67 has 418 edges and 225 intersections, and is characterised topologically by a few small lakes of minor influence and few no-choice-strings, and tactically by an abundance of path choice. Where the medium and sparse grids contain main transport routes, the dense grid has so many alternative paths for a given source and destination that intervening features that would have been obstacles in grids of lesser densities barely impact. Consequently, grouping of nodes is less pronounced in a dense grid.

The wealth of shortest paths between node and depot increases the possibility of other nodes lying on or near such paths and consequently increases the number of nodes that might be serviced by any one vehicle without increasing the distance travelled by that vehicle (subject of course to other constraints). For example, in Figure 5.67, shortest paths between node 11 and the depot pass through nodes 8, 9, 13 and 14 and are only one unit away from nodes 6, 10, 12

and 15. Accordingly, total distances in the dense grid tend to be less than distances in grids of lower density, as shown in Table 5.24.




Figure 5.67: Dense grid.

The dense grid is investigated with high demands at all nodes, and at nodes 9 and 22, located diametrically opposed on opposite sides of the depot and thus unlikely to appear in the same route.

			Sparse grid				Dense grid			
			(Nodes 8 and 21)				(Nodes 9 and 22)			
			splitting		other		splitting		other	
CD	min Routes	R	dist	fleet	dist	fleet	dist	fleet	dist	fleet
50	7	2	1	1	5	2	6	0	3	0
60	8	2	0	0	4	0	4	0	6	0
70	9	3	0	0	0	0	0	0	0	0
80	10	2	8	-1	7	0	9	0	8	-1
90	11	2	8	0	12	1	18	1	9	1
100	12	3	4	0	6	1	4	0	2	0

Table 5.27: Reductions in distance (%) and fleet size using splitting and other forms of coordination in the sparse and dense grids. Slight increases in fleet size are tolerated when they accompany distance reductions >5%.

As expected, the greatest reductions in distance are found with CD values of 80-90 with the dense grid more susceptible to improvement by splitting than by other forms of coordination.

5.17.4 Comparison of splitting with other forms of coordination

In order for splitting deliveries to be an effective method, one or more vehicles must access a customer node, whereas other forms of coordination require only that the vehicles involved access any suitable coordination site, which includes sites coincident with customer nodes. Since forms other than splitting do not require routes to interact near customer nodes, the chances of achieving effective coordination by non-split methods are improved with increasing density of interacting routes. The number of intersections is a measure of the potential for route interaction since routes are proscribed from interacting at locations other than vertices. In this study, the sparse grid is characterised by tendrils of routes with 39 intersections, the medium grid has 65 intersections, and the dense grid has 225 intersections. It can be expected, therefore, that the sparse grid is an environment that is less conducive to coordination by split delivery than the medium grid, with the dense grid providing the most hospitable conditions for splitting.

Since greater distance reduction is achieved where the coordination sites are located away from the depot, it can be expected that intersections close to the depot are less important in the coordination process than other intersections. The region of influence of a depot is probably unique for each layout since the physical layout, the routing structure and the depot location are inter-dependent.

The delivery of goods in split form can be emulated by the establishment of a coordination site coincident with or near to the recipient node. In such cases, the demand is delivered intact to the customer. Such non-split coordination might not be accepted since it may not produce the required reduction in the objective function.

In the medium grid, splitting generally produces a result that is no better than that found using other forms of coordination. In the sparse grid, overall distances tend to be greater and splitting is less effective, in contrast to the dense grid, in which splitting is generally more effective than other forms of coordination. Splitting becomes more effective when several vehicles are able to interact to service the demands of a particular node or group of nodes; other forms do not necessarily require the interaction of multiple vehicles – they can act independently

and may require fewer vehicles. The trends shown in Figure 5.68 relate to systems which differ only in grid density.




Figure 5.68: Grid density and relative effectiveness of coordination methods.

5.18 Order of application of coordination

Three forms of coordination have been investigated which result in a reduction in total distance and/or fleet size. In addition, two further effects have been investigated (fleet reduction by minor changes in route length, and reduction of path duplication) which utilise coordination. As with coordination in the Euclidean scenario, coordination in the rectilinear grid takes several forms depending on the environment and the requirement of the system administrator:

Load splitting

Requires:

- Possible involvement of more than one vehicle
- All vehicles involved in the coordination can access the coordination site(s)
- Total load can be accommodated by the vehicles involved in coordination

Outcome:

- Reduction in fleet size
- Reduction in total distance
- Number of intersections increases

Load swapping

Requires:

- Involvement of more than one vehicle
- Total load can be accommodated by the vehicles involved
- Concurrent presence of vehicles at coordination site

- Vehicle with surplus capacity
- Vehicle with surplus distance

Outcome:

- Reduction in fleet size
- Reduction in total distance
- Number of intersections increases

No inter-route return to depot

Requires:

- Multi-routed vehicle
- Donor and recipient vehicles can access site in correct order

Outcome:

- Fewer routes
- Reduction in total distance

Shortening of routes

Requires:

- Fleet size > minimum fleet size
- Number of routes greater than the minimum required

Outcome:

- Reduction in fleet size

Reduction of path overlap

Requires:

- Path duplication

Outcome:

- Reduced overlap
- Reduction in total distance

These forms of coordination are not mutually exclusive, and more than one type of coordination might be applied to a particular scenario, although the situations which prove potentially rich for coordination generally are not found with equal probability. The application of one form of coordination can result in an environment that is more or less suitable for the application of another form of coordination.

Load splitting allows other types of coordination to exist in situations which would not lend themselves to coordination without splitting, by enabling more efficient use to be made of vehicle carrying capacity. It can be used in conjunction with load swapping and the removal of an inter-route return to depot. Consequently, load splitting should be applied before or at the same time as other forms of coordination.

Any application that reduces the number of multi-routed vehicles may also change the number of inter-route returns to depot, as shown in the generic example of Table 5.28. Primed route identifiers in the Table refer to routes in which the nodes have been re-distributed to form new routes which are subsequently assigned to vehicles in a new scheduling phase. The number of inter-route returns to the depot is decreased whenever there are fewer routes being served by an unchanged number of multi-routed vehicles (Scenario A below), and occasionally when single-routed vehicles assume the payloads of multi-routed vehicles (as in Scenario B). In Scenario A, the nodes comprising route 7 are redistributed to routes 2', 4' and 6' (routes on each of the multi-routed vehicles), resulting in fewer routes being served by the same number of multi-routed vehicles; the number of inter-route returns to the depot decreases by the number of routes no longer served by multi-routed vehicles, in this example, by one. In Scenario B, the nodes on a route (route 2) served by a multi-routed vehicle are redistributed to routes on each of the original multi-routed vehicles, at least one of which is no longer multi-routed.

In Scenario C, the improvement phase results in a multi-routed vehicle becoming single-routed. The same number of routes are serviced by a single-routed vehicle and a reduced number of multi-routed vehicles; the number of inter-route returns to the depot is unchanged. In contrast, Scenario D is an example of the situation where no single-routed vehicle is generated even though the number of routes remains the same; the number of inter-route returns to the depot increases.

Inter-route returns increase in number when the number of routes is unchanged but the number of multi-routed vehicles decreases.



The removal of an inter-route return to depot has the potential to remove paths from the region close to the depot. Since the depot is a compulsory part of each route, this region is necessarily more densely travelled than other areas and has a much higher incidence of path overlap. Consequently, the removal of an inter-route return to the depot might automatically reduce path overlap near the depot but generate overlap in other regions of the grid. Since the removal of path overlap away from the depot has the potential for a much greater reduction in total distance than the removal of path overlap near the depot, there is greater potential benefit to be gained by investigating the removal of inter-route returns to the depot before an attempt is made to reduce path overlap. Consequently, any situation which has the potential to result in changes to the number of multi-routed vehicles

or changes to the number of routes served by multi-routed vehicles should be investigated prior to the reduction of path overlap.

	Vehicle assignments	Result
Original	Veh. A: 1 2 Veh. B: 3 4 Veh. C: 5 6 7	
Scenario A	1 2' 3 4' 5 6'	Decrease
Scenario B	1' 3 4' 5 6 7'	Decrease
Scenario C	1 2' 3' 4' 5 6 7'	No change
Scenario D	1' 2' 3' 4' 5' 6 7	Increase

Table 5.28: Generic examples of possible outcomes after reducing the number of multi-routed vehicles. *Result* refers to the change in the number of inter-route returns to the depot. Numbers are route identifiers.

Unless the acceptance of deliveries is time constrained, splitting is not subject to the same temporal pressures as other forms of coordination which require that the payload involved be delivered to sites no later than it is collected. Since this may introduce waiting times, sequential coordination is difficult to apply where distances and travel times are already large by reason of grid density and geography.

5.19 Quality of initial solution

In this study, there are no inherent limitations imposed on the application of coordination beyond the constraints of vehicle carrying capacity and distance capacity. There are no bounds on the size and number of splits, and the number of coordination sites is not constrained. The critical factor in determining the potential for coordination is the environment created by an initial solution.

Usually, improvement methods are relatively more effective when they are applied to poor solutions, and the conditions that render a solution poor may also be the same conditions that provide a fertile environment for coordination. In the example of Appendix B, shown in Figure 5.69, both initial solutions have environments suitable for coordination (regions of path overlap, paths close to nodes served by other vehicles, spare distance and carrying capacity of vehicles) and each can be improved by coordination. ■

Structures that do not exhibit such environments may not be suitable for coordination. Consider the structure shown in Figure 5.70 in which the depot is located on a central island connected to clusters of nodes by long arms, and where the total demand of each cluster lies between $\text{maxVehCap}/2$ and maxVehCap . Each cluster is, therefore, serviced by a different vehicle, and the inter-cluster arms are so long that there is no advantage in any vehicle travelling between clusters. Coordination is unlikely to improve any feasible solution. ■

The routing scheme itself may prevent coordination. Consider, for example, a scenario, such as that shown in Figure 5.71, in which only one rectilinearly 'convex' route is involved. Single-vehicle coordination is unlikely because the edges are identical, and swapping and transfer will not occur with only one vehicle; there is no inter-route return to the depot because only one route is involved; the route is convex, so there is no path overlap; the fleet size is already at a minimum. There is no environment in such a scenario to encourage improvement by coordination. ■




Figure 5.69: Good (upper, distance 222) and poor (middle, distance 236) solutions can be improved by coordination (lower, distance 198). Some overlapping paths are slightly offset for the sake of clarity.




Figure 5.70: Grid structure prevents coordination.




Figure 5.71: Routing scheme prevents coordination.

For other initial solutions, the routing scheme and structure may combine to create a situation in which only one type of coordination may be appropriate. A scheme with no path overlap, that has no multi-routed vehicles, with a minimum fleet size, and where there is no surplus capacity, can be improved only by splitting, if that is appropriate.

As with other improvement methods, in addition to the environment created by the initial solution, the value of any potential improvement may be a consideration in deciding whether to apply coordination in real-world situations. It has already been shown that coordination can improve solutions substantially but in cases where the initial solution is already reasonable and where workforce and management are comfortable with their current practice, it may be that the extra effort involved in establishing, securing and maintaining coordination sites that are not co-located with customer nodes renders the application of any improvement method undesirable.

5.20 Appraisal of rectilinear coordination

Hall [H01] describes all transportation systems as having the same essential elements of vehicles, guideways and terminals, operating under some control policy. He defines guideways as stationary resources that define feasible paths of travel and provide the physical infrastructure to support vehicles. In the standard Euclidean system, potential guideways consist of the entire Euclidean plane within the coordination hull since movement is unrestricted. In a grid network, the guideways are restricted to those paths connecting nodes to other nodes and to intersections. The grid network coordination hull is not restricted to the footprint of the Euclidean coordination hull of an equivalent Euclidean system: the absence of some edges connecting nodes in the grid forces the grid network coordination hull to include edges whose Euclidean shadow lies outside the Euclidean coordination hull. This is demonstrated in Figure 5.72 where the nodes of a grid network system and a Euclidean system coincide; the dotted line in the grid network indicates a bay. The presence of bays, lakes and streams in the grid forces the grid network coordination hull to adopt a profile that includes concavities.

■




Figure 5.72: Coordination hulls of grid networks (blue) and Euclidean systems (red) do not necessarily coincide.

The multitude of possible paths for every route order and orientation combination means that coordination in the rectilinear grid has been investigated in a simulation environment with the form of coordination decided within the following broad guidelines:

- load swapping takes place before load splitting
- removal of inter-route return to depot occurs before reduction of path overlap
- where one form of coordination has been used successfully, the result is investigated for suitability of other types of coordination.

Generally, load splitting is the most widely used form of coordination since it allows other types of coordination to exist in situations which would not lend themselves to coordination without splitting.

Coordination in the rectilinear grid network is an effective improvement method to reduce the total distance travelled and/or the fleet size. Of 300 ten-node problem instances investigated, 243 (81%) were improved by at least one form of coordination with an average distance reduction of 10% and fleet size reduction of 25%. As the number of nodes was increased to 30, up to 90% of problems were successfully improved by coordination, with similar reductions in distance and fleet size, but computation time increased dramatically. The rate-determining step is the calculation of inter-nodal distances since not only does the number of internodal distances increase with the number of nodes but also the reduced grid becomes more complex with more intersections, each of which slows the distance calculations. When the depot is located at an extreme, the number of 25-node problems that can be improved using coordination approaches 100%, due to the increase in route interaction. The incidence of fleet improvement for medium density grid problems with the depot at an extreme is 50% compared with approximately 30% for problems with a more centrally located depot. In the sparse grid, 75% of problems were improved using coordination to reduce the distance and fleet size; in the dense grid, over 80% of problems resulted in distance reduction and 50% resulted in fleet reduction.

Splitting becomes more effective as grid density increases, whereas other forms of coordination are more effective in regions of low grid density. Experimental results in sparse, medium and dense grids confirm the importance of route interactions in the potential for improvement by all forms of coordination. The presence of demands that are much larger than others has a de-sensitising effect on the splitting process by allowing a wider variation in the size of splits.

In contrast with the Euclidean case, in which nearness of paths and nodes was defined by the near radius value, paths in the rectilinear case are deemed close enough for coordination if the use of a particular site produces a suitable result. The use of surplus tables aids greatly in determining the potential for coordination by indicating the distances that individual vehicles are able to detour and the amount of extra payload they can carry.

Load splitting, load swapping and removing the return to depot for multi-routed vehicles is present in both Euclidean and rectilinear scenarios. Because of structural limitations of the rectilinear grid, the interim metastable state found in the Euclidean scene is absent in the grid network. The establishment of a coordination site, which can be loaded with payload equivalent to the value of

maxVehCap by each vehicle using it as a type 1 site, assumes a high profile in the rectilinear grid. Limitations created by the grid structure enhance the value of such a site.

Reduction of the grid to remove many unnecessary edges greatly improves the speed and efficiency of investigations. Since multiple paths of equal value in the original grid are reduced to one feasible path in the final grid, the actual paths travelled can be considered as representatives of a larger phase space, each component of which is comparable.

Coordination requires no special assumptions other than the acceptable application of the method; the form of route construction is irrelevant to the improvement phase of coordination beyond providing fertile environments in which coordination may occur. Even explicit knowledge of the demands is unnecessary, as will be shown in the next chapter.

CHAPTER 6

AN EXTENSION - DYNAMIC COORDINATION IN THE RECTILINEAR GRID NETWORK

6.1 Introduction

Investigation of coordination thus far has taken place in completely deterministic environments in which all details of customer availability, location and demand have been known before any vehicle has left the depot. In practice, this is not always the case; it is not uncommon for customers to change their demands, the times at which their demands become known, or even their very existence, during a planning period.

This chapter investigates coordination where at least some of the node attributes of location, demand and time are unknown at the start of the planning period, but under the assumption that they will become apparent before the end of the planning period. Routes formed using consideration of initial data may be modified as information is revealed. The chapter aims to show that coordination exists in a dynamic environment and that implementation of coordination can allow a system to be responsive to changes in node attributes.

6.1.1 Dynamic vehicle routing methods

The dynamic vehicle routing problem has been described by Lund et al ([LMR96]) as being any vehicle routing problem in which the decision maker must create routes before all necessary information is available and then must modify those decisions as new information is received. Real-time dynamic decision making (see [LH01]) requires a series of interdependent decisions in a continuously changing environment where the key characteristics are:

- the tasks require a series of decisions

- the decisions are interdependent
- the environment changes autonomously and as a result of decisions, and
- decisions are made in real time.

The information required to make a decision has several attributes including value and quality (see [P00]). Information value allows better decisions to be made; information quality depends on accuracy, timeliness and completeness. For example, if a dispatcher knows with accuracy 1% of customer demands at the start of a work period, the information has quality with little value. If the dispatcher knows that customers will require a particular commodity which the competition will be unable to supply, the information is valuable but has little quality.

Vehicle routing problems can be characterised by the assumptions about both the information content and the time the information becomes available to a dispatcher. Kleywegt et al ([KNS97]) describe deterministic problems as those in which all relevant data are assumed to be known before any decisions are made, stochastic problems as those in which some data are assumed to be unknown (with or without knowledge of the relevant probability distributions), static problems as those in which decisions are made only once with available information, and dynamic problems as those in which a sequence of decisions is made over time using information as it becomes available. Dynamic vehicle routing therefore deals with problems in which information is revealed as the work cycle progresses, and initial route construction occurs with incomplete information (see [Ps95]). Real-world applications include the delivery of petroleum products (e.g. see [B83]), maritime services (e.g. [PT93]), the airline industry (e.g. [BH93]) and postal services (e.g. [PSN88]).

Stochastic and dynamic problems often have poorly defined objective functions, because the models are difficult to define ([DP93]). In a stochastic program, decisions are made before the uncertain quantities are observed ([C98]) and recourse (corrective) actions take place when the uncertain quantities become available, so it is not possible to assume that vehicle routes can be followed as planned without violation of constraints. Such recourse actions include the vehicle returning to the depot to replenish stock and supplying the new customer, or route modification.

In an effort to understand the type of policies that perform well and the value of information about the future, [KNS97] consider a dynamic and stochastic resource minimisation problem where shipments arrive at a depot randomly over time and all deliveries to customers have to be completed before a given deadline, although vehicles are not required to return to the depot by a given time (as they

are in this study). The shipments to be performed at a particular time are determined at a given set of decision points. Since each shipment is a full truckload and all customer nodes are close enough to the depot for vehicles to service them and return before the next decision point, the model involves scheduling without routing. Vehicles are rented from a fleet of unlimited size, with the rental fee covering the entire period, regardless of when the rental occurs. The objective function to minimise is total number of vehicles used. They investigate the situations where there is no knowledge of future shipments, and where the probability distribution of the number of shipments arriving is known, and show that routing and scheduling is more efficient when the future demands can be predicted with more certainty. Their policies are tested on a worst-case scenario (predicated on the assumption that no new shipments will occur) even though the worst-case performance may have little or no bearing on how a policy will perform on average. When there is no knowledge of future shipments, they allocate vehicles based on the average number of shipments remaining. Resource allocation is improved where the arrival distribution has a small or moderate coefficient of variation. When the coefficient of variation is high, the number of vehicles used is reduced by postponing deliveries from high-use periods to low-use periods; when the coefficient of variation is small, it is better to make deliveries as soon as the shipments arrive at the depot. When the arrival of shipments has a known distribution, vehicles can be assigned more efficiently.

Although dynamic vehicle routing has been studied since the 1970s ([WC77], [D78], [S78]), developments in information technology and rapid communications have increased the need for dynamic decision making in the presence of uncertainty by increasing the real-world expectations both of customers and suppliers. Improvements in communications enable dispatchers to respond rapidly to changes in customer requirements and to environmental factors affecting the transport of goods; satellites and global positioning systems in vehicles allow dispatchers to monitor fleets of thousands of vehicles [DP93]. Powell et al [PSC00] describes dynamic routing and scheduling problems as requiring rapid response times and the ability to output multiple recommendations for a particular driver.

Because quick response times to problems with large numbers of variables are required, heuristic solutions are sought in real-time environments. The three main categories of approach [IGP00] are:

- adaptation of static programs, including Tabu Search
- stochastic methods, and
- neural networks [SP95] and genetic programming.

The adaptation of static algorithms is based on the concept of a rolling horizon, where static problems are repeatedly solved over events found within an horizon of length δt that extends from the current time t to $t+\delta t$. For very small δt , strategies can produce near-optimal solutions (see [BDF83], [D95], [P88]). For large δt , the solution is typically weaker, unless a fast and powerful solution procedure is provided [T89]. This approach ignores the potential benefits of considering the probabilistic aspects of the problems and trying to forecast the future. Either the problem is modified every time an update occurs or rapid insertion techniques are used. Although insertion is easy to implement, solutions created by consecutive insertions may require improvement by local search methods such as 2-opt, whereas complete re-ordering of routes may have lead to improved solutions. There is a trade-off between speed of response and quality of solution.

Savelsbergh [SS98] investigates dynamic routing with real-life data in a pickup and delivery mode, and inserts new requests into feasible routes where possible or creates new routes where no feasible route exists. Post-insertion improvement is by inter-route transfer or inter-route exchange.

Stochastic methods (see [P88], [BV91], [BV93]) aim to react to an event in such a way as to minimise the undesirable impact on the whole system while ensuring a good quality of service to customers disturbed by the events, but problems are plagued by a rapid increase in state space with problem size, and assumptions must be made to make the problems more tractable.

The stochastic programming approach [DLT89] involves the construction of a preplanned set of routes that is not reoptimised. A survey on stochastic programming can be found in [B97]. Neural networking [SP95] aims to replicate the dispatcher's decision making process.

Scenario planning is used in stochastic problems [O99] to describe a number of plausible future states that a system may develop, and thus enable a planner to develop solutions that perform well under all scenarios. Aspects of this method will be used in this study of dynamic coordination to develop a set of responses to possible future scenarios, so that an adaptive strategy can be constructed that will allow routes to be modified as information becomes available. The modification may range from an insertion to a complete overhaul but, by pre-planning the strategies, the response can be achieved rapidly.

6.1.2 Dynamic coordination

There does not appear to be any published material on the application of dynamic coordination in the rectilinear grid network, although several researchers have applied limited coordination to dynamic problems in the Euclidean plane.

Regan et al [RMJ94,95] investigate a limited form of dynamic coordination in a pure pickup (or delivery) scenario in a Euclidean setting by allowing a vehicle to be diverted away from its current destination to service a request that just occurred in the vicinity of its current position. In the figure below, a vehicle travelling from node A to node B (and thence to node C) is notified of a new request from node D while at the location labelled A' at time t . The vehicle detours from its planned route to attend node D and then resumes its planned route.




Figure 6.1: Diversion from route ABC at location A' to route AA'DBC. Planned route is shown by small dotted line; diverted section by larger dotted line; full line indicates completed section of route.

Ichoua [IGP00] extends this strategy by permitting the diversion of one or more vehicles to occur at a time other than time t and by allowing multiple loads to be carried simultaneously. Requests are more distant from each other, and activities take place over a longer time horizon (a few days), than in [RMJ94,95]. At the time when a new request occurs, the current system status and the request become inputs into some optimization procedure which could take the form of

- a simple insertion by the immediate diversion of a vehicle
- an insertion followed by a re-optimization procedure, or

- an exact algorithm.

Consequently, on receipt of a new request the vehicles may continue along their planned routes to their current destinations or one or more may change. At the time a new request appears, the vehicle is at position $A'(t)$ but has moved to position $A''(t+\delta t)$ in the time taken to process the request, δt . The strategy is to use the information available at time t to provide a solution that can be applied to the system at time $t+\delta t$, whilst using a value of δt that is both large enough to allow the improvement procedure to be performed, but not so large that the opportunities for diversion are lost.




Figure 6.2: New request occurs at time t but is not actioned until time $t+\delta t$.

[IGP00] use Procedure 6.1 [GGPT96] to accommodate new requests. In [G96]'s procedure, the neighbourhood is generated by the inter-route exchange of variable length segments known as the CROSS exchange [T97]. When a new event occurs, such as the completion of service to a customer or the generation of a new request, the search threads are interrupted while the adaptive memory is updated, and re-started with new solutions from the updated memory. When a new request is received, it is inserted into all solutions in adaptive memory and rejected if there is no feasible position in any solution. If feasible solutions exist, the best solution is re-optimised using a local search heuristic based on CROSS exchange. The current destination of each vehicle is fixed and is not part of the planned route to be optimised.

[IGP00] modify this procedure by fixing the current position of each vehicle rather than the current destination, and allowing the current destination to become part of the optimization process. Thus, they apply a limited form of coordination by allowing routes to deviate from their planned paths at locations that are not necessarily co-located with customer nodes, although they do not extend this technique to allow interaction between vehicles. The optimization phase is allocated a processing time of δt dynamically, using either historical values of request inter-arrival times, or future requests on the planned routes. Using uniformly distributed customer locations in a Euclidean framework of dimension 100x100 with travel times proportional to Euclidean distances, the objective function to minimise is the sum of the total distance travelled and the total lateness over all customer locations. Fleet size is set at the number of routes in the best reported solutions. When only 25% of requests are known in advance, both the number of unserved customers and the objective function are improved, with greater improvements when the value of δt is determined in advance as a moving average of the inter-arrival time of the previous n requests (where n is proportional to the number of requests known in advance). When 50% of requests are known in advance the number of unserved customers is further improved.

Parallel Tabu Search	
Step One:	Construct different initial solutions with a stochastic insertion heuristic (where the choice of the next customer to be inserted is randomised).
Step Two:	Apply tabu search to each solution and store the resulting routes in the adaptive memory [RT95].
Step Three:	While a stopping criterion is not met do:
Step Three (a):	Use the routes stored in the adaptive memory to construct the initial solution.
Step Three (b):	Decompose the problem into subproblems obtained through a decomposition procedure.
Step Three (c):	Apply tabu search to each subproblem.
Step Three (d):	Add the routes of the resulting solution in the adaptive memory.

Procedure 6.1: Parallel tabu search approach.

Another investigation into limited dynamic coordination in the Euclidean setting has been performed by Bard et al (see [BHJD98]), who establish coordination sites for en route replenishment of stock in a version of the inventory routing problem. The Inventory Routing Problem (IRP, see [CCKS97]) is an extended version of the VRP that takes into account inventory levels and fixed storage capacities at customers. The IRP is concerned with the repeated distribution of a single product from a single depot to a set of customer nodes over a given planning horizon of known length, by a fleet of homogeneous vehicles. The objective to minimise is the average distribution costs during the planning period; all customer requirements must be satisfied. The IRP is an extended version of the VRP because it is based on customers' usage rather than customers' orders. In addition to the routing and scheduling phases of the VRP, the IRP has a customer selection phase which identifies the customers to be serviced during a particular planning period. Beale (see [B01] Chapter 9) addresses the problem of selecting which customers to service and includes implications for not servicing customers, or for servicing them poorly.

In practice the IRP is stochastic and dynamic. In the stochastic inventory routing problem the future demand of a customer is uncertain and there is a possibility that a customer's demand may not be satisfied (i.e. a stockout occurs).

The IRP is a long-term dynamic control problem which is modelled by short-term decisions. Stockouts can be reduced by delivering to a customer earlier than is necessary, but this usually means delivering less than a full vehicle load which may lead to higher distribution costs. There is a similarity between this approach and the establishment of type 1 coordination sites, where the coordination site is located at the destination node but where that node is not necessarily on a single-noded route and therefore may be unable to have a full vehicle load delivered to it. When the short-term planning problem consists of a single day, the problem is solved by adapting VRP solution techniques. When the planning problem consists of several days, the short-term problems are formulated as mathematical programs and solved using decomposition techniques such as Lagrangean relaxation. Successful solutions usually assume that each visit delivers the maximum quantity possible and that vehicles are deployed with full loads, consistent with the results found in the previous chapter.

[BHJD98] use a decomposition approach to show that a tradeoff can be made between distance and annual costs by establishing coordination sites as satellite facilities where vehicles can replenish their stocks without having to return to the depot. They prevent stockouts by making an immediate special delivery at high cost. Using a planning period of one week and examining two weeks at a time

with a rolling horizon, only the solution for the first week is implemented and the process is repeated for the next two weeks. Customer selection is based initially on those customers whose optimal delivery day falls within this two week period. Servicing a customer on a day other than the optimal day for that customer incurs an incremental cost; other incremental costs involve customer retention and aspects of customer service. Customer assignments to days of the week are based on a combination of a desire to balance demand and to minimise incremental costs. The inventory routing problem with satellite facilities (IRPSF) is decomposed into two subproblems – clustering of customers, and routing.

The IRPSF is defined by a set of customers with known location, each with a non-negative demand for a given commodity that varies daily in accordance with a known probability distribution. Each customer scheduled for service on a given day must be visited by exactly one of the identical vehicles based at a central depot. The depot and each of the coordination sites hold an unlimited supply of the commodity and can reload any vehicle. The depot serves as the origin and final destination of each vehicle.

A feasible sequence for a vehicle starts at the depot, visits a subset of nodes, delivering the required amount to each, possibly reloads at one of the coordination sites, continues to another subset of nodes, reloads again at a coordination site, and at the end of a route returns to the depot, all within a maximum time constraint.

Inventory Routing Problem with Satellite Facilities	
Step One:	Identify those customers in the service area whose optimal replenishment day falls within the current planning horizon.
Step Two:	Assign each customer identified in Step One to a given day by solving a balanced assignment problem with the objective of minimising incremental costs.
Step Three:	For each day in the planning horizon, try to find a good feasible solution by solving a VRP with satellite facilities (VRPSF). The objective is to minimise total distance.
Step Four:	Attempt to improve the solution by swapping customers between routes.
Step Five:	Examine the tradeoff between incremental costs and route lengths by swapping customers between different days of the planning horizon.

Procedure 6.2: The inventory routing problem with satellite facilities.

In the VRPSF, the coordination site to be inserted into a route is determined from a pre-defined list of sites subject to distance constraints. Initial routes are tested for improvement by intra-route edge exchange and inter-route node exchange, and the two-week schedule is checked for improvement by the re-assignment of nodes to other days, or swapping nodes between days. In the Euclidean plane with up to 500 nodes and up to 5 coordination sites, [BHJD98] report greater efficiency in clustering of customers from using a modified Clarke and Wright (see [CW64]) route construction than with GRASP (see [KB95]) or Sweep (see [GM74]) although the Clarke and Wright heuristic took up to 3 times longer than the Sweep heuristic and about 3 times longer than the GRASP heuristic. The authors report that the establishment of coordination sites for the intra-route replenishment of stock implies that the distance traversed by a route is a greater limitation than the cumulative demand. Permitting swaps between routes one and two days apart leads to relatively large reductions in distance due to the greater number of feasible transitions, but increasing the number of days between swaps increases the chance of violating other constraints. The routing distance is a non-increasing function of the number of coordination sites with an almost directly proportional advantage in having more coordination sites.

In this study, coordination is applied to dynamic vehicle routing problems in the rectilinear grid network by extending the results of static systems discussed in the previous chapter and incorporating these into dynamic systems using techniques similar to those discussed in this section.

6.2 Problem description

This study considers dynamic stochastic coordination in a rectilinear grid, using pre-processing of problems in order to create a decision frontier. Rather than divide the work period into time slices, pre-processing enables the dispatcher to select an appropriate response to a new request as the new request is received, from a selection of responses that have already been determined.

Because the planning period covers one work cycle, there is no guarantee that all demands will be satisfied, and route failures will occur where a request is received from a customer node that cannot be reached in the time available without violating the constraint that returns all vehicles to the depot at the end of each cycle. If the planning horizon were extended to cover more than one work cycle, such late distant demands could become inputs into the next cycle. In practice, delivery companies often state that orders received later than a certain time will be

processed the following day. Such an extension to the solution techniques described in this study may provide material for future research.

Problems are generated with a fixed number of nodes created at the vertices of a rectilinear grid network. Unless stated otherwise, nodes have known location and known delivery demand; all nodes have zero pickup demand. The fleet of identical vehicles, each with a maximum capacity, maxVehCap , and maximum total distance allowance, maxVehDist , begins the work cycle from the single depot and all routes end at the depot. All edges of the rectilinear grid are identified and are two-way with effective infinite carrying capacity in each direction.

Where possible, vehicles are deployed with full loads, deliver the maximum amount possible to customers and to coordination sites, and coordination sites are established as early as possible. At the conclusion of a work cycle, all coordination sites will have been dis-established and all vehicles will have returned to the depot. Although penalties will not be applied explicitly, endeavours will be made to service the feasible requests of all customers, even if this means employing extra vehicles. Thus, a form of inherent penalty is imposed by increasing the fleet size. There is no limit on waiting times beyond the requirement to produce feasible routes. Initial routes are constructed using the techniques of Chapter 5 (largest/smallest neighbourhood farthest/nearest insertion improved by exchange and transfer).

In order to keep the following solutions brief and general, there is no prescribed limit on the number of coordination sites that may be established, and no set-up time for such sites. As a demonstration of solution techniques, the structure shown in Figure 6.3 is used as an indicative example only; node locations and demands are given in Appendix F. This structure has been selected because it demonstrates a basic form of the problem under investigation and the discussion is intended to generate indicative results rather than as an in-depth investigation. The depot is centrally located, and the fifteen nodes are spread across a medium density grid characterised by several lakes and peninsulas. Examples in the following sections are based on weak dynamic systems [LMR96] in which most of the information is known in advance and only one node has unknown attributes. The consequences of applying coordination to systems that have multiple nodes with unknown attributes are discussed later.




Figure 6.3: Example structure.

6.3 Nomenclature

Problems are categorised by the prior knowledge (or lack of it) of the node attributes of location, demand and time at which the demand is received or at which the location becomes known. The significance of the time attribute is context-sensitive since ‘time’ may refer to the time at which information on location or demand becomes available, or the time at which deliveries will be accepted, or the time after which deliveries will not be accepted (deadlines). Because this investigation is included merely as a brief look at the implications of coordination in a dynamic environment, only the simplest context-sensitive version of the time attribute is considered, and the following classification is made.

location	demand	time	
√	√	√	Not dynamic
√	√	×	DC1
√	×	√	DC2
×	√	√	DC3
√	×	×	DC4
×	√	×	DC5
×	×	√	DC6
×	×	×	DC7

Table 6.1: Dynamic coordination classification.

Each of the dynamic situations defined in Table 6.1 is discussed briefly below. The intent is to give an indication of the methods by which such problems may be solved using coordination; an in-depth investigation forms the basis for future research.

Where problems exist with nodes that fit more than one category, the classification is used that is numerically lowest in the classification table, if that is practicable. In practice, problems will exist with aspects that belong to different categories. This situation is discussed in more detail in §6.11.

6.4 DC1 Only time is unknown

In time window constrained routing and scheduling problems (e.g. [SD88]), some nodes impose service deadlines and earliest service time constraints. In the dynamic version of this problem, the locations and demands of all nodes are known before the routing and scheduling phases, but the time that at least one node is able to accept delivery is unknown at the beginning of the work cycle and becomes known at some stage during the planning period. There are four main approaches to applying coordination to this problem:

(a) Conservative

One approach to this type of problem consists of including nodes with unknown time in tentative routes that approach their respective locations at different times. As soon as their acceptance times are revealed, the best of the tentative routes service those nodes and the other tentative routes are amended. Because routes are tentative until acceptance times are known, the later parts of routes may not proceed as planned initially, so it is important to establish coordination sites as early as possible. This fits well with the findings of Chapter 5 where the early establishment of coordination sites was found to be beneficial to the routing scheme.

(b) Optimistic

Another approach is to remove such nodes from consideration in the routing process and to insert them into the most appropriate routes when their acceptance times are revealed. The subsequent modification of routes may result in a quite different routing scheme from the initial one, requiring considerable alteration to the environment and greatly changing the potential for coordination. By ignoring nodes with dynamic time attributes, those nodes are removed from any coordination scheme that may exist prior

to the time of demand revelation (although they may be included in schemes that are developed after revelation).

(c) Co-location

In this approach, routing and scheduling is carried out as though the situation were non-dynamic, and coordination sites are created that are co-located with the problem nodes. When acceptance times are revealed, suitable vehicles service such nodes and subsequent coordination schemes may be changed. The allocation of such vehicles and the dis-establishment of coordination sites may increase the time to determine a recourse action and increase fleet size.

(d) Mop-up

The mop-up approach is also based on a static version of the problem, whereby routing and scheduling is carried out without regard to the nodes that have undetermined acceptance times. When those delivery times are revealed, at least one more vehicle is deployed to service these nodes exclusively. Such vehicles may be unavailable for consideration in the coordination process because their spatial and temporal characteristics are difficult to determine until some stage in the work cycle which may be too late to allow their utilisation.

Each of these approaches carries its own implications for the application of coordination. By omitting some nodes from the initial routes, as in all but the conservative approach, an unnecessary handicap is imposed on the system by the inability to take advantage of the prior knowledge of the node locations and their demands. By using a conservative approach and pre-processing problems, the dispatcher is able to select an appropriate response from existing options. The conservative technique, shown in Procedure 6.3, may include aspects of co-location by the early establishment of coordination sites co-located with target nodes and guaranteeing delivery by imposing a latest delivery constraint on the customer.

Conservative DC1 Procedure	
Step One	Initialise: $U :=$ set of all nodes with unknown time attribute $N :=$ set of all nodes with known time attribute $R :=$ set of all routes.
Step Two	Route and schedule all nodes as for static coordination, except that nodes in U are included in routes that approach them at different times.
Step Three	As the time attribute is revealed for each node, x , in U , 1. remove x from U and add it to N. 2. choose the most appropriate route to service it from R; if no such route exists, go to Step Four. 3. adjust routes in R. 4. if U is not empty, continue with Step Three, otherwise STOP – all nodes have been successfully serviced.
Step Four	If the node can be serviced without violating time constraints, create a single-noded route and service it with a dedicated vehicle, otherwise STOP – the node cannot be serviced.

Procedure 6.3: Conservative DC1.

6.4.1 Time constraint

Assuming a one-to-one relationship between distance and time, the existence of a maximum distance constraint infers the presence of a maximum time constraint of equal magnitude. Consequently, the maximum time during which an order for a node, u , with unknown time attribute can be guaranteed service is

$$\text{maximum available time, } t_{\max} = \text{maxVehDist} - \text{dist}(\text{depot, node } u) .$$

Such service can be guaranteed by placing a dedicated vehicle at the node and returning it feasibly to the depot. Alternatively, any vehicle servicing a route with node u as its ultimate node is able to remain at node u until $t=t_{\max}$. Similarly, when node u is placed at other positions in a route, its waiting time can be adjusted with regard to the travel times between other nodes.

6.4.2 Multiple placement

By constructing routes that contain the node with dynamic demand in more than one position, the situation is created whereby one vehicle can investigate the acceptance status of that node at more than one time. This is shown in Figure 6.4, which consists of a depot, one target node with unknown time attribute at location A, and several other nodes on a route that passes twice through the node at A. If the target node is visited before $t=t_{rev}$ and the vehicle is unable to wait until $t=t_{rev}$, the node cannot be serviced, and the vehicle must continue along the generic path indicated in (a). However, if the node at A can be serviced when the vehicle first passes through A (as in (b)) the remainder of the route may be reorganised so as to avoid subsequent passages through A.




Figure 6.4: Multiple placement in a route of a node with unknown time attribute.

6.4.3 Problem generation

Consider the example in Figure 6.5, where the attributes of all nodes except node 15 are known at the time of routing and scheduling, and are given in Appendix F. Node 15 has known location and demand (as in Appendix F) but its earliest acceptance time, t_{rev} , will not be revealed until the work cycle has begun. Node 15 has been chosen for this example because its relatively central location and proximity to other nodes makes it easier to include in multiple routes than an isolated node, such as nodes 1 or 6. In order to make the problem simple but non-trivial, a demand of 50 has been allocated arbitrarily to node 15; split deliveries are permitted, and only the earliest acceptance time is considered with no regard for the latest acceptance time. As discussed in the following sections, several routing schemes are created that contain node 15 in different positions, one of which will be adopted.

Because node 15 lies 10 units from the depot, along its shortest path, t_{max} has value 40; consequently, any vehicle servicing a route with node 15 as its ultimate node is able to remain at node 15 until $t=40$, at which time it must return to the depot.

For this problem, the minimum number of routes (equivalent to the value of $\text{ceil}(\text{total demand}/\text{maxVehCap})$) is five.




Figure 6.5: DC1. Node 15 has dynamic demand. The three coordination sites indicated are discussed below.

6.4.4 Incorporation of coordination

The schemes shown below are representative of the principles that could be followed when dealing with one node with dynamic demand. The total distances reported are those that would be travelled under circumstances where the acceptance time for node 15 is revealed later rather than earlier. Where the acceptance time is earlier, routes shown may be adjusted so that the total distance is less than that reported here.

6.4.4.1 Routing scheme A (node 15 approached at $t=12$ and $t=16$)

The vehicle servicing route 1 approaches node 15 at $t=12$ before continuing on to deliver part of node 14's requirement and then returning to node 15 (which lies on a shortest path back to the depot) at $t=16$. If node 15 has been serviced, the vehicle will return to the depot; otherwise, it may wait until $t=40$ to receive permission to deliver node 15's load. If such permission is not forthcoming by $t=40$, the vehicle returns to the depot and node 15 is not serviced. The scheme is shown in Table 6.2.

Route #	Distance	Route
1	24	0 15 14' 15 0
2	50	0 14'' 1 3 0
3	50	0 5 8 9 11 0
4	36	0 13 10 7 0
5	36	0 6 12 0
6	20	0 2 4 0
	216	

Table 6.2: Routing scheme A for DC1 example. Five vehicles service these routes.

If node 15 had been serviced at $t=12$, it need not be approached again at a later time and the vehicle servicing route 1 could have returned to the depot by any shortest path between node 14 and the depot (which in this problem leads it past node 15's location anyway). Consequently, the total distance is unchanged regardless of when node 15 can accept delivery. Although the total distance of scheme A is quite high, it has the advantage of early encounters with the target node, one of which allows a vehicle to remain at the node for a considerable time, $16 \leq t \leq 40$, thus increasing the chances of providing prompt service to the customer at node 15.

6.4.4.2 Routing scheme B (node 15 approached at $t=24$ and $t=40$)

Nodes 2 and 14 have split deliveries. The vehicle servicing route 2 approaches node 15 at $t=24$, at which time the node is either serviced or ignored. If servicing of node 15 occurs at $t=24$, the vehicle servicing route 6 is not deployed, and the total distance becomes 196. The latest time at which the vehicle servicing route 6 must leave the depot is $t=30$, so the vehicle servicing route 1 is able to service route 6 if required, without violating time and distance constraints, and no increase in fleet size beyond 5 is incurred. If the vehicle servicing route 1 had not been able to continue as a multi-routed vehicle, a fleet increase would have been required at $t=30$. The scheme is shown in Table 6.3.

Route #	Distance	Route
1	30	0 6 2' 0
2	38	0 2" 4 12 3 15 14' 0
3	42	0 14" 1 0
4	50	0 5 8 9 11 0
5	36	0 13 10 7 0
6	20	0 15 0
216		

Table 6.3: Routing scheme B for DC1 example. Five vehicles service these routes.

6.4.4.3 Routing scheme C (node 15 approached at $t=26$ and $t=30$)

The vehicle servicing route 1 establishes a coordination site at W which may be used by route 2 or dis-established by route 1 no later than $t=44$. If node 15 is able to be serviced by $t=26$, route 6 is adjusted by removing node 15 and reducing the total distance travelled to 218; otherwise, the vehicle servicing route 6 may wait there for permission to service the node from $t=30$ to $t=40$, after which the vehicle returns to the depot and node 15 is not serviced. The scheme is shown in Table 6.4.

Route #	Distance	Route
1	50	0 W ^{del 50} 9 8 5 11 W 0
2	40	0 14 W _{col 50} 15 3 0
3	36	0 13 7 10 0
4	44	0 1 12 0
5	20	0 2 4 0
6	40	0 6 15 0
		230

Table 6.4: Routing scheme C for DC1 example. Six vehicles service these routes.

6.4.4.4 Routing scheme D (node 15 approached at $t=12$ and $t=36$)

Node 15 is approached at $t=12$ on route 3; if servicing is not available at this time, route 3 is traversed until node 15 is again approached at $t=36$ at which time either the node is serviced or the vehicle may wait until $t=40$ before servicing the node and returning to the depot. Nodes 2 and 14 have split deliveries. If node 15 is able to be serviced at $t=12$, the total distance drops to 196. An advantage of this scheme is that it causes no disruption to routes other than the one which services the node with dynamic demand. Also, by the time the second encounter with node 15 occurs, at $t=36$, two routes have been completely traversed and the vehicle servicing a third is already approaching the depot, leaving only one route which is in any position to be re-assessed. Subsequent assessment reveals that this route should not be changed, so that late encounter results in no changes to other routes. The scheme is shown in Table 6.5.

Route #	Distance	Route
1	30	0 6 2' 0
2	38	0 2'' 4 12 3 14' 0
3	46	0 15 14'' 1 15 0
4	50	0 5 8 9 11 0
5	36	0 13 10 7 0
		200

Table 6.5: Routing scheme D for DC1 example. Five vehicles service these routes.

6.4.4.5 Routing scheme E (node 15 approached at $t=12$ and $t=16$)

Node 15 is approached at $t=12$; if servicing is not available at this time, a coordination site is established at V in order for the vehicle servicing route 2 to be able to wait there until $t=40$ or until node 15 is serviced, whichever comes first. At that time, either node 15 is serviced, or the site is dis-established and the vehicle returns to the depot. The scheme is shown in Table 6.6.

Route #	Distance	Route
1	44	0 15 V_{del}^{50} 3 1 0
2	26	0 12' 14 V_{col}^{50} 15 0
3	26	0 12'' 4 2 0
4	50	0 5 8 9 11 0
5	36	0 13 10 7 0
6	28	0 6 0
		210

Table 6.6: Routing scheme E for DC1 example. Six vehicles service these routes.

6.4.4.6 Routing scheme F (Dedicated vehicle)

A dedicated vehicle is assigned to node 15 to wait at that node until either permission to deliver is granted or $t=40$, whichever is earliest. Since node 15 is 10 distance units from the depot, the service time for this node lies between $t=10$ and $t=40$. The scheme, shown in Table 6.7, involves split deliveries to node 2.

Route #	Distance	Route
1	50	0 11 9 8 5 0
2	36	0 7 10 13 0
3	20	0 15 0
4	30	0 6 2' 0
5	30	0 Y_{del}^{50} 3 12 4 2'' 0
6	46	0 14 Y_{col}^{50} 1 0
		212

Table 6.7: Routing scheme F for DC1 example. Five vehicles service these routes.

This scheme provides a customer whose demand is dynamic with a service equal to that which it would have received had its demand been static, with an identical range of delivery times, $10 \leq t \leq 40$, but under scheme F the delivery can be guaranteed to be made intact. Had the demand of node 15 been static, the routing scheme would have been that shown in Table 6.8. Differences between the static and dedicated dynamic schemes are evident in the treatment of nodes 2 and 4 which, under the dedicated dynamic scheme, are serviced later in the work cycle, and in which the demand of node 2 is delivered in split form.

Route#	Distance	Load	Route
1	36	193	0 7 10 13 0
2	50	167	0 11 9 8 5 0
3	20	194	0 2 4 0
4	28	119	0 6 0
5	44	168	0 12 3 1 15 0
6	24	200	0 14 0
	202	1041	

Vehicle	Distance	Routes(s)	Nodes
1	36	1	0 7 10 13 0
2	50	2	0 11 9 8 5 0
3	28	4	0 6 0
4	44	5	0 12 3 1 15 0
5	44	3 6	0 2 4 0 14 0

Table 6.8: Routing scheme for static DC1 example without coordination. Five vehicles service these routes.

6.4.4.7 Routing strategy

The schemes shown above are not exhaustive, and combinations of them may be appropriate. Schemes D and E result in shorter distances than the dedicated vehicle approach of scheme F, and scheme B may be shorter (depending on when node 15 can be serviced). Which scheme is adopted depends on several factors including

- the range of approach times,
- fleet requirements such as fleet size and equitable route distances, and
- the manner in which the adopted scheme affects other nodes which may have their own dynamic considerations.

6.4.5 Conclusions

DC1 problems are solved by generating routes that include the node with unknown time attribute in several different positions, in order that the node is approached by vehicles with enough stock to service its demand. By utilising routes that encounter the target node more than once, as in Figure 6.4, disruption to other routes can be reduced.

Encountering the target node as the ultimate node in a route allows the vehicle to wait at the node (or near it) until either the delivery can be made or the vehicle is forced to return to the depot, without disrupting service to other customers. Where the node is followed by other nodes in a route, waiting at the target node will cause uncertain delays in delivery to those other nodes. Psaraftis (see [PT93]) describes systems which allow vehicles to wait at nodes until the environment is more favourable.

As the number of nodes with dynamic demand increases, it becomes more difficult to provide specific schemes such as those shown in the example above, and routing and scheduling may be more fluid with routes being re-constructed as acceptance times of nodes are revealed and their requirements are added to the active pool of requests. As routes are re-formed, the environment for coordination changes and it becomes more difficult to predict suitable coordination schemes. Consequently, there may be a limit to the use of pre-processing in the improvement by coordination of problems with dynamic demand. This is discussed in more detail in §6.11.

6.5 DC2 Only demand is unknown

In vehicle routing problems with random demand, all node locations are known in advance but the demand of at least one node will not be known either until some pre-specified time, $t=t_{rev}$, or until the node is visited. If the demand has zero value and the demand is revealed only when the node is visited, the node need not be serviced.

If a node's demand exceeds vehicle capacity, a route failure may occur and a recourse action will be needed. Such action may involve the vehicle returning to the depot to replenish stock or, where coordination is available, proceeding to a coordination site to collect the required amount of stock, before returning to the target node. If split deliveries are permitted, replenishment stock may be carried by another vehicle. Recourse action must guarantee, if possible, that the demands of all nodes will be satisfied, and the objective function to be minimised involves the total expected distance. Effective use of coordination in systems with stochastic demand involves the anticipation of route failure and the establishment of coordination sites that can provide replenishment stock.

Approaches to this problem in the literature are based on a priori optimisation and reoptimisation. In a priori optimisation (see [BJO90]), nodes are arranged in pre-determined static routes that are not updated during real-time operations. Bertsimas (see [BS96]) reports on Euclidean schemes that dynamically update initial fixed a priori routes depending on when a customer's demand becomes available: a fixed a priori strategy visits all customers in the same fixed order as under the a priori sequence but replenishes only those that require stock; an adaptive a priori strategy updates routes by skipping nodes with zero demand.

In reoptimisation schemes (e.g. [DLT89] and [DP93]), routing occurs concurrently with service. Where the demand is known on arrival at the node, the decision to service that node is made on arrival, and the route ends when the vehicle is empty or when all nodes have been serviced. Reoptimisation schemes are characterised by computationally intensive solutions, and the first computationally tractable heuristic procedure using this approach is reported by Secomandi (see [S01]), who sequentially improves an a priori solution.

Dror (see [DLT89]) discusses stochastic vehicle routing in a single depot/homogeneous fleet Euclidean setting where customer demand is the only source of uncertainty. Since split deliveries may be permitted and demands may become evident over time, the study represents a form of dynamic stochastic coordination. Routes can be interrupted before demands become known, so that vehicles can replenish stock in anticipation of new demands. A stochastic programming

formulation of this problem does not permit route reoptimisation, whereas formulation of the problem as a Markov decision process allows reoptimisation on receipt of new data at any location. Reoptimisation instructions may be created in the form of a decision tree for each stop in the route.

Stochastic methods attempt to use the stochastic nature of demands to predict future requirements so that customers can then receive reasonable service with a minimum of disruption. The main approaches include Markov decision processes (see [P88]) and stochastic programming (see [DLT89]). Markov decision processes are hindered by the rapid growth in state space with problem size and the need for simplifying assumptions to make solving the model more tractable.

6.5.1 Problem generation

A single depot/fifteen node system (with raw data given in Appendix F) is generated in a rectilinear grid where the attributes of nodes 1 - 14 are fully determined but where node 15 has a known location and a demand that will become apparent at time $t=15$. Node 15 has had only one previous demand of value 50. Although coordination is possible in the form of split deliveries, that technique will not be applied in the initial construction phase since it might serve to cloud the effect of the dynamic nature of the demand. There is no load/unload time at nodes, so the relationship between time and distance is one-to-one as in §6.4, and all vehicles depart the depot at $t=0$. The accommodation of delayed departure times (see [FBSTW93]) is an area for future research; it has not been addressed in this study in order to keep the examples as unobfuscated as possible.

6.5.2 Problem solution

This study employs a strategy similar to the adaptive a priori approach by creating a set of initial routes using an estimation of the value of the unknown demand based on the demand history of that and other nodes. After initial routes are constructed, a coordination plan is assembled and the preliminary estimation is adjusted to a greater value; the sizes of loads delivered to and collected from coordination points are modified. At a given time, the actual demand of all nodes is revealed and the necessity of applying the coordination is determined.

For a problem with only one node with stochastic demand, solution is a four step process, as shown in Procedure 6.4.

DC2	
Step One	Initial routes are constructed and scheduled using an estimated value for the unknown demand.
Step Two	Routes are adjusted so that the node in question is not visited before the time at which its true demand is revealed, $t=t_{rev}$.
Step Three	Coordination points are then established at likely sites so that they will be available for use when the mystery demand is revealed.
Step Four	At the time $t=t_{rev}$, routes are adjusted to accommodate, if possible, the known demand. Coordination sites are emptied before the end of the work cycle.

Procedure 6.4: DC2.

6.5.3 Demand estimation

In the simulation environment, a random number with a normal distribution generates demand. The unknown demand is estimated by considering the number of past occurrences of demand from that particular node:

- >3 : use the average of the last four demands
- <4 : use the latest demand
- no previous demands: use the average of all demands for the known nodes.

Where there is no pattern, using the last four values accommodates extremes of demand.

Since node 15 has had only one previous demand of size 50, the estimated value of Demand(15) is 50.

6.5.4 Initial route construction

Routing and scheduling is carried out as for a non-coordinated system using the estimated demand for node 15, via any method that creates feasible routes. The particular method of routing and scheduling is irrelevant to the concept of coordination, since coordination requires only that routes, nodes and demands exist. In this example, the routes are constructed using the techniques discussed in Chapter 5, and are shown in Table 6.9 and Figure 6.6.

Route#	Distance	Load	Route
1	36	193	0 7 10 13 0
2	50	167	0 11 9 8 5 0
3	20	194	0 2 4 0
4	28	119	0 6 0
5	44	168	0 12 3 1 15 0
6	24	200	0 14 0
	202	1041	

Vehicle	Distance	Routes(s)	Nodes
blue	36	1	0 7 10 13 0
green	50	2	0 11 9 8 5 0
pink	28	4	0 6 0
black	44	5	0 12 3 1 15 0
red	44	3 6	0 2 4 0 14 0

Table 6.9: Initial routes using an estimated value for Demand(15) without coordination.




Figure 6.6: Initial routes generated with an estimated demand for node 15.

6.5.5 Coordination site location

Site location consists first of determining the position of all vehicles and the serviced/unserviced state of nodes using the initial routes, and then preparing for large differences in the values of the estimated and actual demands of node 15.

Since it is known that the true value of Demand(15) will be revealed at time $t_{rev}=15$, the locations of all vehicles at that time can be determined in advance. At $t=15$, the progress of the five vehicles is that shown in Figure 6.7, with each route having been traversed for 15 time units, and some nodes already serviced (indicated in the table by highlighting and in the figure by filled discs). For the sake of clarity in the figure, some edges have been slightly off-set.

Vehicle	Progress at $t=15$
blue	0 7 10 13 0
green	0 11 9 8 5 0
pink	0 6 0
black	0 3 1 15 0
red	0 2 4 0 14 0

Table 6.10: Progress at $t=15$. Serviced nodes are highlighted.




Figure 6.7: Progress at $t=15$. Filled discs indicate nodes whose demands have been fully met.

Since the red vehicle is multi-routed, the progress shown in the figure is only one of two possibilities depending on which order these routes are undertaken. Similarly, each of the multi-noded routes can be traversed in the opposite direction. The time $t=15$ is late enough that slight changes in the order in which some of the nodes are visited has no effect here.

The basis for coordination site location is contingency planning. If the actual demand is no greater than the estimated value, coordination is not necessary, but the possibility of an actual value far greater than the estimated value must be considered and planned for in advance of any scheduling. All routes must be considered, since the application of coordination can have a ripple effect on distant routes. Accommodation of a larger demand for node 15 involves the use of coordination, and schemes are devised that take into account a range of values for node 15's demand.

In the initial route structure, the vehicle servicing node 15 has 32 surplus stock units available which means that the initial scheme can be used without change for values of Demand(15) up to 82; beyond this value, coordination will be applied.

Accommodation of a much greater value of Demand(15) involves the creation of three coordination points, W, X and Y, as follows. The pink vehicle leaves the depot at $t=0$ with a full payload; it delivers $\text{maxVehCap} - \text{demand}(6) = 81$ to the coordination site at X before continuing on to service node 6 and returning to the depot, increasing its distance by 2. The red vehicle leaves the depot on its first route with a full payload and after servicing nodes 2 and 4 empties the coordination site at X and delivers the contents to a coordination site at Y, adding 6 to its original distance. The black vehicle services nodes 12, 3 and 1, collects the 81 stock units from Y and delivers these to node 15 along with the 82 units it is already carrying for that node without incurring any additional distance. Thus, the demand of node 15 is accommodated up to the value of $82+81+6=169$ by using coordination sites at X and Y, and increasing the total distance by 8.

Since node 15 can have demand up to the value of maxVehCap , the use of coordination sites X and Y may not be adequate. Extra stock can be delivered to node 15 by the establishment of a coordination site at W by the green vehicle without incurring extra distance. Since this vehicle can carry an extra 33 units, it delivers these to W while servicing route 2. If the demand for node 15 is revealed as being greater than 169, the blue vehicle may be diverted to W to empty the coordination site and deliver the stock to node 15, adding to its original distance. Because the requirement of node 15 is known before the red vehicle approaches the coordination site at X, and because it will not have enough surplus distance

capacity to visit both X and W, the following scheme is adopted if $\text{demand}(15) > 169$: the green vehicle establishes a coordination site at W with 33 units; the pink vehicle establishes a coordination site at X with 81 units; the red vehicle transfers the contents of the site at X to a new site at Y, along with its surplus payload; after servicing nodes 12, 3 and 1, the black vehicle transfers the contents of the site at Y to the site at Z and returns to the depot; the blue vehicle services node 13 then detours to the site at W which it empties, travels to the site at Z which it also empties and delivers the contents to node 15.

In each of the above scenarios, the dis-establishment of coordination sites must be included. The schemes are combined to create the initial routes shown in Figure 6.8; their implementation is described next.




Figure 6.8: Initial routes incorporating coordination points which may be used if the demand of node 15 exceeds 82.

6.5.6 Incorporation of coordination

The schemes described above are combined and formalised in such a manner that they have identical configurations up to the time $t=t_{rev}$, thus allowing any of them to be adopted when the actual demand of node 15 is revealed.

6.5.6.1 Routing scheme A

No coordination is applied; the routes are those shown in Table 6.11 with 82 units carried for possible delivery to node 15. The coordination sites established at X and W are unused and emptied by the black and green vehicles respectively on their return to the depot. Differences between the schemes represented by Tables 6.9 and 6.11 result from the incorporation of coordination: because scheme A exists in a system in which coordination may be required, it must create the environment for $t \leq t_{rev}$ which allows subsequent coordination to be applied by another routing scheme. The routes are those shown in Figure 6.8 above, except that coordination sites at Y and Z are not established. The sites at W and X are emptied by the green and black vehicles respectively.

Route#	Distance	Route
1	36	0 7 10 13 0
2	50	0 W 11 9 8 5 W 0
3	20	0 2 4 0
4	30	0 X 6 0
5	44	0 12 3 1 15 X 0
6	24	0 14 0
204		

Vehicle	Distance	Routes(s)	Nodes
blue	36	1	0 7 10 13 0
green	50	2	0 W 11 9 8 5 W 0
pink	28	4	0 X 6 0
black	44	5	0 12 3 1 15 X 0
red	44	3 6	0 2 4 0 14 0

Table 6.11: Routing scheme A.

6.5.6.2 Routing scheme B

As shown in Table 6.12:

1. Pink vehicle establishes a coordination site at X with 81 units, services node 6 and returns to the depot.
2. Red vehicle services nodes 2 and 4, collects the stock at X and transfers it to a new coordination site at Y, returns to the depot and then services route 6.

3. Black vehicle services nodes 12, 3 and 1, collects contents at Y and delivers them to node 15 before returning to the depot.
4. Green vehicle establishes a coordination site at W with 33 units, services nodes on route 2 and empties the site at W on its return to the depot.
5. Blue vehicle services nodes on route 1 and returns to the depot.

Route#	Distance	Route
1	36	0 7 10 13 0
2	50	0 W 11 9 8 5 W 0
3	26	0 2 4 X Y 0
4	30	0 X 6 0
5	44	0 12 3 1 Y 15 0
6	24	0 14 0
210		

Vehicle	Distance	Routes(s)	Nodes
blue	36	1	0 7 10 13 0
green	50	2	0 W 11 9 8 5 W 0
pink	30	4	0 X 6 0
black	44	5	0 12 3 1 Y 15 0
red	50	3 6	0 2 4 X Y 0 14 0

Table 6.12: Routing scheme B.

Although a coordination site is established at W, it is unused, and its contents collected by the green vehicle on its return to the depot. The routes are those shown in Figure 6.8 above, except that the coordination site at Z is not established.

6.5.6.3 Routing scheme C

As shown in Table 6.13:

1. Pink vehicle establishes a coordination site at X with 81 units, services node on route 4 and returns to the depot.
2. Red vehicle services nodes 2 and 4, collects the stock at X and transfers it, along with its own surplus, to a new coordination site at Y, returns to the depot and then services route 6.

3. Black vehicle services nodes 12, 3 and 1, empties the site at Y, travels to the coordination site at Z where it delivers its payload before returning to the depot.
4. Green vehicle establishes a coordination site at W with 33 units, services nodes on route 2 and returns to the depot.
5. Blue vehicle travels to the coordination site at W which it empties, then travels to the site at Z which it empties, then delivers its payload to node 15 and returns to the depot.




Figure 6.9: Routing scheme C.

Route#	Distance	Route
1	50	0 7 10 13 W Z 15 0
2	50	0 W 11 9 8 5 0
3	26	0 2 4 X Y 0
4	30	0 X 6 0
5	42	0 12 3 1 Y Z 0
6	24	0 14 0
		222

Vehicle	Distance	Routes(s)	Nodes
blue	50	1	0 7 10 13 W Z 15 0
green	50	2	0 W 11 9 8 5 0
pink	30	4	0 X 6 0
black	42	5	0 12 3 1 Y Z 0
red	44	3 6	0 2 4 X Y 0 14 0

Table 6.13: Routing scheme C.

At $t=t_{\text{rev}}$, the blue vehicle is instructed to change its path between node 13 and the depot to detour to the coordination site at W, where it delivers its surplus payload. It arrives at W at $t=34$, by which time the green vehicle has established the coordination site, and by the time the blue vehicle arrives at Z, the black vehicle has established the second coordination site. This scheme takes advantage of the grid network structure by applying a detour to an alternate shortest path between the ultimate node of a route and the depot at such a time that no vehicle is required to wait. The 14 surplus distance units of the blue vehicle are used in addition to the extra coverage afforded by the alternate shortest route in order to empty coordination sites W and Z and subsequently service node 15.

6.5.6.4 Routing scheme D

As shown in Table 6.14:

1. Pink vehicle establishes a coordination site at X with 81 units, services node 6 and returns to the depot.
2. Red vehicle services nodes 2 and 4, empties the site at X and returns to the depot before servicing route 6.
3. Black vehicle services nodes 12, 3 and 1 before returning to the depot.
4. Green vehicle establishes a coordination site at W with 33 units, services nodes on route 2 and returns to empty the site at W and travels to the depot.
5. Blue vehicle delivers 4 units to the coordination site at W then services nodes on route 1 and returns to the depot.

Route#	Distance	Route
1	36	0 7 10 13 0
2	50	0 W 11 9 8 5 W 0
3	20	0 2 4 0
4	30	0 X 6 0
5	40	0 12 3 1 X 0
6	24	0 14 0
200		

Vehicle	Distance	Routes(s)	Nodes
blue	36	1	0 7 10 13 0
green	50	2	0 W 11 9 8 5 W 0
pink	28	4	0 X 6 0
black	44	5	0 12 3 1 X 0
red	44	3 6	0 2 4 0 14 0

Table 6.14: Routing scheme D.

Although coordination is not actually needed in this scheme, the sites at X and W must be established so that they may be used by other routing schemes at $t=t_{rev}$, even though this increases the distance. The coordination sites are emptied by the black and green vehicles without adding extra distance. This scheme differs from the others at $t=t_{rev}$ by re-organising route 5 to eliminate node 15, and because that node was the ultimate node in the route, the only improvement required is to truncate the route without affecting other nodes.

6.5.7 Limitations and accommodations

Although the routing strategy developed above is not unique, it is nevertheless effective, and progressive in that each routing scheme represents a slight change over its preceding scheme, which makes for easy implementation of such a strategy.

Routing schemes B and C are based on the ability of the pink vehicle to establish a coordination site at X with 33 units. The vehicle has enough surplus distance available to be able to deliver maxVehCap units directly to node 15, increasing its distance by 20, which would make routing scheme C unnecessary were it not that the implementation of scheme B results in a shorter total distance than this would create.

Since coordination sites are deemed to have unlimited capacity, the ability of the pink vehicle to deliver any units to the site at X depends on the demand of node 6:

$$\text{surplus capacity of pink vehicle} = \text{maxVehCap} - \text{demand}(6) .$$

While node 6 has a non-zero demand, there can be no guarantee that the pink vehicle alone will be able to accommodate all of the demand for node 15. The only other vehicles passing close to X are the red vehicle (carrying 6 extra units) and the black vehicle (already fully laden). Consequently, when the actual demand of node 15 exceeds the estimated demand by more than the surplus capacities of the pink and red vehicles, another coordination site must be established by the blue and/or green vehicles, thereby implementing routing scheme C and creating an upper limit for the demand of node 15 using schemes A – D of $\text{maxVehCap} - \text{demand}(6) + 128$. This value exceeds that of maxVehCap for the data in the given example. The routing strategy then becomes that shown in Table 6.16.

Actual value of demand(15)	Routing scheme
0	D
$0 < \text{demand}(15) \leq 82$	A
$82 < \text{demand}(15) \leq \text{maxVehCap} + 88 - \text{demand}(6)$	B
$\text{maxVehCap} + 88 - \text{demand}(6) < \text{demand}(15) \leq \text{maxVehCap} + 128 - \text{demand}(6)$	C

Table 6.16: Routing strategy based on limitations of node 6.

6.5.8 Conclusions

In deterministic coordination, all factors (demand, location of nodes, availability) are known before the start of a work cycle and all vehicles leave the depot with their routes firmly decided. In contrast, dynamic coordination exists in an uncertain atmosphere where at least one of the factors is unknown.

In the example presented in this discussion, coordination has produced a most acceptable result. Coordination has not been applied in the initial routing in this discussion in order to highlight the effect of the dynamic variant. In accommodating the dynamic demand of a single node, the fleet size has not been increased, and the total distance ranges only from 200 to 222. By creating a routing strategy based on the progressive alteration of routing schemes, implementation is a straight forward affair that consists of choosing a suitable response from a pre-determined selection, each choice of which has similar set-up requirements in the $t \leq t_{rev}$ phase.

Because coordination depends heavily on the environment created by routes, extending the methods presented here to a system with more than one node with unknown demand involves understanding the interactions between nodes to determine the manner in which the demand of one node affects the behaviour of routes. Such an investigation has been alluded to in Table 6.16, where the routing scheme is shown to be dependent on the demands of node 15 and node 6; similar relationships exist with other nodes.

6.6 DC3 Only location is unknown

In DC3 problems, the demands of all nodes are known in advance, but the location of at least one node is unknown until a pre-specified time which may differ for each node with unknown location. As in the previous sections, coordination is most effective when advantage is taken of all the known characteristics of a system including attributes of nodes and grid structure. Four approaches to applying coordination to systems in the rectilinear grid network in which the location of at least one node is unknown at the start of the work cycle are:

(a) Optimistic

The system is divided into two sets, a pool of nodes with advance requests and a pool of nodes with real-time requests (as in [LMR96]), but only the nodes in the pool of advance requests are assigned to routes and vehicles. As each new location is revealed, the associated node is transferred from the pool of real-time requests to the pool of advance requests and inserted into a suitable route (either an existing route or a new route that is generated). This approach takes no advantage of the known attributes of nodes in the pool of real-time requests and, because coordination sites tend to be more effective when established early in the work cycle, may remove those nodes from participation in the coordination process.

(b) Conservative

Each node with unknown location is placed provisionally into several routes which are modified as the locations are revealed. This method reduces the risk of route failure by over-subscribing to routes (that is, by including the same node in more than multiple routes, all but one of which will be modified later) but can result in inflated fleet sizes (because in effect the vehicle capacities have been reduced for those vehicles with the redundant nodes). Because routes may be created in regions of the grid that need not be visited when the actual locations are revealed, the temporal and spatial characteristics that are assumed in the initial stages for the application of coordination may not be those that exist in later stages and a false environment for coordination may exist.

(c) Geographical

Advantage is taken of the known grid structure to create coordination sites at positions which may be useful when all locations are known. Such places include bottlenecks, areas close to paths that are heavily used by routes servicing nodes with known locations, and spots on the edge of barriers,

such as lakes, which are likely to be used. The geographical dependence means this approach is more successful in medium density and sparse grids than in dense grids. By utilising paths that are heavily used, this system is more successful in systems where there is a higher number of nodes with known location.

(d) Dedicated vehicle

Nodes with advance requests are assigned to vehicles and routes without regard for those nodes with unknown locations. As each new location is revealed, a vehicle is deployed to service the corresponding node. This approach may remove those vehicles from the coordination process because their relatively late involvement may mean that their temporal and spatial characteristics are incompatible with existing coordination schemes. By requiring such vehicles to leave the depot at the time the location is revealed, route failure may occur where the location is revealed too late for a vehicle to make the return trip from the depot without violating time constraints. A variant on this approach is to deploy mop-up vehicles early in the work cycle so that they may be closer to the appropriate nodes than to the depot when locations are revealed.

By using scenario planning and pre-processing, the problem is akin to one of dynamic deterministic facility location. In locating freight carrier transportation terminals, Campbell (see [C90]) considered costs associated with transportation, location and re-location and found that re-location costs were detrimental to the efficiency of a system when seeking an optimal solution. Schilling (see [S80]) used a different approach by employing a maximal covering problem formulation in the location of emergency services and assumed that once a facility opened it remained so. In this study, the cost of establishing a coordination site is merely the travel cost and increase in fleet requirement (if any), and once a site has been emptied it no longer exists; it is assumed that there are no maintenance costs associated with coordination sites. It will be shown below that a combination of geographic and dedicated vehicle approaches with an allied covering problem provides an effective solution to the problem of servicing nodes with unknown location.

6.6.1 Problem generation

In order to keep the examples as simple as possible, a system comparable to that in previous sections is generated. That is, a single depot/ fifteen node system (with raw data given in Appendix F) is generated in a rectilinear grid where the attributes

of nodes 1 - 14 are fully determined but where node 15 has a known demand of 50 and a location that will become apparent at time $t=10$. Although coordination is possible in the initial construction phase, that technique will not be applied since it might serve to cloud the effect of the dynamic nature of the location. There is no load/unload time at nodes, so the relationship between time and distance is one-to-one as in §6.4, and all vehicles depart the depot at $t=0$, unless specified otherwise.

6.6.2 Problem solution

Because only one node has unknown location with a demand that will be revealed at a known time, $t=t_{rev}$, solution is based on a similar static problem. As in the previous sections, a solution is sought that generates a set of appropriate responses from which one can be chosen to accommodate the actual value of the previously unknown attribute.

A combination of the above schemes is presented here as a solution to the problem with one unknown node location. Rather than making accommodation for the demand of node 15 in each route, several coordination sites are established by $t=t_{rev}=10$ in such a way that they may be used by vehicles to service any area of the grid and still return to the depot by $t=50$. The problem is similar in some respects to the location of emergency vehicles (see [BL93]), in which:

1. There is a set of demand points, each of which represents a geographic area to which service might be provided.
2. There is a set of vehicle locations, where each vehicle location is associated with a set of demand points that can be reached from that location.
3. There is a constraint on the time it takes to service the demand point which may be different for different demands. In the emergency service scenario, emergencies are prioritised (see [HM00]) with different response times from the call-out time to arrival at the node; in this problem the time constraint relates to the ability of a vehicle to service the node and return to the depot before the end of the work cycle.

By requiring the geographic coordination sites to be established by $t=t_{rev}=10$, the sites chosen must lie on vertices along the red lines in Figure 6.11, between the locations marked by red discs and the depot, and yet must be able to service any part of the black edges and still be able to return to the depot no later than $t=50$. In addition, there should be little overlap between the regions that can be serviced by the geographic coordination sites. The region that could be serviced

by a site at S is shown in blue; the region that cannot be covered by a site at S can be serviced by a site at T (which can also service any vertex other than the two indicated in Figure 6.11). If t_{rev} had had a value larger than 10, a vehicle would have had been able to traverse more of the grid before establishing the site but the site would have afforded less coverage than that shown in the figure. If the vehicle were to be dedicated to node 15, it could have stayed at the depot and still been able to service any vertex, but by establishing sites at S and T, those vehicles may be able to service other nodes in addition to node 15.

By creating initial routes that involve only those nodes whose locations are known, the base routing scheme shown in Table 6.17 is generated.

Vehicle	Nodes
blue	0 7 10 13 0
green	0 11 9 8 5 0
red	0 2 4 0 14 0
pink	0 6 0
black	0 12 3 1 0

Table 6.17: Assignments for nodes with known location.

Because there is considerable overlap in the regions that can be serviced by the sites at S and T, the constraint on the establishment time for the sites is relaxed so that S is co-located with node 6. This is feasible because node 6 lies on a single-noded route of such a distance that co-locating a site there will not leave any of the grid exposed, and the single-noded route can still be traversed in time $t=(t_{max}-t_{rev})$. Such a relaxation is not appropriate for the site at T because the green vehicle of Table 6.17 is unable to traverse the route, 0 11 9 8 5 0, feasibly without departing the depot at $t=0$. The routing strategy then incorporates the following decision, Procedure 6.5.

IF	the site at S is not used by node 15, the teal vehicle is sent to service node 6 the pink vehicle is not deployed
ELSE	the teal vehicle is sent to service node 15 the pink vehicle is deployed at $t=t_{rev}$
ENDIF	

Procedure 6.5: Co-locating a geographic coordination site and a node.

This method combines the geographic and dedicated vehicle approaches, and eliminates route failure by accepting a possible increase in fleet size. It differs from the solutions for DC1 and DC2 in that the coordination sites are identified before the initial routes are constructed. Subsequent coordination may still occur when the initial routes are constructed.

DC3	
Step One:	Identify suitable coordination sites for geographic coverage.
Step Two:	Route and schedule all nodes with known attributes as well as the sites identified in Step One (which have provisional demand equal to the demand of the unknown node).
Step Three:	At $t=t_{rev}$, reorganise routes to accommodate the new node. Coordination sites are emptied before the end of the work cycle.

Procedure 6.6: Combination approach for DC3.




Figure 6.11: Upper: Red discs indicate the extent of 10 units from the depot. Lower: Blue lines indicate the coverage of a coordination site established at S at $t=10$; a site established at T at $t=10$ covers all but the two indicated vertices.

6.6.3 Initial route construction

Routes are constructed so that they include the 14 nodes with advance requests and the coordination sites to be established for geographic purposes, using any suitable construction method. A typical result is shown in Table 6.18 and Figure 6.12 in which no coordination has been applied other than the establishment of sites at S and T.

Route#	Distance	Load	Route
1	36	193	0 7 10 13 0
2	50	167	0 11 9 8 5 0
3	20	194	0 2 4 0
4	28	119	0 6 0
5	40	118	0 12 3 1 0
6	24	200	0 14 0
7	20	50	0 S 0
8	20	50	0 T 0
	238	1091	

Vehicle	Distance	Routes(s)	Nodes
blue	36	1	0 7 10 13 0
green	50	2	0 11 9 8 5 0
pink	28	4	0 6 0
black	40	5	0 12 3 1 0
red	44	3 6	0 2 4 0 14 0
teal	20	7	0 S 0
violet	20	8	0 T 0

Table 6.18: Initial routes using coordination sites at S and T but otherwise without coordination.




Figure 6.12: Initial routes generated with geographic coordination sites, S and T.

6.6.4 Behaviour at $t=t_{rev}$

At $t=t_{rev}$, the actual location of node 15 is revealed, Procedure 6.5 is enacted and the routes may be modified to accommodate the new information. The positions of all vehicles and the identities of all serviced nodes, shown in Figure 6.13, become inputs into the design of new routes.

From Figure 6.13 it can be seen that a vehicle positioned at the site at S is unable feasibly to service node 15 where that node is actually located only by accessing one of the edges coloured dark grey; each of the dark grey edges can be traversed feasibly by the violet vehicle from the coordination site at T. In the event that node 15 is serviced by the teal vehicle, the violet vehicle becomes available for servicing other nodes. Since there is no advantage in the violet vehicle servicing nodes 5, 8 and 9, the only candidates are nodes 1, 6, 13 and 14. By deploying the violet vehicle fully laden, it is able to service node 14 on its return to the depot, thereby removing the requirement for the red vehicle to be multi-routed and reducing the total distance travelled. On the other hand, if node 15 is serviced by the violet vehicle, the teal vehicle may service node 6 and remove the requirement for the pink vehicle entirely.




Figure 6.13: Situation at $t=t_{rev}$. Coloured lines indicate edges traversed by vehicles; serviced nodes are identified by filled discs. Heavy grey lines indicate edges that cannot be traversed by the teal vehicle from $t=10$; all edges can be traversed by the teal vehicle from $t=10$.

6.6.5 Incorporation of coordination

At $t=0$, two vehicles (teal and violet) are deployed from the depot bound for coordination sites, and one of them eventually will service node 15. At $t=t_{rev}=10$, the actual location of node 15 is revealed and one of the following schemes is adopted.

6.6.5.1 Routing scheme A

As shown in Table 6.19 and Figure 6.14, the teal vehicle is assigned to service node 15; it services node 6 before or after node 15 as appropriate. The other vehicles modify their routes: the pink vehicle is not deployed; the violet vehicle services node 14, thereby removing the return to depot for the red vehicle; the red vehicle services route 3 and returns to the depot; the black, blue and green vehicles continue without alteration.

Vehicle	Distance	Nodes
blue	36	0 7 10 13 0
green	50	0 11 9 8 5 0
black	40	0 12 3 1 0
red	20	0 2 4 0
teal	$14 + \text{dist}(6,15) + \text{dist}(15,0)$	0 S 6 15 0
violet	30	0 T 14 0
$190 + \text{dist}(6,15) + \text{dist}(15,0)$		

Table 6.19: Teal vehicle services node 15.




Figure 6.14: Routing scheme A – the teal vehicle services node 15. The teal route continues to node 15 (wherever that is) and then returns to the depot.

6.6.5.2 Routing scheme B

As shown in Table 6.20 and Figure 6.15, the violet vehicle is assigned to service node 15. The other vehicles modify their routes: the pink vehicle is not deployed; the teal vehicle services node 6; the red, green, black and blue vehicles continue as planned.

Vehicle	Distance	Nodes
blue	36	0 7 10 13 0
green	50	0 11 9 8 5 0
black	40	0 12 3 1 0
red	44	0 2 4 0 14 0
teal	28	0 S 6 0
violet	$10 + \text{dist}(T, 15) + \text{dist}(15, 0)$	0 T 15 0
	$208 + \text{dist}(T, 15) + \text{dist}(15, 0)$	

Table 6.20: Violet vehicle services node 15.




Figure 6.15: Routing scheme B – the violet vehicle services node 15. The violet route continues to node 15 (wherever that is) and then returns to the depot.

6.6.5.3 Routing scheme C

As shown in Table 6.21 and Figure 6.16, a vehicle other than the teal or violet vehicles is assigned to service node 15. Of the five vehicles, only the black and blue vehicles are in a position to divert from their planned routes in order to service a node with demand 50 because the pink vehicle is not deployed; the red vehicle, having serviced nodes 2 and 4, does not have the capacity; and there is no advantage in using the green vehicle over the violet vehicle. Regardless of whether the blue or black vehicles services node 15, the pink vehicle is not deployed because the teal vehicle services node 6, and the violet vehicle services node 14, thereby removing the inter-route return to depot for the red vehicle.

Vehicle	Distance	Nodes
blue	$10 + \text{dist}(10,15) + \text{dist}(15,0)$	0 7 10 15 0
green	50	0 11 9 8 5 0
black	40	0 12 3 1 0
red	46	0 2 4 0 13 0
teal	28	0 S 6 0
violet	30	0 T 14 0
$204 + \text{dist}(10,15) + \text{dist}(15,0)$		

Table 6.21: Blue vehicle services node 15. Scheme C1.




Figure 6.16: Routing scheme C1 – the blue vehicle services node 15. The blue route continues to node 15 (wherever that is) and then returns to the depot.

Unlike the other vehicles, the black vehicle has the capacity to accommodate the demand for node 15 without losing any of its customer nodes to another vehicle, and is able to insert the new node into the route at one of two places. In doing so, the teal vehicle is available to service node 14 and thus remove the inter-route return to depot for the red vehicle; it is the distance reduction associated with this removal that renders this as a desirable option, as will be seen in §6.6.5.4.

Vehicle	Distance	Nodes
blue	36	0 7 10 13 0
green	50	0 11 9 8 5 0
black	21 + dist(1, 15) + dist(15,0)	0 12 3 1 15 0
	or	or
	29 + dist(3,15) + dist(15,1)	0 12 3 15 1 0
red	20	0 2 4 0
teal	28	0 S 6 0
violet	30	0 T 14 0
	185 + dist(1,15) + dist(15,0)	
	or	
	193 + dist(3,15) + dist(15,1)	

Table 6.22: Black vehicle services node 15. Scheme C2.

6.6.5.4 Routing strategy

By comparing total distances in Tables 6.19-22, the following routing strategy, Procedure 6.7, is formulated based on the location of node 15. Figure 6.17 simplifies the routing strategy by indicating which scheme should be used dependent on the actual location of node 15. The uses to which vehicles other than that servicing node 15 are put determines the suitability of any scheme. Because use of the teal or black vehicles allows the removal of the return to depot for the red vehicle, schemes A and C2 are the most widely preferred options. Retention of the return to depot involving node 14 is so undesirable that the only scheme using it, B, is not employed, and scheme C1, which replaces that return to depot with a shorter one, is employed only once from one vertex which would as readily use scheme A.

6.6.6 Conclusions

Dynamic facility location problems of the type DC3 can be solved effectively using pre-processing and scenario planning where the number of dynamic variables is low. The uncertainty surrounding future conditions imbues problems with complexity and thus improves the prospect of heuristic solutions over optimal ones. Daskin (see [DHM90]) recommends delaying decision making in order to gather more information and thus improve forecasting; this concept has been used in the discussion above by delaying the departure of one vehicle in order to better utilise another.

While the grid could have been covered by one extra vehicle, using two creates a robustness for the solution which one vehicle would not have rendered, without greatly increasing the fleet size. The implications of this for the situation in which more than one node exhibits DC3 qualities is discussed later.

Of the two vehicles added to the fleet to establish geographic coordination points, one took the place of another vehicle (pink) and the other was not used to service the unknown node. The fleet increase of one was used in the modification of routes after $t=t_{rev}$ to remove or improve an inter-route return to depot. The replacement of one vehicle by another is facilitated by the delayed departure of the replaced vehicle.

The maximum total distance, under schemes A and C2, of 225 occurs when node 15 is located at the point X in Figure 6.17. Had the system been fully deterministic, the routing scheme would have been that shown in Figure 6.18 and Table 6.23, using a different coordination point, W, and yielding a distance of 208 and a fleet size of five.

The fleet increase of one and the distance increase demonstrated in the dynamic system with coordination is a measure of the dynamism and the need to prepare for the eventual location of the unknown node at any point on the grid.

A similar technique could be used in a system with more than one node with unknown location, as discussed later.

Vehicle	Distance	Nodes
blue	36	0 7 10 13 0
green	50	0 W 11 9 8 5 0
black	40	0 12 3 1 0
pink	48	0 6 0 2 4 0
red	34	0 14 W 15 0
	208	

Table 6.23: Deterministic routing for system with node 15 at vertex X (in Figure 6.18) and demand 50.




Figure 6.18: Deterministic routing scheme for a node at X with demand 50.

6.7 DC4 Only location is known

The variants of DC4 problems depend on when the demands of nodes become known. Where the input parameters vary deterministically over time, the problems can be classed as dynamic (see [LMR96]); where the input parameters vary probabilistically over time, the problems are stochastic. Real-world scenarios typically contain stochastic and dynamic elements ([BS96]), as routing involves a continuous process of receiving demands and creating routes using delayed deployment of vehicles. Although considerable effort has been applied to solutions in the Euclidean scene, the rectilinear grid has received little attention and the involvement of coordination does not appear to have been addressed. This study considers the dynamic problem by pre-processing comparable versions of the static version and developing a routing strategy based on their outcomes.

6.7.1 Problem generation

In order to keep the example as simple as possible without loss of generality, the system created is similar to that used in other parts of this chapter. A single depot/fifteen node system identical to that above (see Figure 6.5), with data in Appendix F, is generated with the exception that fourteen nodes have pre-determined demands and node 15 has a demand that will become apparent at time $t=t_{rev}$, although that time is unknown at the start of the work cycle.

6.7.2 Problem solution

Using Procedure 6.8, initial routes are constructed using a preliminary estimation for demand, calculated as in DC2 above. For each of several time values, a coordination plan is made and the estimated demand is modified to be as large as possible within the plan. When the actual demand is revealed at $t=t_{rev}$ the coordination plan for the value of t close to t_{rev} is considered for adoption. Usually, a coordination plan for smaller values of t can be modified for larger values. Each coordination plan is determined as for DC2 where the time attribute had value 15.

Lund ([LMR96] page 8) reaffirms [Ps95] in stating that

“Near term events are more important than long term events”

and discourages resource allocation for distant requests that may be superseded by intermediate term events. With this in mind, schemes for a particular value of t are

designed initially to be self-contained and will not pave the way for later schemes. When they have been compiled, an umbrella strategy will be generated (in §6.7.9) that will accommodate all the schemes.

DC4	
Step One	Initial routes are constructed and scheduled using an estimated value for the unknown demand.
Step Two	The maximum service time for the DC4 node, u , is determined by $\text{maxService} = \text{maxVehDist} - \text{dist}(u,0) $.
Step Three	The time interval from $t=0$ to $t=\text{maxService}$ is divided into sensible intervals.
Step Four	For each time interval selected in Step Three follow Steps Five – Seven.
Step Five	Adjust the routes so that node u is not departed from before the time at which its true demand is revealed, $t=t_{rev}$.
Step Six	Establish coordination points at likely sites so that they will be available for use when the mystery demand is revealed.
Step Seven	At the time t_{rev} , adjust routes to accommodate, if possible, the known demand. Ensure coordination sites are emptied before the end of the work cycle.
Step Eight	Modify the schemes to provide consistency so that earlier schemes allow later schemes to exist, and later schemes acknowledge the existence of earlier schemes.

Procedure 6.8: DC4. Only the location is known.

6.7.3 Demand estimation

As in DC2, the demand is estimated stochastically using the demand history of the unknown nodes and, if necessary, the demand histories of other nodes. As in the examples above, node 15 has had only one previous demand (of value 50), so the preliminary estimate of Demand(15) is 50, but this is increased after the initial routes have been constructed to be the maximum that can be accommodated by the vehicle assigned to service node 15.

6.7.4 Choosing time values

As in §6.4.1, the maximum distance constraint infers a maximum time constraint and $t_{\max}$ has the same magnitude as $\maxVehDist$. Consequently, times vary from $t=0$ to $t=50$ but the maximum time at which node 15 can be serviced feasibly has value

$$t_{\max,15} = |\maxVehDist - \text{dist}(15,0)|$$

which allows the node to be serviced and return to the depot without violating the time constraint. Since node 15 is 10 units away from the depot by the shortest path, there is no reason to consider any time greater than $t=40$.

Because the demand for node 15 may be revealed at any time after the start of the work cycle, there is no limit on the minimum time at which node 15 can be serviced feasibly.

The coordination plan for $t=15$ was constructed in the discussion of DC2 above and will be re-used here. By considering times $t=25$ and $t=35$, no revealed time will be more than 5 time units away from one of the considered scenarios.

Generally, a suitable initial value for t is one at which all vehicles are deployed on their initial routes. An alternative method of selecting values of t is discussed in §6.10.3, and involves adaptive time values rather than the fixed values used here.

6.7.5 Initial route construction

As in the discussions above, initial routes are constructed using any system that produces the best feasible routes from nearest/farthest insertion / largest/nearest neighbour with a neighbourhood size of 3. The routing scheme is shown in Table 6.9 and Figure 6.6.

6.7.6 Time $t=15$ scheme

The discussion on DC2 above yielded routing schemes A-D for the time $t_{rev} = 15$, and they are re-used here.

6.7.7 Time $t=25$ scheme

The location of all vehicles is determined at $t=25$, shown in Figure 6.19. Most nodes have already been serviced, leaving only nodes 5, 14 and 15, and all vehicles are still en route. Three routing schemes are created to provide stock to node 15 in

varying amounts, and each must accommodate the requirements of the others so that any of them may be chosen from the same initial configuration.




Figure 6.19: Situation at $t=25$. Coloured lines indicate edges traversed by vehicles; serviced nodes are identified by filled discs.

6.7.7.1 Routing scheme E

The black vehicle, to which node 15 has been assigned, is capable of carrying 82 units beyond its requirements for the other nodes that it services, all of which it may deliver to node 15. Consequently, if the demand of node 15 is no greater than 82, coordination will not be required for the purposes of servicing that node, although it may be used for other reasons. The routing scheme is shown in Table 6.24 and is consistent with other schemes for $t=25$.

Vehicle	Distance	Routes(s)	Nodes
blue	50	1	0 V 0 7 10 13 0
green	50	2	0 11 9 8 5 0
pink	28	4	0 6 0
black	44	5	0 12 3 1 15 V 0
red	44	3 6	0 2 4 0 14 0
	216		

Table 6.24: Routing scheme E.

6.7.7.2 Routing scheme F

By leaving the depot fully laden, the pink vehicle is able to carry 81 surplus stock units which it may deliver to node 15 before returning to the depot. Node 15 receives 82 units from the black vehicle and up to 81 units from the pink vehicle. The routing scheme, shown in Table 6.25, thus provides up to 163 units to node 15 and, although no coordination site is required under this scheme, one is established to provide consistency with routing scheme G.

Vehicle	Distance	Routes(s)	Nodes
blue	50	1	0 V 0 7 10 13 0
green	50	2	0 11 9 8 5 0
pink	40	4	0 6 15' 0
black	44	5	0 12 3 1 15" V 0
red	44	3 6	0 2 4 0 14 0
	228		

Table 6.25: Routing scheme F.

6.7.7.3 Routing scheme G

With a distance of 36, the blue vehicle is able to establish a type 1 coordination site up to 7 units from the depot. The point V in Figure 6.19 indicates a distance of 7 units along a shortest path from the depot to node 15, within reach of other vehicles, and is the best location for such a site, were it to be used for supplying the demand of node 15. Rather than supplementing the stock delivered to node 15 by the black vehicle, the use of a type 1 coordination site by a vehicle that will, the use of a type 1 coordination site by a vehicle that will service no other nodes means

that up to maxVehCap of material can be delivered to the node with unknown demand, and consequently the stock carried by the black vehicle for node 15 is not needed. An appropriate routing scheme is shown in Table 6.26. Since this scheme requires the early establishment of a coordination site, other schemes for $t=25$ must also provide that site in order to maintain consistency of potential, and thus scheme G supersedes scheme F.

Vehicle	Distance	Routes(s)	Nodes
blue	50	1	0 V 0 7 10 13 0
green	50	2	0 11 9 8 5 0
pink	40	4	0 6 V 15 0
black	40	5	0 12 3 1 0
red	44	3 6	0 2 4 0 14 0
	224		

Table 6.26: Routing scheme G.

6.7.7.4 Routing scheme H

If the actual demand of node 15 is zero, it need not be serviced and may be eliminated from the black vehicle’s route. In doing so, schemes E and F must be accommodated so the site at V must still be established. The scheme is shown in Table 6.27.

Vehicle	Distance	Routes(s)	Nodes
blue	50	1	0 V 0 7 10 13 0
green	50	2	0 11 9 8 5 0
pink	28	4	0 6 0
black	40	5	0 12 3 1 V 0
red	44	3 6	0 2 4 0 14 0
	212		

Table 6.27: Routing scheme H.

6.7.7.5 Routing strategy $t=25$

The routing scheme that is adopted depends on the value of the demand of node 15 that is revealed at $t=t_{rev}$. In the absence of earlier routing strategies, that shown in Table 6.28 would be applied, but since a strategy exists for $t=15$, that earlier strategy must be accommodated; the resultant strategy is discussed later.

Actual value of demand(15)	Routing scheme	Total distance
0	H	212
$0 < \text{demand}(15) \leq 82$	E	218
$82 < \text{demand}(15) \leq \text{maxVehCap}$	G	224

Table 6.28: Routing strategy for $t_{rev}=25$.

6.7.8 Time $t=35$ scheme

In the absence of other coordination schemes, by $t=35$, the pink vehicle has returned to the depot, the blue vehicle is nearly there, and the red and green vehicles are returning; the black vehicle has been waiting at node 15 for 2 time units, as shown in Figure 6.20. With only node 15 to service, the black vehicle delivers the goods and returns to the depot. If the demand of node 15 is greater than that which the black vehicle alone can deliver, coordination is employed. Consistency between schemes for different values of t will be applied in §6.7.9.




Figure 6.20: Situation at $t=35$. Coloured lines indicate edges traversed by vehicles; serviced nodes are identified by filled discs.

6.7.8.1 Routing schemes

As in the above sections, schemes are required for the instances where the demand of node 15 is revealed as zero, no greater than 82 (in which case the black vehicle is able to accommodate fully the demand) and between 82 and maxVehCap . Because the routes are almost completely traversed by $t=35$, the routing schemes are the same as those for $t=25$.

6.7.8.2 Routing strategy $t=35$

In the absence of earlier routing strategies, the routing strategy for $t=35$, shown in Table 6.29, is the same as that for $t=25$.

Actual value of demand(15)	Routing scheme	Total distance
0	H	212
$0 < \text{demand}(15) \leq 82$	E	218
$82 < \text{demand}(15) \leq \text{maxVehCap}$	G	224

Table 6.29: Routing strategy for $t_{\text{rev}}=35$.

6.7.9 Routing strategy

Having devised strategies for different values of t , consistency must now be applied so that the involvement of earlier schemes does not handicap or prevent the application of later schemes. This is done by first recognising, for each value of t discussed, that part of the scheme that is enacted in the first t time units, as shown in Table 6.30.

Vehicle	$t=15$	$t=25$	$t=35$
blue	0 7 10	0 Z 0 7 10	0 Z 0 7 10
green	0 W 11	0 11 9 8	0 11 9 8 5
pink	0 X	0 6	0 6 Z 15
black	0 12 3	0 12 3 1	0 12 3 1 Z
red	0 2 4	0 2 4 0	0 2 4 0 14

Table 6.30: Rendering schemes consistent.

Then, each earlier scheme is incorporated into each later scheme, and later schemes are incorporated into earlier schemes, and final schemes created, as shown in Table 6.31 with the strategy shown in Table 6.32. Note that scheme C is

no longer viable because the added distance incurred by the creation of the coordination site at V violates the distance constraint.

	A	B
blue	0 V 0 7 10 13 0	0 V 0 7 10 13 0
green	0 W 11 9 8 5 W 0	0 W 11 9 8 5 W 0
pink	0 X 6 0	0 X 6 0
black	0 12 3 1 15 X 0	0 12 3 1 Y 15 V 0
red	0 2 4 0 14 0	0 2 4 X Y 0
	D	E
blue	0 V 0 7 10 13 0	0 V 0 7 10 13 0
green	0 W 11 9 8 5 W 0	0 W 11 9 8 5 W 0
pink	0 X 6 0	0 X 6 0
black	0 12 3 1 X 0	0 12 3 1 15 V 0
red	0 2 4 0 14 V 0	0 2 4 0 14 0
	G	H
blue	0 V 0 7 10 13 0	0 V 0 7 10 13 0
green	0 W 11 9 8 5 W 0	0 W 11 9 8 5 W 0
pink	0 X 6 V 15 0	0 X 6 0
black	0 12 3 1 0	0 12 3 1 V 0
red	0 2 4 0 14 0	0 2 4 0 14 0
	I	J
blue	0 10 7 13 0	0 10 7 13 0
green	0 11 9 8 5 0	0 11 9 8 5 0
pink	0 6 0	0 6 0
black	0 12 3 1 0	0 12 3 1 15 0
red	0 2 4 0 14 0	0 2 4 0 14 0
	K	
blue	0 V 0 10 7 13 0	
green	0 11 9 8 5 0	
pink	0 6 V 15 0	
black	0 12 3 1 0	
red	0 2 4 0 14 0	

Table 6.31: Consistent schemes.

time, t	Demand(15), D			
	$D=0$	$0 < D \leq 82$	$83 < D \leq 169$	$169 < D \leq \text{maxVehCap}$
$t=0$	I	J	K	K
$0 < t \leq 20$	D	A	B	G
$20 < t \leq 30$	H	E	G	G
$30 < t \leq 40$	H	E	G	G
$40 < t$	H	H	H	H

Table 6.32: Routing strategy for DC4 example.

6.7.10 Limitations and accommodations

Inferring a maximum time constraint guarantees route failure whenever $t_{\text{rev}} > 40$. If this constraint were to be relaxed, the above scheme could be modified by allowing a vehicle to wait at the node with unknown demand until $t = t_{\text{rev}}$.

By performing DC2 calculations at arbitrary times, the system adopts the limitations of the DC2 system and does not guarantee that the times chosen are the most suitable. Instead of using regular intervals, another system might use times that were responsive to the system under investigation so that, for example, the $t=25$ scheme for $20 < t \leq 30$ may be appropriate for $18 < t \leq 32$ after which another scheme is more appropriate than the $t=35$ scheme above.

Rather than solving DC2 problems at different times in isolation, better results may be obtained by a more integrated approach which would remove the need to provide consistency between schemes. The complexity of the current approach may make problems with large numbers of nodes more suitable to other forms of solution.

6.7.11 Conclusions

Coordination is a valid technique for vehicle routing where the demand of a particular node is unknown until some unpredictable time after the vehicles have left the depot. It consists of estimating the demand, modifying the estimation, and developing coordination schemes at a selection of different times, in order to be able to judge the worth of coordination when the actual demand is revealed. The situation involving more than one node with DC4 qualities is discussed in §6.11.

6.8 DC5 Only demand is known

In DC5 problems, all demands are known in advance but the location of at least one node and the time at which that location will be revealed is unknown at the start of the work cycle. As in DC4 problems, these can be approached by solving DC3 problems at a selection of different times and then adjusting the final scheme to provide compatibility between schemes developed at different times.

6.8.1 Problem generation

As in §6.6 DC3, the 15 node system of Appendix F is generated with the exception that the location of node 15 is unknown at the start of the work cycle but will be revealed at some unknown time, $t=t_{rev}$; assume that node 15 has demand 100. This demand has been chosen because it is too large to be accommodated on any one vehicle and allows coordination to be demonstrated as a viable solution technique. In keeping with other problems in this study, the node will not be co-located with any other node.

6.8.2 Problem solution

Because the location of node 15 is unknown, the distance between it and the depot is not determined at the start of the work cycle, and the latest time at which a vehicle can leave the node feasibly for the depot, t_{max} , is not known until $t=t_{rev}$. Therefore, the value of t_{max} is set initially at the magnitude of $maxVehDist$, 50, but will be altered at $t=t_{rev}$. After constructing initial routes, the time period, $t \in [0, 50]$, is divided into smaller intervals and a DC3 system is solved for each of these. Then all the schemes are combined to produce a consistent routing strategy that allows early schemes to exist at the same time as the preparation for later schemes, in case they are required.

DC5	
Step One	Initial routes are constructed and scheduled.
Step Two	The time interval from $t=0$ to $t=50$ is divided into sensible intervals.
Step Three	For each time interval selected in Step Two, solve the corresponding DC3 problem.
Step Four	Modify the schemes to provide consistency so that earlier schemes allow later schemes to exist, and later schemes acknowledge the existence of earlier schemes.

Procedure 6.9: DC5. Only the demand is known.

The DC3 problem in Step Three of Procedure 6.9 will be solved in the same manner as in §6.6, at times $t=10, 20, 30$ and 40 .

6.8.3 Initial route construction

As in other examples in this study, the initial routes are chosen from the best solutions from a selection of smallest neighbour/ farthest insertion, smallest neighbour/nearest insertion, largest neighbour/farthest insertion, and largest neighbour/nearest insertion, all improved using k-opt, and constructed without the unknown node. The results are shown in Table 6.33 and Figure 6.21.

Vehicle	Distance	Nodes
blue	36	0 7 10 13 0
green	50	0 11 9 8 5 0
black	40	0 12 3 1 0
pink	20	0 4 2 0 6 0
red	24	0 14 0
198		

Table 6.33: Initial routes for DC5 example.

No one vehicle is able to accommodate the demand for node 15, although several may combine in a coordination scheme to achieve that end. The scheme chosen will depend on the location of node 15 and the time at which its location is revealed.




Figure 6.21: Initial routes.

6.8.4 $t_{rev}=10$

At $t=10$, all vehicles are en route, with coordination sites to be established at X, Y and W as shown in Figure 6.21. Several schemes are presented to accommodate the demand of node 15 using coordination, following which a routing strategy will be developed. The routing schemes described below rely on coordination sites being established at points X, Y and W and being dis-established before the end of the work cycle. To preserve compatibility, each scheme must use the same route fragments for the first 10 time units, at which time the choice of scheme adopted by the routing strategy determines the later fragments.

6.8.4.1 Routing scheme A

Scheme A, shown in Table 6.34, utilises the coordination site at point W; the red vehicle dis-establishes the unused coordination site at X on its return to the depot. The black vehicle delivers the demand for node 1 to the coordination site at Y, and travels to W; there it collects the 51 units dropped there by the green vehicle, before travelling to node 15, delivering its demand and returning to the depot. The red vehicle, having serviced node 14, collects the demand for nodes 1, 3 and 12 from the coordination site at Y, and services them before returning to the depot. The vertices covered by this scheme are shown in Figure 6.22.

Vehicle	Distance	Nodes
blue	46	0 $X^{del\ 200}$ 0 7 10 13 0
green	50	0 11 $W^{del\ 51}$ 9 8 5 0
black	$16 + \text{dist}(W,15) + \text{dist}(15,0)$	0 $X^{del\ 33}$ $Y^{del\ 118}$ $W_{col\ 51}$ 15 0
red	46	0 14 $Y_{col\ 118}$ 1 3 12 X_{col} 0
pink	48	0 2 4 0 6 0
$206 + \text{dist}(W,15) + \text{dist}(15,0)$		

Table 6.34: Routing scheme A for DC5 example.




Figure 6.22: Vertices for which scheme A is effective at $t=10$.

6.8.4.2 Routing scheme B

The blue vehicle establishes a type 1 coordination site at X containing maxVehCap of product, some of which is used by the red vehicle to service node 15 and some used by the pink vehicle to obviate a return to depot and service node 6 directly. Because the site at W is unused, it is emptied by the green vehicle on its return to the depot. This scheme, shown in Table 6.35, is appropriate in the event that node 15 is identified at $t=10$ as being located at one of the vertices in Figure 6.23, indicated by a filled disc.

Vehicle	Distance	Nodes
blue	46	0 $X^{\text{del } 200}$ 0 7 10 13 0
green	50	0 11 $W^{\text{del } 51}$ 9 8 5 $W_{\text{col } 51}$ 0
black	43	0 $X^{\text{del } 33}$ Y 1 3 12 0
red	$19 + \text{dist}(X,15) + \text{dist}(15,0)$	0 14 $X_{\text{col } 100}$ 15 0
pink	40	0 2 4 $X_{\text{col } 133}$ 6 0
$198 + \text{dist}(X,15) + \text{dist}(15,0)$		

Table 6.35: Routing scheme B for DC5 example.




Figure 6.23: Vertices for which scheme B is effective.

6.8.4.3 Routing strategy

Taken together, schemes A and B cover the entire grid and are identical from $t=0$ until $t=10$, at which time the black vehicle either deposits 118 stock units at the coordination site at Y, or backtracks to service nodes 1, 3 and 12, and the red vehicle is en route to service node 14. The scheme selected depends on the actual location of node 15 that is revealed at $t=10$, as shown in Procedure 6.10. If the location is not revealed at this time, the routing strategy for a later time will be enacted.

Routing Strategy, $t=10$	
IF	the location of node 15 has been revealed
THEN	
	$D := \min(8 + \text{dist}(W,15), \text{dist}(X,15))$
	CASE D
	$8 + \text{dist}(W,15)$ scheme A
	$\text{dist}(X,15)$ scheme C
	ENDCASE
ELSE	use the strategy relating to a later time
ENDIF	

Procedure 6.10: Routing strategy based on the location of node 15 at $t=10$.

6.8.5 $t_{\text{rev}}=20$

If the location of node 15 has not been revealed at $t=10$, the schemes for $t=20$ must be prepared for use and the routing strategy in Procedure 6.10 abandoned. To preserve backwards compatibility, schemes for $t=20$ must accommodate the earlier schemes up to the point of their abandonment, that is, $t=10$. Consequently, all schemes for $t=20$ include the route fragments:

0 $X^{\text{del } 200}$ 0 7 10

0 11 $W^{\text{del } 51}$

0 $X^{\text{del } 33}$ Y

0 14

0 2 4

where the coordination sites X, Y and W hold 233, 0 and 51 stock units respectively.

Schemes A and B cannot be applied simultaneously because they involve the same nodes, so if one were to be accepted for continued use at $t=20$ the other would be rejected at $t=10$. Although scheme A covers most of the grid, its prime value is that its coverage includes the outcrop that includes nodes 5, 8 and 9, a region that existing vehicles do not reach feasibly. Scheme B also covers vertices that are not accessed feasibly, but the most distant vertex in this area is closer to the depot than the most distant vertex in the outcrop above. Also, by accepting scheme A for use at $t=20$, the vehicle servicing the outcrop region is already in the vicinity (at point S) at that time, whereas scheme B does not allow for the addition of another node without assistance. Consequently, scheme A is continued through $t=20$ as scheme C, and scheme D is employed to cater for the situation where node 15 is revealed to be located at one of the vertices not covered by scheme C.

6.8.5.1 Routing scheme C

Scheme C, shown in Table 6.36, adopts the principles of scheme A and applies them at $t=20$, which reduces the coverage from that of scheme A to that of Figure 6.24. If the location of node 15 is not revealed at $t=10$, scheme C is implemented and reassessed at $t=20$. Because the black vehicle will have reached point S at that time, fewer vertices are covered by scheme C, as shown in Figure 6.24. If node 15 is revealed as being located at a vertex not covered by scheme A, scheme B will be implemented, otherwise the black vehicle will remain stationed at point S (scheme D).

Vehicle	Distance	Nodes
blue	46	0 X_{del}^{200} 0 7 10 13 0
green	50	0 11 W_{del}^{51} 9 8 5 0
black	$20 + \text{dist}(S,15) + \text{dist}(15,0)$	0 X_{del}^{33} Y_{del}^{118} W_{col}^{51} S 15 0
red	46	0 14 Y_{col}^{118} 1 3 12 X_{col}^{114} 0
pink	40	0 2 4 X_{col}^{119} 6 0
$202 + \text{dist}(S,15) + \text{dist}(15,0)$		

Table 6.36: Routing scheme C for DC5 example.




Figure 6.24: Vertices for which scheme C is effective at $t=20$.

6.8.5.2 Routing scheme D

Scheme D, shown in Table 6.37, is based on a dedicated vehicle approach, whereby another vehicle is dispatched to service the relevant nodes, while the black vehicle (that would be servicing node 15 under scheme C) is returned to the depot. This scheme can be used for vertices at any point on the grid, but the increase in fleet size may render its implementation undesirable unless the accompanying reduction in total distance is acceptable to the dispatcher.

Vehicle	Distance	Nodes
blue	46	0 $X^{\text{del } 200}$ 0 7 10 13 0
green	50	0 11 $W^{\text{del } 51}$ 9 8 5 0
black	32	0 $X^{\text{del } 33}$ $Y^{\text{del } 118}$ $W_{\text{col } 51}$ S 0
red	46	0 14 $Y_{\text{col } 118}$ 1 3 12 $X_{\text{col } 114}$ 0
pink	40	0 2 4 $X_{\text{col } 119}$ 6 0
violet	$2 \times \text{dist}(15,0)$	0 15 0
$214 + 2 \times \text{dist}(15,0)$		

Table 6.37: Routing scheme D for DC5 example.

6.8.5.3 Routing scheme E

Because of the difficulty in servicing node 15 were it to be located on the outcrop, the black vehicle may be stationed at S until either node 15 is identified as being there (in which case scheme C is implemented), or node 15 is identified as being located elsewhere (when another scheme will be implemented), or time expires and all vehicles are returned to the depot. Scheme E, shown in Table 6.38, allows the black vehicle to wait at the vertex S.

Vehicle	Distance	Nodes
blue	46	0 $X^{\text{del } 200}$ 0 7 10 13 0
green	50	0 11 $W^{\text{del } 51}$ 9 8 5 0
black	20	0 $X^{\text{del } 33}$ $Y^{\text{del } 118}$ $W_{\text{col } 51}$ S
red	46	0 14 $Y_{\text{col } 118}$ 1 3 12 $X_{\text{col } 114}$ 0
pink	40	0 2 4 $X_{\text{col } 119}$ 6 0
	202	

Table 6.38: Routing scheme E for DC5 example.

6.8.5.4 Routing strategy

If the location of node 15 was not revealed before $t=20$, one of schemes C-E will be implemented. Schemes C and D are initiated exclusively where only one of them covers the vertex containing node 15, and, where they both cover the vertex, a decision is made as to which will be used based on the actual location of node 15, as in Procedure 6.11. If the location is not evident at $t=20$, scheme E is used to prepare the routes for a later routing strategy.

Routing Strategy, $t=20$	
IF	the location of node 15 has been revealed
THEN	
	$D := \min(\text{dist}(S, 15), 12 + \text{dist}(15, 0))$
	CASE D
	dist(S,15) scheme C
	12 + dist(15,0) scheme D
	ENDCASE
ELSE	scheme E followed by the routing strategy for a later time
ENDIF	

Procedure 6.11: Routing strategy based on the location of node 15 at $t=20$.

6.8.6 $t_{rev}=30$

If the location of node 15 has not been revealed at $t=20$, scheme E and the $t=30$ routing strategy will be implemented. At $t=30$, only nodes 3, 12, 13 and 15 remain to be serviced, as shown in Figure 6.25, with the blue vehicle en route to node 13, the red vehicle en route to node 3, the black vehicle stationary at S, and the pink and green vehicles returning to the depot. The coordination site at X contains 114 stock units; the sites at Y and W have been dis-established by the red and black vehicles respectively.




Figure 6.25: Vehicle locations and serviced nodes at $t=30$. Arrows indicate vehicle positions, filled discs indicate serviced nodes.

Nodes 3, 12 and 13 can be serviced by the red and blue vehicles without the addition of extra distance, since those vehicles are closer to the nodes than any alternative routing scheme would place other vehicles. Therefore, no other scheme is presented for $t=30$ than that presented for $t=20$.

6.8.6.1 Routing strategy

If the location of node 15 is revealed at $t=30$ and it is at a vertex unreachable by the black vehicle, an extra vehicle is deployed to service the node as in scheme D; if the node is located on the outcrop containing nodes 5, 8 and 9, the black vehicle services it as in scheme C; if the node is located at a vertex that can be serviced

feasibly by either scheme C or D, a decision as to which scheme to employ is made, as in Procedure 6.12. If the location is not revealed at $t=30$, scheme E is implemented until $t=38$, which has the black vehicle standing by to service node 15 and the other routes continuing without interruption, after which all vehicles must return to the depot.

Routing Strategy, $t=30$					
IF	the location of node 15 has been revealed				
THEN					
	$D := \min(\text{dist}(S,15), 12 + \text{dist}(15,0))$				
	CASE D				
	<table> <tr> <td>$\text{dist}(S,15)$</td><td>scheme C</td></tr> <tr> <td>$12 + \text{dist}(15,0)$</td><td>scheme D</td></tr> </table>	$\text{dist}(S,15)$	scheme C	$12 + \text{dist}(15,0)$	scheme D
$\text{dist}(S,15)$	scheme C				
$12 + \text{dist}(15,0)$	scheme D				
	ENDCASE				
ELSE					
	apply scheme E until $t=38$ followed by the routing strategy for $t=40$				
ENDIF					

Procedure 6.12: Routing strategy based on the location of node 15 at $t=30$.

6.8.7 $t_{\text{rev}} = 40$

At $t=38$, the black vehicle is compelled to return to the depot in order that it arrives there no later than $t=t_{\text{max}}=50$. At $t=40$, all vehicles are either returning to the depot or already there, and only node 12 remains to be serviced, as shown in Figure 6.26. Since this node lies on a shortest path between the red vehicle and the depot, and because the red vehicle is carrying enough stock to accomplish the task, it services node 3 as planned and empties the coordination site at X whilst passing through it. Although there is very little scope for route modification, some vertices can be covered feasibly without resorting to an increase in fleet size, as indicated in the schemes below.




Figure 6.26: Vehicle locations and serviced nodes at $t=40$. Arrows indicate vehicle positions, filled discs indicate serviced nodes.

6.8.7.1 Routing scheme F

The black vehicle services node 15 while returning to the depot along a shortest path from the vertex indicated by the black arrow in Figure 6.26. Because this vehicle has no surplus time available, it may not deviate from the path to service feasibly node 15, and node 15 is limited therefore to only those vertices that lie on a shortest path between the arrowed vertex and the depot. The scheme is shown in Table 6.39.

Vehicle	Distance	Nodes
blue	46	0 $X^{del\ 200}$ 0 7 10 13 0
green	50	0 11 $W^{del\ 51}$ 9 8 5 0
black	$20 + \text{dist}(S,15) + \text{dist}(15,0)$ =32	0 $X^{del\ 33}$ $Y^{del\ 118}$ $W_{col\ 51}$ S 15 0
red	46	0 14 $Y_{col\ 118}$ 1 3 12 $X_{col\ 114}$ 0
pink	40	0 2 4 $X_{col\ 119}$ 6 0
	214	

Table 6.39: Routing scheme F for DC5 example.

6.8.7.2 Routing scheme G

Under this scheme, shown in Table 6.40, the red vehicle services node 15 after emptying the coordination site at X using its 4 surplus distance units. Although the number of vertices covered by this scheme is small, it covers vertices that would otherwise require a fleet increase to service feasibly, should node 15 be located there.

Vehicle	Distance	Nodes
blue	46	0 $X^{\text{del } 200}$ 0 7 10 13 0
green	50	0 11 $W^{\text{del } 51}$ 9 8 5 0
black	32	0 $X^{\text{del } 33}$ $Y^{\text{del } 118}$ $W_{\text{col } 51}$ S 0
red	$41 + \text{dist}(X,15) + \text{dist}(15,0)$	0 14 $Y_{\text{col } 118}$ 1 3 12 $X_{\text{col } 114}$ 15 0
pink	40	0 2 4 $X_{\text{col } 119}$ 6 0
	$209 + \text{dist}(X,15) + \text{dist}(15,0)$	

Table 6.40: Routing scheme G for DC5 example.

6.8.7.3 Routing scheme H

The pink vehicle services node 15 from the depot using its 10 surplus distance units. Because it departs directly for the depot, it does not use any coordination sites. Scheme H, shown in Table 6.41, covers all vertices reached by the blue vehicle from the depot using a comparable scheme and thus supersedes such a scheme.

Vehicle	Distance	Nodes
blue	46	0 $X^{\text{del } 200}$ 0 7 10 13 0
green	50	0 11 $W^{\text{del } 51}$ 9 8 5 0
black	32	0 $X^{\text{del } 33}$ $Y^{\text{del } 118}$ $W_{\text{col } 51}$ S 0
red	46	0 14 $Y_{\text{col } 118}$ 1 3 12 $X_{\text{col } 33}$ 0
pink	$40 + 2 \times \text{dist}(15,0)$	0 2 4 $X_{\text{col } 119}$ 6 0 15 0
	$214 + 2 \times \text{dist}(15,0)$	

Table 6.41: Routing scheme H for DC5 example.

6.8.7.4 Routing strategy

Because of the limited additional manoeuvrability of vehicles at $t=40$, where node 15 is revealed as being located at most vertices, scheme D remains the only option as shown in Figure 6.27. At other vertices, one or more of the other schemes is applicable and, where more than one scheme can be applied, the choice of scheme depends on the relative sizes of the distance increase incurred by the adoption of each scheme.




Figure 6.27: Effective schemes for vertices at $t=40$.

Routing Strategy, $t=40$	
IF	the location of node 15 has been revealed
THEN	
IF	node 15 lies on a white vertex in Figure 6.26
THEN	scheme D
ELSEIF	node 15 lies on a black vertex in Figure 6.26
THEN	scheme F
ELSEIF	node 15 lies on a red vertex in Figure 6.26
THEN	scheme G
ELSEIF	node 15 lies on a pink vertex in Figure 6.26
THEN	scheme H
ELSEIF	node 15 lies on a blue vertex in Figure 6.26
THEN	
	$D := \min(\text{dist}(X, 15), 5 + \text{dist}(15, 0))$
	CASE D
	$\text{dist}(X, 15)$ scheme G
	$5 + \text{dist}(15, 0)$ scheme H
ELSEIF	node 15 lies on a yellow vertex in Figure 6.26
THEN	
	$D := \min(\text{dist}(S, 15), 12 + \text{dist}(15, 0))$
	CASE D
	$\text{dist}(S, 15)$ scheme F
	$12 + \text{dist}(15, 0)$ scheme H
ELSE	
	$D := \min(\text{dist}(S, 15), 7 + \text{dist}(X, 15), 12 + \text{dist}(15, 0))$
	CASE D
	$\text{dist}(S, 15)$ scheme F
	$7 + \text{dist}(X, 15)$ scheme G
	$12 + \text{dist}(15, 0)$ scheme H
	ENDCASE
ENDIF	
ELSE	
	route failure
ENDIF	

Procedure 6.13: Routing strategy based on the location of node 15 at $t=40$.

6.8.8 Routing strategy for DC5

Having compiled routing strategies for the state of the system as it exists at different times, an overall strategy is created that encompasses the isolated strategies above and ensures compatibility. Thus, the earlier stages of the strategy prepare for both the imminent identification of node 15 and the much later identification. The overall strategy is shown in Procedure 6.14.

Routing strategy DC5	
IF	the location of node 15 is revealed at $t=10$
THEN	
	$D := \min(8 + \text{dist}(W, 15), \text{dist}(X, 15))$
	CASE D
	8 + dist(W,15) scheme A
	dist(X,15) scheme C
	ENDCASE
ELSEIF	the location of node 15 is revealed at $t=20$
THEN	
	$D := \min(\text{dist}(S, 15), 12 + \text{dist}(15, 0))$
	CASE D
	dist(S,15) scheme C
	12 + dist(15,0) scheme D
	ENDCASE
ELSEIF	the location of node 14 is revealed at $t=30$
THEN	
	$D := \min(\text{dist}(S, 15), 12 + \text{dist}(15, 0))$
	CASE D
	dist(S,15) scheme C
	12 + dist(15,0) scheme D
	ENDCASE
ELSEIF	the location of node 14 is not revealed and $30 < t \leq 38$
THEN	scheme E
ELSEIF	the location of node 15 is revealed at $t=40$
	IF node 15 lies on a white vertex in Figure 6.26
	THEN scheme D
	ELSEIF node 15 lies on a black vertex in Figure 6.26
	THEN scheme F
	ELSEIF node 15 lies on a red vertex in Figure 6.26

```

THEN      scheme G
ELSEIF    node 15 lies on a pink vertex in Figure 6.26
THEN      scheme H
ELSEIF    node 15 lies on a blue vertex in Figure 6.26
THEN
    D:=min(dist(X,15), 5 + dist(15,0))
    CASE D
        dist(X,15)      scheme G
        5 + dist(15,0)  scheme H
ELSEIF    node 15 lies on a yellow vertex in Figure 6.26
THEN
    D:=min(dist(S,15), 12 + dist(15,0))
    CASE D
        dist(S,15)      scheme F
        12 + dist(15,0)  scheme H
ELSE
    D:= min(dist(S,15), 7 + dist(X,15), 12 + dist(15,0))
    CASE D
        dist(S,15)      scheme F
        7 + dist(X,15)  scheme G
        12 + dist(15,0)  scheme H
    ENDCASE
ENDIF
ELSE      route failure
ENDIF

```

Procedure 6.14: Routing strategy for DC5 example.

6.8.9 Limitations and accommodations

The quality of solution depends to some extent on the choice of time intervals selected, which are initially fixed but, as demonstrated above for $30 \leq t < 40$, show some flexibility. By using times that are closer together, the system may be more responsive to change, but will require much greater processing.

The inference of a maximum time constraint is in keeping with real-world situations, although in this example, no time allowance has been made for loading/unloading and administration at nodes and coordination sites, which is unlike real-world situations. The lack of additional time costs employed in the

schemes above is without loss of generality since the concepts demonstrated are not dependent on activities at nodes. Similarly, there is no time allowance for traversing paths that include bends, intersections or high traffic volumes, scenarios that slow real-world traffic.

6.8.10 Conclusions

The pre-processing of scenarios before the deployment of vehicles allows systems to respond to change quickly, but the effectiveness of the response naturally depends on the closeness of predicted scenarios to actual events. In solving problems where only the demand of a node is known at the start of a work cycle, scenarios are created at pre-determined times and a routing strategy developed to direct the vehicles in the event that the location of the unknown node is revealed at or near those times. The effectiveness of the routing strategy depends to some extent on how far apart the actual time is to one of the predicted times, and the system will be more responsive in cases where the number of predicted times is greater, although this will greatly increase the processing required.

In developing the routing strategy, accommodation must be made in earlier schemes for the possible introduction of later schemes, should the location of the unknown node not be revealed until later. Similarly, later schemes must include the remnants of schemes that precede them. Such compatibility issues increase the computation and complexity with the number of times chosen for pre-processing.

As the number of nodes and routes increase, there may be more options and greater scope for coordination. Even the fifteen node system of the example used four coordination points (X, Y, W, S) all located in the lower half of the grid, so it could be expected that with more nodes in the upper half of the grid, there would be more route interaction there and accompanying coordination sites.

The method of solution presented in this discussion will increase in complexity with more nodes of unknown location. Nevertheless, as demonstrated in the DC5 example above, coordination is an effective tool in dealing with dynamic location problems.

6.9 DC6 Only time is known

At the start of the work cycle, some or all of the locations are unknown and some or all of the demands are unknown, but the times at which this information will be revealed is known. There are two main approaches to the solution of DC6 problems:

(a) Unknown location problem for different demands

Solution involves solving DC3 problems for different demands and integrating the results to obtain a routing strategy based on the location and demand of each unknown node. This process is computationally more intensive than the corresponding DC3 problem because, as each demand level is solved, the range of demands for which each solution is feasible is calculated, as in DC2 problems.

(b) Unknown demand problem at different locations

The DC6 problem solution involves the early establishment of coordination sites at strategic geographic locations and the development of options for different ranges of demand. It is similar to solving DC2 problems at different locations and integrating the results to obtain a suitable routing strategy, but more computationally intensive because the range of locations for each solution must be established.

Coordination can be applied in each approach since each involves an early identification of routes and vehicle assignments, thereby creating environments for the early establishment of coordination sites.

This discussion will present a technique for solving DC6 problems using coordination and will demonstrate this with reference to an example problem that is similar to the examples of earlier sections of this chapter. For the sake of simplicity, the example is set in a medium density grid with six routes, and incorporates the results of the discussions above.

6.9.1 Problem generation

As in the above sections, a fifteen node system is generated where nodes 1-14 have all attributes known at the start of the work cycle (see Appendix F). The demand and location of node 15 will be revealed at some known time, $t=t_{\text{rev}}=10$. This time has been chosen as all vehicles are still en route and there is scope for coordination to be demonstrated as an effective tool for the solution of DC6

problems, which may not be the case at later times when vehicles have exhausted their distance allowance. In addition, by choosing the same time as that used in the example for DC3 above (§6.6), the results obtained for that section can be used in this example.

6.9.2 Problem solution

Solving the DC3 problem for different demands, as in Procedure 6.15, requires the initial construction of routes and the identification of serviced nodes and vehicle locations at $t=t_{revr}$ so that the environment for coordination can be assessed and coordination sites established as early as possible. By using an estimated demand and creating an initial solution, the strategy can establish the range of demand for which the solution is applicable, and thereby assist in the selection of another demand to be investigated. This continues until all possible demands for the unknown node have been dealt with, at which time an overall routing strategy is determined.

DC6	
Step One:	Create initial routes and a surplus table.
Step Two:	Select as the initial demand range [1, the minimum non-zero surplus capacity].
Step Three:	For each demand, solve corresponding DC3 problem, then find the maximum demand for which the strategy applies.
Step Four:	Choose another demand and repeat Step Three until the demand range [1, maxVehCap] is accommodated.
Step Five:	Consider the possibility of zero demand.
Step Six:	Combine solutions into a integrated routing strategy

Procedure 6.15: DC6 solution.

6.9.3 Initial route construction

As in other examples in this study, the initial routes are chosen from the best solutions from a selection of smallest neighbour/ farthest insertion, smallest neighbour/nearest insertion, largest neighbour/farthest insertion, and largest neighbour/nearest insertion, all improved using k-opt, and constructed without the unknown node. The results are shown in Table 6.33 and the corresponding surplus table is given as Table 6.42. Routes may be traversed in the opposite direction and

the multi-routed vehicle may use the alternative route order. If the demand and location of the unknown node allow it to be accommodated in the initial scheme, no other scheme will be implemented.

Vehicle	Distance	Load	Surplus distance	Surplus capacity	Nodes
blue	36	193	14	7	0 7 10 13 0
green	50	149	0	51	0 11 9 8 5 0
black	40	118	10	82	0 12 3 1 0
pink	48	194/119	2	6/81	0 4 2 0 6 0
red	24	200	26	0	0 14 0
	198				

Table 6.42: Surplus table for DC6 example.

6.9.4 Determining demand

The maximum demand for any routing strategy is the maximum demand that can be accommodated by every aspect of that strategy, thus ensuring that the strategy will be feasible over the entire demand range. Portions of the strategy for any demand range that are feasible over an extended range may be included in the strategies for the demand range that encompasses the extension.

For example, the surplus table above indicates that a demand of 7 can be satisfied by every vehicle (except red) as long as the node for which the demand is intended is accessible feasibly by any of the non-red vehicles. If the relevant routing scheme applied to a region no greater than that, the maximum demand for which the scheme applied would be 7, even though up to 82 units could be delivered by a vehicle operating under that routing scheme (see Figure 6.28). The next demand range considered would begin at 8 and may include aspects of the scheme that operates for the range $[1,7]$ other than for the blue vehicle.

6.9.5 Routing strategy

Having determined routing schemes that can be applied to different regions for different ranges of demand, they are combined into an overall routing strategy that incorporates them into a cohesive, compatible union. For example, the scheme indicated in §6.9.4 for the demand range $[1,7]$ can be succeeded by a scheme for the demand range $[8,82]$ that includes schemes A and C2 from §6.6. Similar

schemes are developed for the demand range $[9, \text{maxVehCap}]$, which may be further sub-divided, as indicated in §6.9.4.




Figure 6.28: The vehicle to service node 15, with non-zero demand less than 8, depends on the vertex at which the node is located. Unfilled discs (other than nodes) require a dedicated vehicle.

6.9.6 Conclusions

Coordination assists in problems in which both location and demand are unknown by the early establishment of coordination sites that can be used by en route vehicles as reservoirs to replenish stock, or as sites at which to swap loads, and by enabling vehicles to share responsibility for a previously unknown demand by the use of split deliveries. Development of an overall strategy is computationally intensive and the result is not unique.

6.10 DC7 Nothing is known

In problems of this type, there is at least one node about which there is no information. In one variant of DC7 problems, even the existence of nodes about which nothing is known is unsuspected, and when demand, time and location are all unknown, estimates can be made only if the appearance of a node is anticipated. In real-world DC7 problems, the objective function to be minimised often involves the time that a node waits for service, rather than the total distance travelled, and are represented in the repair industry and emergency services [BL93], but this study will continue to use an objective function that includes only the total distance travelled and the fleet size, in order to maintain consistency with earlier sections.

Approaches to incorporating coordination in the solution to DC7 problems in which the appearance of such nodes is anticipated include:

(a) Conservative estimation

Tentative routes are created that include potential (imaginary) nodes as well as the nodes with advance requests, so that, when new nodes are revealed, the most appropriate of the tentative routes are assigned to them and the other routes are modified. By allocating the unknown demand to too many vehicles, their effective capacities are reduced, and route numbers and fleet sizes may be increased unnecessarily thereby creating false environments for coordination. If new nodes are not revealed until late in the work cycle, the coordination process will be handicapped because material that could have been used as replenishment stock in coordination is tied up as assigned stock for potential nodes.

(b) Optimistic

Initial routing and scheduling is carried out with only those nodes with advance requests and, as each new node is revealed, it is inserted into the most appropriate route along with a possible modification of routes. This approach relies on the continued availability of coordination sites throughout the grid and for the duration of the work cycle, and is more effective in situations of high grid density (to reduce inter-nodal travel times) and high route density (to increase the potential for coordination).

(c) Geographical

The appearance of new nodes is anticipated by the establishment of coordination sites as transient replenishment depots which vehicles may access, rather than by reserving space on the vehicles themselves. Site location takes advantage of known grid features to facilitate ready access by

a variety of vehicles, thus allowing each vehicle to service more nodes than the conservative approach above, and is likely to be more successful in medium density and sparse grids than the optimistic approach because of the greater influence of geographical features.

(d) Dedicated vehicle

Nodes with advance requests are assigned to vehicles and routes without regard for the existence of nodes with real-time requests and, as each new node is revealed, a vehicle is deployed to service it. As the work cycle progresses, the appearance of new nodes introduces into the routing scheme vehicles that may not be able to participate in the coordination process because their unpredicted temporal and spatial characteristics may be incompatible with existing schemes. Because dedicated vehicles are deployed from the depot, route failure will occur where new nodes are located at such a distance from the depot that no vehicle may feasibly service them. In order to overcome this problem, mop-up vehicles may be deployed from coordination sites established earlier in the work cycle.

(e) Periodic DC6

This approach is similar to the approaches used in other discussions in this chapter, where less complex problems are solved at a variety of times and the results integrated into an overall routing strategy. For DC7 problems, the solution of DC6 problems at suitably selected times is computationally intensive and may be intractable for systems other than those with only a few unknown nodes.

As in earlier investigations in this study, a combination of approaches may be appropriate so that, for example, where geographic coordination sites are unsuitable, a dedicated vehicle may be deployed to service a new request, and this may occur in the context of a DC3 problem that is being solved as part of a periodic DC6 approach.

6.10.1 Problem generation

A single depot/ fourteen node system is generated similar to that employed above with time, demand and location given in Appendix F. At some unknown time, a fifteenth node will appear with an as yet unknown demand. The situation is identical to the example in §6.9, with the exception that the time at which the attributes of the unknown node will be revealed is also unknown at the start of the work cycle.

6.10.2 Problem solution

Using the periodic DC6 approach indicated above, the time period is sub-divided into smaller periods and, for each, a DC6 problem is solved, as in Procedure 6.16. Compatibility is achieved by ensuring cohesive schemes at each level of the solution, that is, for each demand, the corresponding DC3 problems are integrated into an harmonious strategy, and at each time, the corresponding DC6 problems are integrated into a similarly compatible structure, so that the overall strategy consists of sub-strategies that fit well together.

DC7	
Step One:	Select a suitable value of $t \in [1, t_{\max}]$.
Step Two:	Solve the corresponding DC6 problem.
Step Three:	Choose another value of t and solve the corresponding DC6 problem.
Step Four:	Repeat Step Three until no vehicles may service new nodes feasibly.
Step Five:	Combine the solutions into an integrated routing strategy.

Procedure 6.16: DC7 solution

6.10.3 Determining time

As in discussions above, the time period may be divided into smaller periods by choosing time values at which to investigate the system, such as $t=10, 20, 30$, etc, but times may also be chosen adaptively so that the next time chosen depends on the progressive state of the current system. Since the former method has already been demonstrated in this chapter, the latter method will be described.

Initial routes are constructed as in §6.9, and a scheme is developed for any initial time value. Having used an initial time value, successive values of t can be determined by noting the time range over which each scheme is applicable in the case that the new node is not revealed in that time period. Suitably strategic times include those where a multi-routed vehicle is replenishing stock at the depot, vehicles are replenishing stock at coordination points, nodes have just been serviced, routes have just been completed, and points at which diversion results in less route modification. In the example, the scheme developed for $t=10$ is adopted and, if the new node is not revealed then, the vehicles may be allowed to proceed until $t=20$, at which time the red vehicle will be at the depot collecting a full payload and will be able to travel to large regions of the grid while still leaving time for its second route to be accommodated by other methods.

As time progresses, vehicles become more committed to their planned routes and the opportunity for coordination reduces so there is a latest time at which existing vehicles may feasibly service new nodes, and, as in the discussions above, the inferred time constraint introduces a latest time at which new nodes may be serviced using dedicated vehicles.

6.10.4 Conclusions

While coordination may be applied in the solution of DC7 problems, in much the same way as the solution of problems in the other categories discussed above, its application is computationally more intensive with an added dimension of uncertainty. Computations increase dramatically with the number of time values chosen, and there is a trade-off between planning time and suitability of the pre-planned routing strategy. By increasing the number of time periods considered, the relationship between the scenario planned for and the scenario that eventuates is strengthened.

Because each route may be involved in servicing a new node, solution is enhanced when the initial scheduling produces vehicles with surplus capacities so that they have potential to detour and to participate in the coordination process. Coordination sites are more useful when they are located either close to established routes (so that they may be used by more than one vehicle with little disruption to service), or at geographical features that take advantage of the structure of the rectilinear grid. By establishing sites early, they may be used over a wide time frame, and their dis-establishment can be achieved more readily.

When the new node is further from the depot, there is more chance of a satisfactory solution where there are more vehicles in the scenario since there is usually less distance for one or more of them to detour to the new node. Similarly, there is greater success where the new node lies close to established routes that have still to be travelled at the time of disclosure since the detour distance is reduced.

6.11 Multiple nodes with dynamic attributes

Thus far, the examples in this chapter have dealt with systems in which attributes of a single node exhibit dynamism, and the solutions involving coordination have been shown to be computationally intensive yet effective. The level of complexity increases enormously as the number of nodes with dynamic attributes increases. Where the same type of dynamism (as in Table 6.1) is present in more than one node, the solution technique recommended for a single node may be applicable with little modification, but, where more than one type is present in nodes, the solution may be much more difficult and may be infeasible. These situations are demonstrated in the following sections.

6.11.1 Same category of dynamism

The solution techniques recommended above for DC1 – DC3 problems include high levels of redundancy that may be taken advantage of in systems with more than one node exhibiting the same type of dynamism; for example, several geographic coordination sites are established in order to service only one DC3 node.

Procedure 6.3 used in the solution of singly dynamic DC1 (unknown time) problems involves the inclusion of the dynamic node in routes that approach it at different times, and the technique can be extended to include several such nodes. Because each node is supplied with only one allotment of its demand (either intact or split), the inclusion of several nodes at several places in any one route does not adversely affect the total carried, but the paths taken by vehicles become convoluted as they wind back on themselves to re-visit more nodes. Although the same technique can be applied to each DC1 node, the interaction between routes becomes more difficult to forecast as the routes themselves are less certain.

Procedure 6.4 used in the solution of singly dynamic DC2 (unknown demand) problems involves the estimation of the unknown attribute and the creation of routing schemes that accommodate ranges of demand for that node. This technique can be applied to more than one node, but, as indicated in §6.5.7, the level of demand of one node may have serious consequences for the surplus capacities of vehicles, and thus for the ability of vehicles to participate in a coordination process designed to service any other node.

Procedure 6.6 used in the singly dynamic DC3 (unknown location) problems involves the establishment of coordination sites located at strategic positions, in order that they may provide replenishment stock for vehicles that subsequently service the unknown node. In the case where more than one node exhibits DC3

qualities, the system created for a single node may be applied directly if those nodes are located in regions that are serviced by different coordination sites, or if the quantity of replenishment stock available and the routing scheme is such that more than one node may be serviced by vehicles using the same coordination site.

Just as the inclusion of coordination becomes more complex as the number of dynamic attributes for a single node increases, obtaining solutions becomes even more difficult as the number of nodes with dynamic attributes increases, and DC4 – DC7 problems involving more than one node with dynamic attributes are very difficult to solve. The potential of any system for improvement by coordination depends on the interactions between vehicles which, in turn, rely on their surplus capacities and the spatial and temporal qualities of the paths they traverse. By creating strategies composed of embedded layers of possible routes, the actual interactions between vehicles are clouded in a fog of uncertainty, and the potential for coordination is difficult to ascertain.

6.11.2 Different categories of dynamism

Systems that involve nodes that exhibit different types of dynamism are usually unable to take advantage of the redundancy built into the singly dynamic systems discussed above, and the level of uncertainty is greater even than in those systems. Consider, for example, the system shown in Figure 6.29, in which the locations and demands for nodes 1 - 14 are those in Appendix F. In addition, node 13 has unknown time attribute (and therefore exhibits DC1 qualities), node 15 has known location but unknown demand or time (DC4) and node 16 is of type DC7. By creating initial routes that involve only those nodes whose locations and demands are known, the base routing scheme shown in Table 6.43 is generated.

Vehicle	Nodes					
blue	0	7	10	13	0	
green	0	11	9	8	5	0
red	0	2	4	0	14	0
pink	0	6	0			
black	0	12	3	1	0	

Table 6.43: Assignments for nodes with known location and demand.




Figure 6.29: Three types of dynamism exist in this structure. In addition to the two nodes indicated, a 16th node has all attributes unknown.

Solution of the DC1 system of node 13 involves including that node in several routes that approach it at different times, but, because the node is isolated with respect to other nodes, there are no existing routes other than that serviced by the blue vehicle that are able feasibly to pass close to it. Consequently, the blue vehicle is assigned to wait at node 13 until either the node can accept delivery, or the vehicle is forced to return to the depot.

By maintaining vigil at node 13, the blue vehicle is then unable to participate in further coordination until either it services the node, or it is en route to the depot. Solution of the DC4 system involving node 15 cannot then continue as indicated in §6.7 because it relies on an embedded DC2 system that uses the blue vehicle to establish a type 1 coordination site. Consequently, the DC4 system cannot be solved if the DC1 system is given priority. On the other hand, if the DC4 system is solved in preference to the DC1 system, node 13 is removed from the route serviced by the blue vehicle, and an extra vehicle is assigned to service that node. The resulting changes in the surplus capacities of the blue vehicle allow it to establish the type 1 coordination site at a much greater distance from the depot, and the previous DC4 solution is re-designed. The assignment of a dedicated vehicle to service node 13 renders a subsequent DC1 solution irrelevant, but without the extra vehicle, the solution is infeasible.

At the same time that the DC1/DC4 conflict is occurring, the problem includes a DC7 system, which requires the solution of a series of DC6 systems. Since each DC6 problem involves the solution of several DC3 problems, the total system now involves the concurrent solution of multiple versions of each category of dynamic coordination except DC5. The resulting confusion of possibilities makes determining the actual locations of vehicles and their surplus capacities at any instant considerably difficult, and infeasibility can happen easily.

6.12 Appraisal of dynamic coordination in the rectilinear grid network

This chapter considers vehicle routing problems in the rectilinear grid network where at least one of the nodal attributes of demand, location and time is unknown at the start of the work cycle, offers a classification scheme, and suggests solutions based on adaptive a priori scheduling and improvement involving coordination, where scenario planning yields route strategies that allow the implementation of suitable routing schemes at the time that more information becomes available. Several such schemes are demonstrated with reference to a fifteen node, six route system set in a medium density grid that includes various geographical features, and where the attributes of one node are unknown at the start of the work cycle. In addition, a sixteen node system is considered in which two nodes have differing forms of unknown attribute and a third node has all attributes unknown. Pre-planning allows a rapid response to changes in the system, and routing strategies attempt to forecast plausible future states based on static systems so that, when all attributes are revealed, a suitable course of action can be chosen from the selection planned for.

Coordination is demonstrated in principle to be an effective tool in dealing with dynamic systems, and several general procedures are suggested when considering specific forms of dynamism. Because of the complexity of dynamic systems and the difficulty in forecasting future states, it is not the aim of this study to produce algorithmic responses, rather to highlight the potential of coordination as a useful technique that may provide assistance in improving dynamic VRPs.

Incorporation of coordination is computationally intensive and requires the identification of intersections and geographical features of the grid in addition to the locations and demands of nodes. It is more effective when coordination sites can be established early in the work cycle, and when advantage can be taken of all known nodal attributes and of the grid structure.

DC1 (unknown time) problems are solved by incorporating nodes with unknown time attribute into routes at several positions so that the routes approach those nodes' locations at different times. DC2 (unknown demand) problems are solved by initially estimating the value of the unknown demand and developing routing schemes that accommodate it; then, schemes are developed for demands outside the initial range, until all demands have been catered for. DC3 (unknown location) problems are solved by establishing coordination sites at strategic locations in order that they may provide replenishment stock when the actual location is revealed. DC4 – DC6 problems contain more than one unknown attribute, and are solved by decomposing the original problem into a series of DC1 – DC3 problems. DC4 (unknown demand and time) problems are solved by repeatedly solving DC2 problems at a variety of different times; DC5 (unknown location and time) problems are solved by repeatedly solving DC3 problems at a variety of times; and DC6 (unknown location and demand) problems are solved by repeatedly solving a corresponding DC3 problem, and adapting the maximum demand range for which the solution is feasible, until all possible demands and locations have been accommodated. The solution of DC7 (nothing known) problems is more complex, and involves the repeated solution of DC6 problems at different times.

As the number of dynamic attributes increases, and as the number of nodes with dynamic attributes increases, the environments that foster the successful application of coordination become more difficult to predict. Consequently, coordination strategies based on the pre-planning of scenarios are more fragile and there develops a greater dependence on an increased fleet size in order to prevent infeasibility.

Although many dynamic vehicle routing problems contain an element of stochasticity, this has not been addressed in this study and remains an area for future research. Similarly, a more thorough investigation of the dynamic aspects of coordination is called for, but is beyond the scope of this study which has shown that coordination has a valuable role to play in the solution of dynamic vehicle routing problems.

CHAPTER 7

CONCLUSIONS AND RECOMMENDATIONS

7.1 Conclusions

Coordination has been shown in this study to be an effective improvement technique in the vehicle routing problem on the Euclidean plane and on rectilinear grid networks for a single depot pure delivery system. Various types of coordination have been identified as load splitting, load swapping, deferral (or obviation) of return to depot, and metastable interim step (leading to solutions that are difficult to achieve otherwise), which occur in different, but not disjoint, environments. In any particular problem, the environment dictates the most suitable form of coordination to use. Coordination in the Euclidean plane has been investigated at nodes, and in the rectilinear grid network at nodes and grid vertices.

Load swapping at a node is likely to be most beneficial when the coordination point lies close to a route other than the route servicing that node, the fleet is able physically to accommodate the total load and, because swapping requires that all vehicles involved be present at the same time, there is an acceptable waiting time. Where load swapping occurs at a point not serviced by any vehicle involved in the exchange, that point must lie close to the routes serviced by each vehicle involved. Fleet capacity and waiting time must also be suitable.

Distance reduction resulting from the deferral of a return to the depot can occur only when the routes serviced by a multi-routed vehicle are close enough to interact. Areas of high route density include the region around the depot (since all vehicles use the single depot) and possible choke-points created by the geographic structure of the grid (where appropriate), and it follows that suitable coordination points for the deferral of a return to the depot are likely to be near these areas. The use of surplus tables provides a means for the rapid identification of possible sites in a grid. In the Euclidean plane, potential sites have been considered at nodes; for coordination to occur at locations other than nodes, regions of the plane would be

identified, with infinite possibilities. It is possible that continuous optimisation problems could be formulated to identify 'best' site locations, but this was not pursued here; it could form a basis for worthwhile future research. Whereas sites created to facilitate load swapping require vehicles to arrive simultaneously, or at least not too long apart, those created to obviate a return to depot for a multi-routed vehicle require only that vehicles arrive at the site consecutively in the correct sequence. Simultaneous arrival, although not required, is found to be acceptable.

While splitting has been investigated as an accepted form of route improvement for many years, other researchers have restricted their investigations to carrying split loads between depots and nodes, without allowing loads to be split or recombined for part of the journey. This study has extended the work of earlier researchers by allowing vehicles to transport part-demands of any size for unrestricted sections of routes. In so doing, it has confirmed earlier results in regard to the advantage to be gained by employing the technique, the influence of demand, and the effect of geography, and has shown that splitting in the rectilinear grid is more effective as grid density increases and with compact structures. Splitting has been identified as an invaluable weapon in the arsenal of coordination techniques.

Since the spatial and temporal characteristics of the systems determine the type of coordination that can be achieved, any alterations to either route order or orientation may substantially affect the outcome. Investigating each combination of route order and orientation is therefore necessary but results in a huge increase in computation time. With ever-faster computers available, this is not seen as a major handicap.

For a given problem, its environment dictates the form(s) of coordination that can be considered. When several suitable environments are present in the same problem, it is not clear which effect should be considered first although empirical evidence suggests that, for a grid of medium density, load splitting be applied before the removal of any return to the depot, which should be investigated before any method to reduce path overlap. The removal of route intersections reduces total distance very effectively in systems that do not involve coordination, whereas load swapping actively creates intersections. Since there is no theoretical advantage in using one form of coordination over another, each problem needs to be considered on its own merits. For each problem, different combinations of route order and orientation can produce conditions which, because of the temporal and spatial inter-dependence of coordination, can favour one method over another or prevent the potential for the existence of coordination.

Because split deliveries do not require the establishment of coordination sites at locations other than destination nodes, this form of coordination can be used during, or as part of, the routing phase. The other forms of coordination require the existence of routes, so they are applied at the scheduling phase. For many problems, more than one type of coordination are shown to be appropriate and circumstances may allow the use of multiple forms of coordination. Because the application of some forms of coordination may involve the imposition of waiting times, sequential coordination is more difficult to apply where distances and travel times are already large by reason of grid density and geography.

Three types of coordination site have been identified in the rectilinear grid, depending on their use: type 1 sites are established on single-noded routes by vehicles before they begin their other routes, and consequently often hold stock in the quantity of multiples of maxVehCap ; type 2 sites are used by vehicles detouring to visit them; type 3 sites are used by vehicles without detouring because the sites lie along their intended paths. Any coordination site can function as more than one type. Suitable locations for the establishment of type 1 coordination sites can be identified by the use of surplus tables that indicate the maximum distance that any vehicle may travel from the depot for such a purpose.

In addition to identifying a variety of types of coordination and investigating their use, a procedure has been presented that reduces the edges in a rectilinear grid network to a representative member of its phase space whilst maintaining connectedness of edges within a hull created by the depots, nodes and their appropriate edges. Such a grid reduction improves the efficiency of all calculations involving paths within the grid. In a deterministic system, where the locations of all essential points are known, removing all edges outside the coordination hull could reduce the grid further.

The results of grid reduction are shown to depend on the layout of depots and nodes, with changing node locations resulting in a diversity of reduced grids varying from compact to broadly distributed. Similarly, the density of the reduced grid also depends to some extent on the locations of essential points. Subtle changes in the existence of edges of the initial (unreduced) grid can have major effects on the final grid. Consequently, it is very difficult to predict the shape and density of the grid network before the reduction process is applied.

Having reduced a grid, the existence of features such as lakes and peninsulas indicates regions of high route density, and the presence of choke-points is an indication of regions more likely to contain suitable coordination sites.

Because coordination is a consequence of route interaction, the number of intersections in a grid is an indication of the potential for improvement by

coordination, with empirical evidence suggesting that the chance of lowering total distance travelled by coordination other than splitting is less with increasing grid density. Coordination by splitting delivery requires one or more vehicles to access a customer node, and is more likely with increasing grid density and with more compact grids. As grid density increases, splitting is shown to be more effective than other forms of coordination.

The application of coordination in fully deterministic systems has been demonstrated here to reduce both fleet size and total distance travelled while still permitting the minor objectives of balanced routes and a streamlined solution.

The use of coordination in dynamic vehicle routing in the rectilinear grid network requires contingency planning. This study has investigated, in general terms, the principles of improvement by coordination of dynamic vehicle routing problems for representative systems involving fifteen nodes where a single node exhibits dynamism. A classification scheme has been introduced for the investigation of every combination of dynamic attribute of node location, node demand, and time (at which the information becomes available or at which delivery can be made), and routing strategies developed that involve the possible use of coordination. Brief consideration has been given also to similar systems in which more than one node exhibits the same type of dynamism, and to a sixteen-node system in which different types of dynamism are exhibited by three nodes. Solutions are effective but computationally intensive, and involve the pre-planning of responses to plausible future states based on static systems so that, as attributes are revealed, suitable courses of action can be selected from the possibilities planned for.

In applying coordination to dynamic systems, coordination sites are established early in the work cycle by estimating unknown attributes and creating alternative schemes for the occasions when those estimates are incorrect. Where the value of only one attribute is unknown, its estimate is based on historical data; where two attributes have unknown values, several iterations of the one-unknown system are made using a selection of values for the other unknown. Because actual values of attributes are not known at the start of the work cycle, their possible effects are accommodated by establishing coordination sites (within routing schemes) to be used in the event that the attributes are revealed to lie within specific ranges. When those values become evident, appropriate routing schemes are enacted and unused coordination sites are eventually emptied. Consequently, the resulting routing strategies incorporate a high level of redundancy.

As in the deterministic VRP, the potential for coordination in the dynamic VRP is dictated by the environment. As the number of nodes with dynamic

attributes increases, and as the number of dynamic attributes themselves increase, it becomes more difficult to predict future environments, and the potential for coordination becomes less clear.

The procedures developed here have considerable practical value in the commercial world. Where a customer base is well established with known inter-nodal distances and demand histories, deterministic and dynamic coordination can be applied with little initial outlay and small maintenance costs with the prospect of substantial reductions in both fixed and variable operating costs.

The five practical problems of Chapter 1 can be addressed in the light of their potential for coordination:

Problem 1: The roading company supplying its road crews with bitumen from a distant source is able to reduce the total distance travelled by establishing a coordination site at a suitable location at which the bitumen can be transferred from the main carrier to other suitably equipped vehicles. By obviating the return to depot (see §4.13) for most vehicles, only the main carrier is required to travel to the supply depot for replenishment of stock, while the other vehicles are required only to travel to the coordination site for replenishment of their stock. In addition to a large reduction in total distance, the area of operations of the company is increased, allowing it to compete more effectively for further work. There may also be savings in the type of vehicle that is used, since most of the fleet can be made up of vehicles that do not have to carry large volumes of hot bitumen and do not have to keep their payloads hot for as long, resulting in both smaller bodies and lighter heating coils. Selection of suitable coordination sites is made prior to tendering for work since their location will have a direct influence on the operating costs and time involved, and the sites are incorporated into the routing scheme as nodes with zero demand to be serviced by each vehicle. As work progresses and reaches a conclusion at each job site, the demand of the corresponding nodes diminishes and eventually is eliminated. The changing state of node demand is not entirely predictable because of the influence of weather on the availability and accessibility of the work sites but the weak dynamic nature of the problem can be accommodated by treating the problem as deterministic. Consequently, when adverse weather conditions prevent a job being tackled on any particular day, the vehicle that would have serviced that task is temporarily idle. Since all vehicles in the fleet are custom built and owned by the roading company, no hiring fee is involved and therefore this does not represent a disadvantage.

Problem 2: The cross-town delivery company is able to reduce its fleet size without greatly increasing its travel distance and service time by incorporating a combination of non-split coordination techniques (see §5.13 and §5.14) into an essentially rectilinear routing scheme. By establishing coordination sites at suitable customer nodes and at appropriate locations away from customers, demands are able to be swapped and transferred between vehicles resulting in more efficient use of a smaller fleet. At coordination sites, vehicles with surplus distance capacity that have delivered much of their payloads are able to accept loads from other vehicles that have limited distance capacity, and can thence deliver these loads to customers that otherwise would have been serviced (if at all) by an extra vehicle. Thus, the fleet size is reduced without adverse effect on customers.

Problem 3: Although pre-sales carriage of animals forms part of the class of problems known as 'pickup' and only the post-sales carriage is a delivery problem, the techniques developed in this study (see §5.9.4) can be applied to the pre-sales problem enabling the stock transport company to reduce operating costs by better utilisation of its trailers without increasing the road taxes associated with the conveyance of multi-axle trailers on the open road. This is achieved by establishing type 1 coordination sites at which trailers are left for later uplift, while the main truck travels to collect stock from farmers. By not including a trailer in service visits that do not involve it, the trailer does not incur the accompanying road user charges, and the truck is able to negotiate roads and accessways that would be difficult or impossible with an attached trailer. In addition, travel without the trailer is faster and the speedier response may go some way to offsetting the time required to detach and reattach the trailer. As a type 1 site, the trailer is available for use by any vehicle from the same company. While the trailer is detached, stock can be transferred between the main truck and the trailer by reversing the vehicle up to the rear of the trailer and allowing animals to move from one to the other. Thus, trailers at the coordination site can be used as temporary storage units. Judicious use of baffles separates animals from different farms. Although similar techniques can be used in the post-sales delivery

Problem 4: The rural bus company extends its service by creating a coordination site at one of the schools at which many-way transfer of pupils is allowed to occur. In the morning, the routing scheme is in two parts: initially, children are delivered to a coordination site; then they are delivered from this site to their

schools. In the afternoon, the process is reversed with children being delivered to the coordination site and then taken from there to their homes. Before school, buses travel to distant areas to collect schoolchildren whom they deliver to the local secondary school (the coordination site). This location is chosen because it is able physically to accommodate all the vehicles involved, has undercover areas in which the children can wait, and the school supplies a staff member to supervise the passengers before and after school. When all the buses have delivered their charges to the coordination site, they collect children from the assembled mass and deliver them to the appropriate school. In the afternoon, buses collect pupils from the schools, deliver them to the secondary school, collect children and take them home. The home-coordination site section of the routing scheme takes much longer than the school-coordination site section. An advantage of this method is that, because the secondary school is attended by far more children than the rural primary schools, establishing the site at that school removes a requirement for buses to service the secondary school in the coordination site-school section. In addition, the scheme can be extended easily to include other schools in the area already covered because the commitment has already been made to collect children from outlying areas and adding extra stops in the home-coordination site section does not have much impact on the total time taken. An unexpected advantage is that young children from tiny rural areas meet other children whom they would not expect to encounter otherwise until their teenage years. By utilising a coordination site at a customer node, the bus company has been able to increase its service without changing its fleet.

Problem 5: Coordination is of no direct benefit to the fashion clothing manufacturer because its payloads are very small and are already delivered directly to selected retail outlets as high priority packages. Staff are employed in highly specialised teams, and workflow within the manufacturing concern is linear with no overlapping paths or intersections. The environment does not allow for improvement by coordination.

7.2 Recommendations for future research

Further research into coordination as an improvement procedure for the static and dynamic VRP should include the incorporation of factors such as pickup, asymmetrical distances, multiple depots, deadlines, time windows, commodity classes and multiple vehicle types. Each of these modifications can be made by adjusting the internodal distance tables and enhancing the procedures developed in this study and are not expected to alter the concept of coordination as presented here.

While an indication has been presented in this study of the suitability of coordination as an improvement procedure in dynamic vehicle routing, much remains to be done in investigating the process in detail. This should include extending the planning horizon to cover more than one work cycle, thereby reducing the possibility of route failure and enabling the introduction of priority demands and subset selection. The effect on the potential for coordination of increasing both the number of nodes with dynamic attributes and the number of dynamic attributes is a suitable area for research; because the incorporation of coordination involves contingency planning, and is heavily influenced by environment, the ability to predict future environments in a system is critical to the effectiveness of coordination as an improvement method for dynamic VRPs. It appears unlikely that any automated umbrella scheme for dynamic coordination would be effective.

The involvement of coordination in stochastic systems is expected to be a fertile area for research since it has been effective in deterministic and dynamic systems.

Because of the temporal and spatial inter-dependence of coordinated systems, the introduction of delayed departure times can be expected to have a significant impact on the potential for coordination.

Psaraftis (see [Ps88]) identifies one of the major differences between static and dynamic routing as the possibility of congestion, and Carey and Srinivasan (see [C93]) have shown that because the peaks in congested traffic flows most affect costs, the rate of change in the level of congestion strongly influences the appropriateness of any model. The possible use of coordination sites to reduce congestion in networks by assisting in regulating flow should be investigated as the suitability of applying coordination in a congested dynamic system might be found to have benefits in real-world scenarios.

While much of this study has involved the establishment of coordination sites at locations other than nodes, the actual identification of those off-node sites

could have been made more effectively by applying optimization to the small regions containing potential sites. Heuristic procedures, similar to those presented here, could identify a selection of suitable coordination sites from which optimization procedures would detect the best site(s). The process of determining the suitability of coordination would then involve pockets of optimization within a heuristic framework.

The concept of coordination has been examined in this study by permitting the establishment of coordination sites only at nodes (in the Euclidean plane) and at nodes and vertices (in grid networks), and is expected to be valid when the entire frame of reference is available for the establishment of coordination sites, although this has yet to be investigated. This study has already shown, in the rectilinear grid network, that coordination sites may be established productively along strings of edges at vertices other than intersections, particularly where each vehicle involved detours, or where type 1 sites are employed. Future investigations may reveal that this benefit can be extended when non-vertex locations are available.

Although this study has determined the environment as the critical factor in determining whether, and what form of, coordination should be applied to a system, other investigations may reveal different factors to be more important. This research has shown that initial solutions are important only to the extent to which they create conditions conducive (or otherwise) to improvement by coordination. Empirical evidence gathered in this study from the application of coordination to hundreds of initial solutions generated using different methods in both the Euclidean plane and rectilinear grid networks has failed to indicate any general pattern dependent on the actual method of creating the initial solution.

In introducing the concept of coordination as an effective improvement tool for deterministic and dynamic VRPs, and identifying and investigating several types of coordination, it would have been desirable to have been able to predict the most effective form of coordination to apply to any given vehicle routing problem. Unfortunately, the heavy dependence on the spatial and temporal environment and on the interaction of vehicles has rendered such prediction optimistic beyond the broad guidelines already discussed. The classification of problem types is an area for future research and could include a range of representative problems to show the effectiveness of different types of coordination, as well as the difficulty of applying coordination. Because of the sensitivity of the results of the grid reduction technique to subtle changes in initial grids, such problems should involve grids that already have been reduced. Each form of coordination applied to a system alters the environment, possibly changing the suitability of other forms of coordination (or

repeated applications of the same form) to improve the solution. Selecting the most effective sequence of multiple applications of coordination for problems may require more in-depth investigation involving pre-processing to determine suitable schemes. Such look-ahead methods may best be included in the incorporation of the various forms of coordination as a variety of move structures in meta-heuristics.

While this study has identified five forms of coordination (splitting, deferral of return to depot, swapping, transfer, metastable state) and three types of coordination site, there may be other forms that are effective in improving solutions to the VRP. The metastable state evident in problems in the Euclidean plane is not present in the rectilinear grid, because of its rectilinear structure, and investigation of other frames of reference may produce other differences. The spherical frame of reference, already discussed as important in the real-world application of mid-air refuelling, may yield other types of metastable state not found in the Euclidean plane, or entirely new forms of coordination.

The fifteen vehicle limit imposed by the MATLAB programming language has been suitable, with appropriate values of maximum vehicle capacity and maximum demand, for investigating the concept of coordination thus far, but much larger systems with high numbers of interacting routes may provide other valuable insights into the application of coordination as an improvement tool. A suitable development medium would be any non-specific, general purpose, lower level, object-oriented programming language, such as C++, which has the advantages of greater speed and flexibility.

APPENDIX A

REDUCTION OF RECTILINEAR GRIDS

This Appendix contains stepwise examples for reduction of rectilinear grids of sparse, medium, high and full density.

NCS = "No Choice String", as described in the text.

Reduction of Sparse Density Grid




Figure A.1: Sparse density grid, 438 edges, 10 nodes, 1 depot.




Figure A.2: Cul de sacs removed.




Figure A.3: NCS removed without considering NCS involving essential points. Cul de sacs created.




Figure A.4: Cul de sacs removed.




Figure A.5: NCS removed without considering NCS involving essential points. Since no cul de sacs have been created, edges involving essential points can now be considered.




Figure A.6: Removed NCS with essential points.




Figure A.7: Removed NCS, thus creating other NCS.




Figure A.8: Final grid has 96 edges.

Reduction of medium density grid




Figure A.9: Medium density grid, 1 depot, 30 nodes, 643 edges.




Figure A.10: Cul de sacs removed.




Figure A.11: First reduction by removal of NCS. Other NCS have been created.




Figure A.12: Second reduction by removal of NCS. Other NCS have been created.




Figure A.13: Removal of remaining NCS not involving essential points.




Figure A.14: Removal of NCS involving nodes.




Figure A.15: Removal of NCS created in previous step.




Figure A.16: Removal of NCS created in previous step. Final grid.
Total number of edges: 341.

Reduction of high density grid




Figure A.17: High density grid. 790 edges. 1 depot, 30 nodes.




Figure A.18: Removal of NCS not involving essential points.




Figure A.19: Removal of NCS.




Figure A.20: Removal of NCS.




Figure A.21: Removal of NCS.




Figure A.22: Removal of NCS, including loop.




Figure A.23: Final reduced high density grid. 451 edges.

Reduction of full density grid




Figure A.24: Initial full density grid. 840 edges.




Figure A.25: Removal of NCS not involving essential points.




Figure A.26: Removal of NCS involving essential points.




Figure A.27: Removal of NCS created in previous step.




Figure A.28: Removal of NCS created in previous step.




Figure A.29: Final grid. 475 edges. Each outside corner houses a node.

APPENDIX B

CHANGING DEPOT LOCATION

This Appendix contains a worked example for the effect of depot location on coordination. For a given data set, the situation before and after coordination is considered with the depot in each of two locations, [12,12] and [10,8].

Node	Demand	Location	
1	43	19	4
2	166	17	15
3	12	16	8
4	28	18	14
5	114	3	8
6	119	18	18
7	118	8	17
8	3	2	4
9	32	5	3
10	17	8	18
11	18	9	8
12	72	16	10
13	58	4	13
14	200	13	7
15	150	14	8

Table B.1: Demands and node locations

Depot at [10,8]

When the depot is located at [10,8], the route structure is that shown in Table B.2.

Without coordination

Route#	Distance	Load	Route
1	38	167	0 11 9 8 5 0
2	40	193	0 7 10 13 0
3	36	193	0 15 1 0
4	32	194	0 2 4 0
5	18	84	0 12 3 0
6	16	200	0 14 0
7	<u>42</u>	<u>119</u>	0 6 0
	222	1150	
Vehicle#	Distance	Route(s)	
1	38	1	
2	40	2	
3	36	3	
4	50	4 5	
5	42	7	
6	16	6	

Table B.2: Route structure without coordination, depot at [10,8].

Applying coordination

Construct the surplus table:

vehicle	dist	load	Surplus	
			dist	capacity
1	38	167 0 11 9 8 5 0	12	33
2	40	193 0 7 10 13 0	10	7
3	37	193 0 15 1 0	13	7
4	32	194 0 2 4 0	0	6/116
	18	84 0 12 3 0		
5	16	200 0 14 0	34	0
6	42	119 0 6 0	8	81

Table B.3: Surplus table for depot at [10,8]

If vehicle 4 is involved in coordination it cannot add distance to its route, although it can carry more payload.

Vehicle 1 creates a type 1 site at $X=[14,8]$, coincident with node 15, leaving 193 demand units there. Vehicle 1 has added 12 distance units. Of the 193 units, 150 units is given to node 15. This is a bonus since the site is coincident with a node.

Vehicle 4 services the nodes on route 5 without returning to the depot. It then travels to the site created by Vehicle 1, collects enough demand for node 1, delivers this to node 1 and returns to the depot.

Vehicle 2 creates a type 1 site at $Y=[14,9]$, leaving 194 demand units there. The vehicle has increased its distance by 10 units to the maximum distance allowed.

Vehicle 6 services node 14 then travels to the site created by vehicle 2 without returning to the depot. It collects the 194 units and delivers these to nodes 4 and 2 before returning to the depot. Vehicle 6 has travelled a total of 36 distance units.

Vehicle 3 has no route 3 to service since the nodes comprising this route, 15 and 1, have been serviced by two other vehicles.

The route structure is shown in Table B.4.

With coordination

Vehicle#	Distance	Load	Route
1	50	193/167	0 X 15 0 11 9 8 5 0
2	50	194/193	0 Y 0 7 10 13 0
4	44	84/43	0 12 3 X 1 0
5	18	119	0 6 0
6	36	200/194	0 14 Y 0 4 2 0
	198		

Vehicle#	Distance	Route(s)
1	50	0 X 15 0 11 9 8 5 0
2	50	0 Y 0 7 10 13 0
4	44	0 12 3 X 1 0
5	18	0 6 0
6	36	0 14 Y 0 4 2 0

Table B.4: Route structure with coordination, depot at [10,8].

Both the total distance and fleet size have been reduced. Distance has been reduced by 12% in the affected routes.

The possible locations for type 1 sites are shown in the diagram below. The discs represent the outer reaches for sites for each vehicle. Virtually the entire map is accessible by vehicle 5 without breaching distance constraints since that vehicle has a surplus range of 17 distance units.




Figure B.1: Locations for type 1 sites. Pink discs: outer limit for vehicles 1 and 3; green discs: outer limit for vehicle 2; yellow discs: outer limit for vehicle 6.

The diagram below shows the final routes.




Figure B.2: Final routes involving coordination, depot at [10,8].

Depot location at [12,12]

When the depot is located at [12,12], the route structure is that shown in Table B.5.

Without coordination

Route#	Distance	Load	Route
1	36	193	0 7 10 13 0
2	50	149	0 9 8 5 0
3	20	194	0 2 4 0
4	28	180	0 11 15 3 0
5	40	106	0 12 1 0
6	28	119	0 6 0
7	<u>24</u>	<u>200</u>	0 14 0
	226	1141	

Vehicle	Distance	Route(s)	Nodes
1	36	1	0 7 10 13 0
2	50	2	0 9 8 5 0
3	48	3 4	0 2 4 0 11 15 3 0
4	40	5	0 12 1 0
5	28	6	0 6 0
6	24	7	0 14 0

Table B.5: Route structure before coordination, depot at [12,12].

The surplus table is shown in Table B.6.

vehicle	dist	load		spare	
				dist	capacity
1	36	193	0 7 10 13 0	14	7
2	50	149	0 9 8 5 0	0	51
3	48	194/180	0 2 4 0 11 15 3 0	2	6/20
4	40	106	0 12 1 0	10	94
5	28	119	0 6 0	22	81
6	24	200	0 14 0	26	0

Table B.6: Surplus table, depot at [12,12].

Normally, one would attempt to obviate the return to the depot for the multi-routed vehicle, but in this case the depot lies between the last node of the earlier route and the first node of the later route, so obviation is not likely to be productive.

Two routes (6 and 7) are close to being capable of sharing the same vehicle without violating the distance constraint. A good option is to reduce slightly the distance of at least one of them so that they may share the same vehicle. This can be done by the establishment of a type 1 site by a third vehicle.

The spare distance of Vehicle 1 allows it to establish such a site within 7 units of the depot. If routes 6 and 7 are to be combined, a natural site would be at a point on the shortest path between nodes 6 and 14, closest to the depot. This reduces the extra distance travelled by Vehicle 1.

Vehicle 1 establishes a type 1 site, X, at [15,12] with load equal to Demand(14), ie 200 units, adding 10 distance units to its total distance. Vehicle 5 services node 6, travels to X, collects the load, delivers it to node 14 and returns to the depot.

Vehicle	Distance	Route(s)	Nodes
1	46	1	0 X 0 7 10 13 0
2	50	2	0 9 8 5 0
3	48	3 4	0 2 4 0 11 15 3 0
4	40	5	0 12 1 0
5	<u>44</u>	6	0 6 X 14 0
	228		

Table B.7: Route structure after coordination, depot at [12,12].

The fleet size is reduced by one and the total distance increases by 2%.




Figure B.3: Routes after coordination, depot at [12,12].

Comparison

Depot location	Without coordination		With coordination		
	Total dist	Fleet size	Total dist	Fleet size	No. coordn sites
[12,12]	226	6	228	5	1
[10,8]	222	6	198	5	2

Table B.8: Effect of depot location.

APPENDIX C

LOAD PLANS AND ROUTE DATA FOR EXAMPLES IN CHAPTER 4

Figure 4.9

MaxVehDist = 220, maxVehCap=200, 10 nodes – one with overlarge demand.

Without coordination

Initial feasible solution created using modified savings.

Route#	Distance	Load	Route
1	194.684	190	0 6 5 3 10 0
2	144.470	153	0 8 7 4 0
3	57.301	118	0 1 0
4	102.315	220	0 2 0
5	<u>45.443</u>	94	0 9 0
	544.214		

Table C.1: Without coordination.




Figure C.1: Without coordination.

Split deliveries at node

Three coordination sites: nodes 1, 2 and 6.

Route#	Distance	Load	Route
1	123.656	200	0 5 3 2 0
2	152.579	200	0 8 7 4 1 0
3	184.379	200	0 10 2 6 0
4	<u>67.934</u>	175	0 6 9 1 0
	528.549		

Table C.2: Split deliveries to nodes 1, 2 and 6.




Figure C.2: With coordination.

Total distance is reduced 3%.

Figure 4.10

maxVehCap=200, maxVehDist=220, number of nodes=30.

Without coordination

Initial feasible solution created using modified savings algorithm.

Route#	Distance	Load	Route
1	166.126	185	0 19 25 27 9 22 0
2	148.588	171	0 26 21 30 0
3	94.247	132	0 10 28 14 29 0
4	104.052	193	0 8 3 23 0
5	125.610	176	0 18 7 4 0
6	144.777	157	0 16 15 0
7	73.574	194	0 24 11 0
8	54.907	170	0 5 2 0
9	37.095	173	0 13 12 0
10	27.007	133	0 1 0
11	95.116	129	0 6 0
12	107.803	147	0 17 0
13	<u>53.467</u>	130	0 20 0
	1232.369		

Table C.3: Initial routes.




Figure C.3: Initial routes.

Split deliveries at node

Route#	Distance	Load	Route
1	153.388	200	0 19 27 9 25 15 6 0
2	149.481	200	0 26 21 30 16 0
3	107.703	200	0 19 17 0
4	141.217	151	0 20 23 22 6 0
5	88.870	200	0 8 3 24 0
6	82.642	200	0 5 2 24 11 0
7	59.334	139	0 20 2 0
8	46.003	200	0 12 13 5 0
9	46.477	200	0 18 1 0
10	147.430	200	0 18 16 7 4 0
11	<u>100.606</u>	139	0 12 10 28 14 29 0
	1123.152		

Table C.4: Split deliveries.




Figure C.4: Split deliveries at node.

Total distance has reduced from 1232 to 1123 units, a reduction of 9%.
Number of routes has reduced from 13 to 11.
Distances are balanced more evenly.
The range of loads is approximately the same.

Figure 4.11

Test problem of [CMT79]. maxVehCap=140, maxVehDist= ∞ , number of nodes=75

Without coordination

Initial feasible solution created using metaheuristics of [H95 p208].

Route#	Distance	Load	Route
1	83.08	139	0 53 66 65 10 58 0
2	52.93	132	0 17 40 9 39 12 26 0
3	120.16	138	0 72 31 25 55 18 50 32 44 3 0
4	92.69	140	0 6 33 63 23 56 24 49 16 0
5	100.01	138	0 62 22 64 42 41 43 1 73 51 0
6	74.38	140	0 30 74 21 61 28 2 68 0
7	107.07	139	0 48 47 36 69 71 70 20 37 27 0
8	89.53	139	0 52 19 54 13 57 15 5 29 45 0
9	29.08	126	0 75 4 34 46 67 0
10	<u>86.33</u>	133	0 7 11 59 14 35 8 0
	835.26		

Table C.5: Without coordination.




Figure C.5: Initial routes.

Route#	Distance	Load	Route
1	83.08	139	0 53 66 65 10 58 0
2	52.93	132	0 17 40 9 39 12 26 0
3	120.16	138	0 72 31 25 55 18 50 32 44 3 0
4	92.69	140	0 6 33 63 23 56 24 49 16 0
5	100.01	138	0 62 22 64 42 41 43 1 73 51 0
6	74.38	140	0 30 74 21 61 28 2 68 0
7	107.07	139	0 48 47 36 69 71 70 20 37 27 0
8	89.53	139	0 52 19 54 13 57 15 5 29 45 0
9	29.08	126	0 75 4 34 46 67 0
10	<u>86.33</u>	133	0 7 11 59 14 35 8 0
835.26			

Table C.6: Coordination.




Figure C.6: Routes including coordination.

Figure 4.13

Test problem of [CMT79]. maxVehCap=200, maxVehDist=220, number of nodes=50. Improvement by interim step formation.
(PercImp = percentage improvement)

Initial feasible solution created using parallel savings.

Without coordination

Route#	Distance	Load	Route
1	147.746	154	0 24 43 7 26 31 28 3 36 35 20 22 1 0
2	115.170	158	0 14 25 13 41 40 19 42 44 37 46 0
3	138.177	155	0 16 9 50 29 21 34 30 39 33 45 15 17 0
4	88.320	149	0 32 11 2 38 49 10 5 12 0
5	63.057	86	0 6 23 48 8 27 0
6	<u>39.625</u>	75	0 18 4 47 0
	592.094		

Vehicle#	Distance	Route(s)
1	187	1 6
2	203	2 4
3	201	3 5

Table C.7: Initial routes.




Figure C.7: Initial routes.

Near radius = 2

Route#	Distance	Load	Route
1	147.746	154	0 24 43 7 26 31 28 3 36 35 20 22 1 0
2	115.170	158	0 14 25 13 41 40 19 42 44 37 46 0
3	138.177	155	0 16 9 50 29 21 34 30 39 33 45 15 17 0
4	88.320	149	0 32 11 2 38 49 10 5 12 0
5	63.057	86	0 6 23 48 8 27 0
6	<u>39.625</u>	75	0 18 4 47 0
	592.094		

Table C.8: Near radius=2.

Either no nodes near intersections or no intersections.

elapsed_time = 21.8210

Candidate = []




Figure C.8: Near radius=2.

Near radius = 5

Route#	Distance	Load	Route
1	147.746	154	0 24 43 7 26 31 28 3 36 35 20 22 1 0
2	115.170	158	0 14 25 13 41 40 19 42 44 37 46 0
3	138.177	155	0 16 9 50 29 21 34 30 39 33 45 15 17 0
4	63.057	86	0 6 23 48 8 27 0
5	39.625	75	0 18 4 47 0
6	<u>85.466</u>	149	0 32 2 11 38 49 10 5 12 0
	589.241		
Vehicle#	Distance	Route(s)	
1	187	1 5	
2	201	2 6	
3	201	3 4	

Table C.9: Near radius=5.

PercImp = 0.4819
elapsed_time = 204.8040
Candidate = 11




Figure C.9: Near radius=5.

Near radius = 10

Route#	Distance	Load	Route
1	147.746	154	0 24 43 7 26 31 28 3 36 35 20 22 1 0
2	115.170	158	0 14 25 13 41 40 19 42 44 37 46 0
3	63.057	86	0 6 23 48 8 27 0
4	39.625	75	0 18 4 47 0
5	91.320	155	0 32 11 16 50 9 49 10 38 5 12 0
6	<u>121.396</u>	149	0 2 29 21 34 30 39 33 45 15 17 0
	578.314		

Vehicle#	Distance	Route(s)
1	148	1
2	218	2 3 4
3	213	5 6

Table C.10: Near radius=10.

PercImp = 2.3273
elapsed_time = 663.3640
Candidate = 2 9 11 16 38




Figure C.10: Near radius=10.

Near radius = 15

Route#	Distance	Load	Route
1	147.746	154	0 24 43 7 26 31 28 3 36 35 20 22 1 0
2	115.170	158	0 14 25 13 41 40 19 42 44 37 46 0
3	63.057	86	0 6 23 48 8 27 0
4	39.625	75	0 18 4 47 0
5	121.396	149	0 2 29 21 34 30 39 33 45 15 17 0
6	<u>91.320</u>	155	0 32 11 16 50 9 49 10 38 5 12 0
	578.314		
Vehicle#	Distance	Route(s)	
1	148	1	
2	218	2 3 4	
3	213	5 6	

Table C.11: Near radius=15.

PercImp = 2.3273

elapsed_time = 1.4502e+003

Candidate = 1 2 5 9 11 16 29 32 38 46 49 50




Figure C.11: Near radius=15.

Near radius = 20

Route#	Distance	Load	Route
1	147.746	154	0 24 43 7 26 31 28 3 36 35 20 22 1 0
2	115.170	158	0 14 25 13 41 40 19 42 44 37 46 0
3	63.057	86	0 6 23 48 8 27 0
4	39.625	75	0 18 4 47 0
5	96.215	152	0 38 9 50 34 21 29 16 2 11 32 0
6	<u>108.215</u> 570.027	152	0 5 49 10 30 39 33 45 15 17 12 0

Vehicle#	Distance	Route(s)
1	148	1
2	218	2 3 4
3	204	5 6

Table C.12: Near radius=20.

PercImp = 3.7270
elapsed_time = 2.3453e+003
Candidate = 1 2 5 9 11 12 16 21 22 27 29 30 32 34 38 46
49 50




Figure C.12: Near radius=20.

Near radius = 25

Route#	Distance	Load	Route
1	147.746	154	0 24 43 7 26 31 28 3 36 35 20 22 1 0
2	115.170	158	0 14 25 13 41 40 19 42 44 37 46 0
3	63.057	86	0 6 23 48 8 27 0
4	39.625	75	0 18 4 47 0
5	96.215	152	0 38 9 50 34 21 29 16 2 11 32 0
6	<u>108.215</u>	152	0 5 49 10 30 39 33 45 15 17 12 0
	570.027		
Vehicle#	Distance	Route(s)	
1	148	1	
2	218	2 3 4	
3	204	5 6	

Table C.13: Near radius=25.

PercImp = 3.7270

elapsed_time = 2.4829e+003

Candidate = 1 2 3 5 9 10 11 12 16 20 21 22 27 29
30 32 34 37 38 46 47 49 50




Figure C.13: Near radius=25.

Near radius = 30

Route#	Distance	Load	Route
1	147.746	154	0 24 43 7 26 31 28 3 36 35 20 22 1 0
2	115.170	158	0 14 25 13 41 40 19 42 44 37 46 0
3	63.057	86	0 6 23 48 8 27 0
4	39.625	75	0 18 4 47 0
5	96.215	152	0 38 9 50 34 21 29 16 2 11 32 0
6	<u>108.215</u>	152	0 5 49 10 30 39 33 45 15 17 12 0
	570.027		
Vehicle#	Distance	Route(s)	
1	148	1	
2	218	2 3 4	
3	204	5 6	

Table C.14: Near radius=30.

PercImp = 3.7270

elapsed_time = 3.7582e+003

Candidate = 1 2 3 5 6 8 9 10 11 12 15 16 17 20 21
22 27 28 29 30 31 32 34 35 37 38 39 46 47 48 49
50




Figure C.14: Near radius=30.

Figure 4.14

Route#	Distance	Load	Route
1	176.890	174	0 10 8 1 0
2	123.782	185	0 7 6 0
3	97.868	139	0 4 2 0
4	129.495	82	0 9 3 0
5	50.260	129	0 5 0

Vehicle#	Route(s)	Distance
1	1	177
2	2	124
3	3	98
4	4 5	180

Table C.15: Initial routes.

Orientation	Route Order (vehicle 4)	number of close approaches
-1 -1 -1 -1 0	4 5	36
-1 -1 -1 -1 0	5 4	36
1 1 1 1 0	4 5	36
1 1 1 1 0	5 4	34
1 -1 -1 -1 0	4 5	37
1 -1 -1 -1 0	5 4	30
-1 1 -1 -1 0	4 5	38
-1 1 -1 -1 0	5 4	31
-1 -1 1 -1 0	4 5	38
-1 -1 1 -1 0	5 4	31
-1 -1 -1 1 0	4 5	36
-1 -1 -1 1 0	5 4	31
1 1 -1 -1 0	4 5	35
1 1 -1 -1 0	5 4	30
1 -1 1 -1 0	4 5	35
1 -1 1 -1 0	5 4	30
1 -1 -1 1 0	4 5	37
1 -1 -1 1 0	5 4	30
-1 1 1 -1 0	4 5	37
-1 1 1 -1 0	5 4	30
-1 1 -1 1 0	4 5	35
-1 1 -1 1 0	5 4	30
-1 -1 1 1 0	4 5	37
-1 -1 1 1 0	5 4	30
1 1 1 -1 0	4 5	37
1 1 1 -1 0	5 4	30
1 1 -1 1 0	4 5	35
1 1 -1 1 0	5 4	30
1 -1 1 1 0	4 5	37
1 -1 1 1 0	5 4	30
-1 1 1 1 0	4 5	35
-1 1 1 1 0	5 4	30

Table C.16: Orientation and route order.

Figure 4.16

maxVehCap=200, maxVehDist=220.

Node	Location		Demand
0	60	70	
1	19	70	85
2	68	54	125
3	30	44	159
4	54	69	144
5	15	62	57
6	69	79	39
7	37	95	68
8	86	52	81
9	85	88	71
10	59	17	129
11	49	97	66
12	89	27	119
13	82	25	120
14	64	87	1
15	81	73	15
16	66	13	160
17	34	1	25
18	28	89	105
19	34	19	141
20	53	29	80
21	72	66	52
22	30	28	21
23	83	46	92
24	56	6	200
25	37	98	5

Table C.17: Locations and demands.

1. Without coordination

Route#	Distance	Load	Route
1	153.199	187	0 22 19 17 0
2	95.650	142	0 5 1 0
3	86.914	178	0 18 25 7 0
4	71.573	173	0 23 8 0
5	129.514	199	0 20 12 0
6	109.875	192	0 6 15 9 14 11 0
7	104.941	172	0 21 13 0
8	35.777	125	0 2 0
9	79.398	159	0 3 0
10	12.166	144	0 4 0
11	106.019	129	0 10 0
12	114.630	160	0 16 0
13	<u>128.250</u>	200	0 24 0
	1227.904		
Vehicle#	Distance	Route(s)	
1	153	1	
2	201	4 5	
3	220	7 12	
4	220	9 10 13	
5	216	6 11	
6	218	2 3 8	

Table C.18: Initial routes.




Figure C.15: Initial routes.

2. Interim step




Figure C.16: Interim step.

Route#	Distance	Load	Route
1	153.199	187	0 22 19 17 0
2	95.650	142	0 5 1 0
3	86.914	178	0 18 25 7 0
4	108.667	200	0 12 8 0
5	120.958	200	0 20 13 0
6	109.875	192	0 6 15 9 14 11 0
7	68.716	144	0 21 23 0
8	35.777	125	0 2 0
9	79.398	159	0 3 0
10	12.166	144	0 4 0
11	106.019	129	0 10 0
12	114.630	160	0 16 0
13	<u>128.250</u>	200	0 24 0
	1220.219		
			continued...
Vehicle#	Distance	Route(s)	
1	201	1 8 10	
2	217	2 5	
3	215	3 13	
4	219	4 6	
5	183	7 12	
6	185	9 11	

Table C.19: Interim step.

3. Load splitting after interim step

New routes:

- 21 23 13* 0
where $\max(\text{Demand}(13^*)) = \text{maxVehCap} - [\text{Demand}(21) + \text{Demand}(23)]$
 $= 56$
 $\min(\text{Demand}(13^*)) = \text{Demand}(13) - [\text{maxVehCap} - \text{Demand}(2)]$
 $= 45$
- 2 13** 0
where $\text{Demand}(13^{**}) = \text{Demand}(13) - \text{Demand}(13^*)$

Route#	Distance	Load	Route
1	153.199	187	0 22 19 17 0
2	95.650	142	0 5 1 0
3	86.914	178	0 18 25 7 0
4	108.667	200	0 12 8 0
5	108.018	200	0 20 10* 0
6	109.875	192	0 6 15 9 14 11 0
7	106.589	189	0 21 23 13* 0
8	100.181	200	0 2 13** 0
9	79.398	159	0 3 0
10	12.166	144	0 4 0
11	118.386	169	0 16 10** 0
12	<u>128.250</u>	200	0 24 0
	1207.293		
Vehicle#	Distance	Route(s)	
1	165	1 10	
2	175	2 9	
3	215	3 12	
4	219	4 6	
5	215	5 7	
6	219	8 11	

Table C.21: Load splitting after interim step.

Distance reduced by ~1%.

4. Load splitting the non-coordinated (1)

New routes:

- 0 20 10* 0
where $\max(\text{Demand}(10^*)) = \text{maxVehCap} - \text{Demand}(20)$
 $= 120$
 $\min(\text{Demand}(10^*)) = \text{Demand}(10) - [\text{maxVehCap} - \text{Demand}(16)]$
 $= 89$
- 0 16 10** 0
where $\text{Demand}(10^{**}) = \text{Demand}(10) - \text{Demand}(10^*)$
- 0 21 8 23* 0

$$\begin{aligned}
 \text{where } \max(\text{Demand}(23^*)) &= \max\text{VehCap} - [\text{Demand}(21) + \text{Demand}(8)] \\
 &= 67 \\
 \min(\text{Demand}(23^*)) &= \text{Demand}(23) - [\max\text{VehCap} - \\
 &\quad \text{Demand}(12) - \text{Demand}(13^*)] \\
 &= \text{Demand}(2) + \text{Demand}(12) + \text{Demand}(13) + \\
 &\quad \text{Demand}(23) - 2\max\text{VehCap} \\
 &= 56
 \end{aligned}$$

- 0 23** 12 13* 0

$$\begin{aligned}
 \text{where } \text{Demand}(23^{**}) &= \text{Demand}(23) - \text{Demand}(23^*) \\
 \text{Demand}(13^*) &= \text{Demand}(13) - \text{Demand}(13^{**})
 \end{aligned}$$

- 0 2 13** 0

$$\begin{aligned}
 \text{where } \max(\text{Demand}(13^{**})) &= \max\text{VehCap} - \text{Demand}(2) \\
 &= 75 \\
 \min(\text{Demand}(13^{**})) &= 3\max\text{VehCap} - \text{Demand}(2) - 2\text{Demand}(12) \\
 &\quad - \text{Demand}(13) - \text{Demand}(23) \\
 &= 25
 \end{aligned}$$




Figure C.18: Load splitting the non-coordinated routes.

Route #	Distance	Load	Route
1	153.199	187	0 22 19 17 0
2	95.650	142	0 5 1 0
3	86.914	178	0 18 25 7 0
4	118.386	169	0 16 10 0
5	108.018	200	0 20 10 0
6	109.875	192	0 6 15 9 14 11 0
7	72.218	200	0 21 8 23 0
8	110.537	200	0 23 12 13 0
9	79.398	159	0 3 0
10	12.166	144	0 4 0
11	100.181	189	0 2 13 0
12	128.250	200	0 24 0
	1177.792		

Vehicle	Distance	Route(s)
1	165	1 10
2	218	3 12
3	219	4 11
4	219	5 8
5	206	2 6
6	152	7 9


Table C.22: Load splitting the non-coordinated routes.

5. Apply interim step to (4)


No improvement – the intersection involves a fully laden single-noded vehicle.

§4.14.1


$\text{Demand}(12) + \text{Demand}(23) > \text{maxVehCap}$




Node 13 is near an intersection




Nodes involved in creating the metastable state




metastable state




Node 23 is near an intersection




Nodes involved in creating the metastable state




metastable state



Node 12 is near an intersection



Nodes involved in creating the metastable state



metastable state




Figure C.19: Metastable state configurations.

§4.12.1

Node:	0	1	2	3	4	5	6	7	8	9	10
Location:	60,	26,	75,	30,	21,	25,	40,	75,	18,	63,	17,
	70	1	71	77	30	92	40	47	7	41	22
Demand:		13	2	20	70	75	100	6	107	3	1

Table C.23: Locations and demands.

Options are given in Table C.24, and show the routes in Figure C.20 as being shortest.




Figure C.20: Routes yielding greatest reductions.

The two part-routes following the coordination can be allocated to either vehicle.

	fixed	result
option 1	0 2 7 ...	Too long
Node 7 Orientation -1 -1	0 4 10 8 1 9 7 ...	Not accepted
option 2	0 2 7 ...	0 2 7 9 1 8 10 4 0
Node 7 Orientation -1 +1	0 7 ...	0 7 6 3 5 0
		Total 336.67
		Reduction 7.9%
option 3	0 6 7 ...	Too long
Node 7 Orientation +1 -1	0 4 10 8 1 9 7 ...	Not accepted
option 4	0 6 7 ...	0 6 7 2 5 3 0
Node 7 Orientation +1 +1	0 7 ...	0 7 9 1 8 10 4 0
		Total 336.06
		Reduction 0.7%
option 5	0 2 7 9 ...	Too long
Node 9 orientation -1 -1	0 4 10 8 1 9 ...	Not accepted
option 6	0 2 7 9 ...	0 2 7 9 6 3 5 0
Node 9 orientation -1 +1	0 9 ...	0 9 1 8 10 4 0
		Total 344.44
		Reduction 5.8%
option 7	0 6 9 ...	0 6 9 7 2 3 5 0
Node 9 orientation +1 -1	0 4 10 8 1 9 ...	0 4 10 8 1 9 0
		Total 367.94
		Not accepted
option 8	0 6 9 ...	0 6 9 7 2 3 5 0
Node 9 Orientation +1 +1	0 9 ...	0 9 1 8 10 4 0
		Total 367.43
		Not accepted
option 9		
No coordination		Total 365.49

Table C.24: Options.

APPENDIX D

DATA AND WORKED EXAMPLE FOR THE STANDARD PROBLEM IN
CHAPTER 5

Standard layout




Figure D.1: The standard layout for medium grid problems with depot located at D, (10,10).

Node locations and demands

Node	Location	Demand
1	10 9	43
2	9 11	166
3	8 10	12
4	7 8	28
5	6 8	114
6	6 10	119
7	6 11	118
8	7 14	3
9	5 15	32
10	9 18	17
11	6 6	18
12	8 4	72
13	5 3	58
14	13 7	200
15	14 8	150
16	3 12	25
17	12 3	42
18	16 3	78
19	19 6	50
20	20 10	120
21	17 11	134
22	13 14	76
23	19 15	80
24	16 16	40
25	14 16	100

Table D.1: Node locations and demands.

Application of coordination

Calculations for the system with maxVehCap=220; maxVehDist=50.

The original plan is shown in Table D.2:

Route#	Distance	Load	Route	Vehicle#	Distance	Route(s)
1	34	186	0 2 8 10 0	1	34	1
2	26	216	0 22 25 24 0	2	28	2
3	36	214	0 21 23 0	3	36	3
4	24	187	0 3 7 16 9 0	4	36	6
5	24	219	0 6 4 12 0	5	44	7
6	36	218	0 14 11 0	6	30	8
7	44	214	0 13 5 17 0	7	30	9
8	30	200	0 15 19 0	8	32	10
9	30	121	0 1 18 0	9	48	4 5
10	32	120	0 20 0			
	316	1895				

Table D.2: Initial load plan.

If the routes had not been improved, the fleet size would have been smaller.

Improved with split deliveries (Table D.3):

total load/maxVehCap = 1895/220 = 8.614, ie minimum of 9 routes required.

Route#	Distance	Load	Route	Surplus dist	Surplus load
1	34	186	0 2 8 10 0	16	34
2	26	216	0 22 25 24 0	24	4
3	36	214	0 21 23 0	14	6
4	24	187	0 3 7 16 9 0	26	33
5	24	219	0 6 4 12 0	26	1
6	36	218	0 14 11 0	14	2
7	44	214	0 13 5 17 0	6	6
8	30	200	0 15 19 0	20	20
9	30	121	0 1 18 0	20	99
10	32	120	0 20 0	18	100

Table D.3: Surplus table for initial plan improved with split deliveries.

1. Create surplus table.
2. Find max surplus dist 26
3. Of these find max surplus load 0 3 7 16 9 0 has 33
4. OK? Yes
5. Find node close to this route node 6
6. := target node
7. Demand(target node)? $D(6)=119$
8. surplus load-Demand(targetnode)>0? $119-33=86>0$
9. Adjust demand*(targetnode)
10. Find next max 26
11. surplus load? 1, ie No real use
12. OK? No
13. Find next max 22
14. surplus load? 4, ie No real use
15. OK? No
16. Find next max 20
17. surplus load 0 1 18 0 has 99
18. OK? Yes
19. Can target node be reached? Yes
20. surplus load-demand*(targetnode)>0? Yes
21. New routes

0 3 6* 7 16 9 0	dist 24	load 220
0 6* 1 18 0	dist 38	load 207
0 4 12 0	dist 18	load 100
19. New surplus table (Table D.4) (Same total distance)

Route#	Distance	Load	Route	Surplus dist	Surplus load
1	34	186	0 2 8 10 0	16	34
2	26	216	0 22 25 24 0	24	4
3	36	214	0 21 23 0	14	6
4	22	220	0 3 6* 7 16 9 0	28	0
5	18	100	0 4 12 0	32	120
6	36	218	0 14 11 0	14	2
7	44	214	0 13 5 17 0	6	6
8	30	200	0 15 19 0	20	20
9	38	207	0 6* 1 18 0	12	13
10	32	120	0 20 0	18	100

Table D.4: Surplus table.

Repeat process.

Max surplus distance = 32 for route 0 4 12 0.

Node close to route (preferably away from depot) = node 13 or node 11. Node 13 is on same route as node 5, which are both quite distant from other nodes, so try node 11 as target node.

Demand(11)=18. It can fit onto vehicle servicing route 0 4 12 0.

New route 0 4 11 12 0 dist 26 load 112

0 14 0 dist 20 load 200

Demand(13)=58 and can also fit onto this route.

New route 0 4 11 13 12 0 dist 28 load 170

0 14 0 dist 20 load 200

0 5 17 0 dist 40 load 156

New surplus table (Table D.5) has total distance 306:

Route#	Distance	Load	Route	Surplus dist	Surplus load
1	34	186	0 2 8 10 0	16	34
2	26	216	0 22 25 24 0	24	4
3	36	214	0 21 23 0	14	6
4	22	220	0 3 6* 7 16 9 0	28	0
5	28	170	0 4 11 13 12 0	22	108
6	20	200	0 14 0	30	20
7	40	156	0 5 17 0	10	6
8	30	200	0 15 19 0	20	20
9	38	207	0 6* 1 18 0	12	13
10	32	120	0 20 0	18	100

Table D.5: Surplus table with total distance 306.

Swap nodes 5 and 17 to get new routes:

0 12 17 0 dist 26 load 114

0 4 11 5 13 0 dist 34 load 218

New surplus table (Table D.6):

Route#	Distance	Load	Route	Surplus dist	Surplus load
1	34	186	0 2 8 10 0	16	34
2	26	216	0 22 25 24 0	24	4
3	36	214	0 21 23 0	14	6
4	22	220	0 3 6* 7 16 9 0	28	0
5	26	114	0 4 11 5 13 0	24	6
6	20	200	0 14 0	30	20
7	34	218	0 12 17 0	16	2
8	30	200	0 15 19 0	20	20
9	38	207	0 6* 1 18 0	12	13
10	32	120	0 20 0	18	100

Table D.6: Surplus table with total distance 298.

Total distance 298 (8 vehicles)

Repeat process.

Max surplus dist with 0 14 0 (30), etc.

Close to node 15.

Demand(15) = 150

Demand(15)-surplus load = $150 - 20 = 130$ [suggests 0 14 15 0 leaving 20 at node 15]

targetnode := 15, demand*(15) := 130

Other routes that can reach node 15 with max surplus load: 0 2 8 10 0

surplus load (0 2 8 10 0) = 34

[suggests 0 15 0 2 8 10 0 leaving 34 at coord site (co-located with node 15)]

demand*(15) := $130 - 34 = 96$

Other routes that can reach node 15 with max surplus load: 0 20 0

surplus load (0 20 0) = 100

demand*(15) - surplus load = -4 [suggests 0 15 20 0 leaving 96 at node 15]

New routes:

0 14 15* 0 dist 20 load 220

0 15* 0 2 8 10 0 dist 50 load 220

0 15* 20 0 dist 34 load 216

0 19 0 dist 30 load 50

New surplus table (Table D.7):

Route#	Dist	Load	Route	Surplus dist	Surplus load
1	50	220	0 15* 0 2 8 10 0	0	0
2	26	216	0 22 25 24 0	24	4
3	36	214	0 21 23 0	14	6
4	22	220	0 3 6* 7 16 9 0	28	0
5	26	114	0 4 11 5 13 0	24	6
6	20	220	0 14 15* 0	30	0
7	34	218	0 12 17 0	16	2
8	30	50	0 19 0	20	170
9	38	207	0 6* 1 18 0	12	13
10	34	216	0 15* 20 0	16	4

Table D.7: Surplus table with total distance 316.

Total distance 316 ie no improvement, so backtrack to previous surplus table.

No distance improvement by split delivery, so look at reducing distance by obviating return to depot for multi-routed vehicles.

Eight vehicles include 2 multi-routed vehicles, so create surplus table (Table D.8) for vehicles:

Vehicle#	Distance	Load	Route	Surplus dist	Surplus load
9	38	207	0 6* 1 18 0	12	13
3	36	214	0 21 23 0	14	6
1	34	186	0 2 8 10 0	16	34
7	34	218	0 12 17 0	16	2
10	32	120	0 20 0	18	100
8	30+20=50	200/200	0 15 19 0 14 0	20/30, ie 0	20/20
2	26+22=48	216/220	0 22 25 24 0 3 6* 7 16 9 0	24/28, ie 2	4/0
5	34	114	0 4 11 5 13 0	16	6

Table D.8: Surplus table for vehicles.

Try to create type 1 site. This involves finding a location that can be reached by a vehicle with surplus distance and spare carrying capacity, so that enough payload can be deposited to prevent the return to depot or can shorten another route.

0 18 20 0 dist 42 load 198
0 6* 1 0 dist 10 load 43+Demand(6*)=129

Alters 6* capability. Adjust. See Table D.9.

Vehicle#	Distance	Load	Route	Surplus dist	Surplus load
9	10	129	0 6* 1 0	40	91
3	36	214	0 21 23 0	14	6
1	34	186	0 2 8 10 0	16	34
7	34	218	0 12 17 0	16	2
10	42	198	0 18 20 0	8	22
8	30+20=50	200/200	0 15 19 0 14 0	20/30	20/20
2	26+22=48	216/187	0 22 25 24 0 3 7 16 9 0	24/28	4/33
5	34	114	0 4 11 5 13 0	16	6

Table D.9: Surplus table.

Vehicle 9 delivers 119/119 to node 6; Vehicle 2 delivers 0 to node 6 and is available to service node 8 without incurring additional distance; reschedule (Table D.10):

Route#	Distance	Load	Route	Surplus dist	Surplus load
1	32	183	0 2 10 0	18	37
2	24	190	0 3 7 16 9 8 0	26	30
3	36	214	0 21 23 0	14	6
4	26	216	0 22 25 24 0	22	4
5	34	114	0 4 11 5 13 0	16	6
6	30	200	0 15 19 0	20	20
7	34	218	0 12 17 0	16	2
8	20	200	0 14 0	30	20
9	10	162	0 6 1 0	40	58
10	42	198	0 18 20 0	8	22

Table D.10: Surplus table.

Now 7 vehicles. See the vehicle surplus table (Table D.11):

Vehicle#	Distance	Load	Route	Surplus dist	Surplus load
10	42	198	0 18 20 0	8	22
3,9	36 +10=46	214/162	0 21 23 0 6 1 0	4	6/58
5	34	114	0 4 11 5 13 0	16	6
7	34	218	0 12 17 0	16	2
1	32	183	0 2 10 0	18	37
6,8	30+20=50	200/200	0 15 19 0 14 0	0	20/20
4,2	26+24=50	216/190	0 22 25 24 0... 3 7 16 9 8 0	0	4/30

Table D.11: Surplus table with total distance 288.

Total distance: 288

Compare with original (see Table D.12). Yellow indicates changed route. Only three of the original routes are unchanged. No split deliveries. No type 1 sites.

Route#	Distance	Load	Route	Surplus dist	Surplus load
1	34	186	0 2 8 10 0	16	34
2	26	216	0 22 25 24 0	24	4
3	36	214	0 21 23 0	14	6
4	24	187	0 3 7 16 9 0	26	33
5	24	219	0 6 4 12 0	26	1
6	36	218	0 14 11 0	14	2
7	44	214	0 13 5 17 0	6	6
8	30	200	0 15 19 0	20	20
9	30	121	0 1 18 0	20	99
10	32	120	0 20 0	18	100

Table D.12: Comparison with original.

Now begin again. Try to create type 1 site to prevent return to depot.
Create surplus table (Table D.13).

Vehicle#	Distance	Load	Route	Surplus dist	Surplus load
10	42	198	0 18 20 0	8	22
3,9	36 +10=46	214/162	0 21 23 0 6 1 0	14/40	6/58
5	34	114	0 4 11 5 13 0	16	6
7	34	218	0 12 17 0	16	2
1	32	183	0 2 10 0	18	37
6,8	30+20=50	200/200	0 15 19 0 14 0	20/30	20/20
4,2	26+24=50	216/190	0 22 25 24 0... 3 7 16 9 8 0	22/26	4/30

Table D.13: Surplus table for re-constructed routes.

Three multi-routed vehicles. Look for any multi-routed vehicle whose total distance would be reduced if it did not have to return to the depot. For example, the path 0 21 23 0 6 1 0 travels almost through the depot, so there is no great advantage to made in altering this.

Consider 0 22 25 24 0 3 7 16 9 8 0. Because each route could be traversed in either direction, this can be read as any of

- (a) 0 22 25 24 0 3 7 16 9 8 0,
- (b) 0 24 25 22 0 3 7 16 9 8 0,
- (c) 0 22 25 24 0 8 9 16 7 3 0, and
- (d) 0 24 25 22 0 8 9 16 7 3 0, and

because the routes can be serviced in any order, (a)-(d) can be reversed.

Option (d) places the last node of the first route closer (in rectilinear terms) to the first node of the second route than the other options, so it will be considered first. The aim is to shorten the path 0 24 25 22 0 8 9 16 7 3 0 by creating a type 1 site somewhere close enough to the path to replace the 22 0 8 section with a shorter piece. The candidates for this include any route that has sufficient surplus distance to create a type 1 site as well as any route that has sufficient surplus capacity to create a type 2/3 site.

The route 0 2 10 0 has surplus distance of 18 and surplus load of 37, so it can create both type of coordination site: type 1 no more than 9 units from the depot, and type 2/3 site holding no more than 37 units no more than 9 units from its original path. The type 1 site is more valuable, since a full load could be deposited (190 is required), so such a site is created along the intersecting section, closest to the depot. This location is chosen because it uses a minimum amount of

the surplus distance and leaves a maximum surplus available for other coordination. The prime site is labelled X in the figure. The routes are then:

0 X 0 2 10 0	dist $8+32=40$
0 24 25 22 X 8 9 16 7 3 0	dist 46

Total distance is increased by 4, so there is no advantage in obviating this return to depot.

Consider the multi-routed path, 0 15 19 0 14 0, which has a vehicle servicing more than one route in the same general region. The possible paths are (a) 0 15 19 0 14 0, (b) 0 19 15 0 14 0, (c) 0 14 0 15 19 0, and (d) 0 14 0 19 15 0.

Again, route 0 2 10 0 can establish a type 1 site no more than 9 units from the depot. The maximum distance of advantage to the multi-routed vehicle places coordination sites at Y1 and Y2 in the figure. Y1 is of value in option (b) and (c); no use of Y2 creates as much advantage. Paths 0 19 15 Y1 14 0 and 0 14 Y1 15 19 0 have length 36. Coincidentally, Demand(14) has the same value (200) as Demand(15) + Demand(19), so a load of 200 will be deposited in the type 1 site at Y1 regardless of which path is adopted. If the values had been different, this might have had consequences for the further use of this site. The routes are then:

0 Y1 0 2 10 0	dist $18+32=50$
0 19 15 Y1 14 0	dist 36

Total distance is increased by 4 so site Y1 and the above paths are not adopted.

Generally, it is advantageous to have a coordination site between the end of one route and the start of another to reduce the amount of travel between routes, but the effect on the vehicle creating the type 1 site must also be considered. In the above case, type 1 sites can be created by any vehicle with surplus distance. Some, such as 0 18 20 0 with a surplus of 8 are not very valuable. Others, such as 0 2 10 0 have the capacity to create sites in between the target routes but the effect on the vehicle creating the site is onerous in comparison; overall distance increases. In this example, for the return to depot of the target vehicle to be removed, a type 1 site must be created that is outside the prime area, that is, away from the region of the grid between the end of one route and the beginning of the other route. The other surplus distances will now be checked.




Figure D.2: Coordination sites are indicated by arrows.

The surplus table is shown in Table D.14.

Vehicle#	Distance	Load	Route	Surplus dist	Surplus load	checked
10	42	198	0 18 20 0	8	22	-
3,9	36 +10=46	214/162	0 21 23 0 6 1 0	4	6/58	y
5	34	114	0 4 11 5 13 0	16	6	-
7	34	218	0 12 17 0	16	2	-
1	32	183	0 2 10 0	18	37	-
6,8	30+20=50	200/200	0 15 19 0 14 0	0	20/20	
4,2	26+24=50	216/190	0 22 25 24 0... 3 7 16 9 8 0	0	4/30	y

Table D.14: Surplus table.

In order to benefit from the obviation of the rtd for 0 19 15 0 14 0 by the creation of a type 1 site at, say, Y, it is necessary for

$$\text{distance}(0\ 19\ 15\ Y\ 14\ 0) + \text{distance}(0\ Y\ 0) < \text{distance}(0\ 19\ 15\ 0\ 14\ 0)$$

ie for

$$\text{distance}(15\ Y\ 14) + \text{distance}(0\ Y\ 0) < \text{distance}(15\ 0\ 14)$$

where

$$\max(\text{distance}(0\ Y\ 0)) = \text{surplus distance}$$

and $\text{distance}(15\ 0\ 14) = 18.$

There is no location that satisfies this.

Consider traversing the multi-noded route in the opposite direction and obviating the rtd for 0 15 19 0 14 0 by creating a type 1 site at Y. Then

$$\text{distance}(19\ Y\ 14) + \text{distance}(0\ Y\ 0) < \text{distance}(19\ 0\ 14)$$

where $\max(\text{distance}(0\ Y\ 0)) = \text{surplus distance}$

and $\text{distance}(19\ 0\ 14) = 25.$

$$\begin{aligned} \text{Using the site at Y3, } \text{distance}(19\ Y3\ 14) + \text{distance}(0\ Y3\ 0) &= 11 + 14 \\ &= 25 \end{aligned}$$

so Y3 cannot act as a type 1 site.

Other multi-routed combinations could exist with 0 6 1 0 and one of 0 4 11 5 13 0, 0 12 17 0, and 0 2 10 0, without increasing the number of vehicles. Attempting to obviate return to depot with any of these is handicapped by the adjacency of node 1 to the depot. Any vehicle that had to add distance to establish a type 1 site would increase the overall distance.

Since there are no other multi-routed vehicles, the final routes are shown in Table D.15, with total distance 288.

Vehicle#	Distance	Load	Route	Surplus dist	Surplus load
10	42	198	0 18 20 0	8	22
3,9	36 +10=46	214/162	0 21 23 0 6 1 0	4	6/58
5	34	114	0 4 11 5 13 0	16	6
7	34	218	0 12 17 0	16	2
1	32	183	0 2 10 0	18	37
6,8	30+20=50	200/200	0 15 19 0 14 0	0	20/20
4,2	26+24=50	216/190	0 22 25 24 0... 3 7 16 9 8 0	0	4/30

Table D.15: Final routes and their surplus table, with total distance 288.



APPENDIX E

COMPARISON OF GRID DENSITIES

Primes indicate split deliveries. X^{del} indicates a delivery to a coordination site at X ; X_{col} indicates a collection from a coordination site at X ; where ambiguity may otherwise arise, the subscripts and superscripts include quantities.

Sparse grid

maxVehCap=235, maxVehDist=50.

The investigation includes systems of two nodes with demand 150 and all other nodes with equal demand:

- All demands=150
- Demand(8)=Demand(21)=150.




Figure E.1: Sparse grid layout. The coordination sites indicated are determined during the improvement process and are not pre-selected.

All demands=150

Route	Load	Distance
0 1 0	150	20
0 2 0	150	12
0 3 0	150	26
0 4 0	150	28
0 5 0	150	30
0 6 0	150	28
0 7 0	150	16
0 8 0	150	28
0 9 0	150	38
0 10 0	150	30
0 11 0	150	22
0 12 0	150	14
0 13 0	150	28
0 14 0	150	36
0 15 0	150	34
0 16 0	150	24
0 17 0	150	20
0 18 0	150	16
0 19 0	150	16
0 20 0	150	24
0 21 0	150	28
0 22 0	150	16
0 23 0	150	10
0 24 0	150	6
0 25 0	<u>150</u>	<u>8</u>
	3750	558

Table E.1: No coordination. Fleet size=12.

Splitting			Other forms	
Route	Load	Distance	Route	Distance
0 20 23' 0	235	24	0 20 $V^{\text{del } 85}$ 0	24
0 21 23" 0	215	28	0 21 $V_{\text{col } 85}$ 22 0	30
0 22 0	150	10	0 23 0	10
0 14 12' 0	235	36	0 14 $R^{\text{del } 85}$ 0	36
0 13 12" 24' 0	235	28	0 13 $R_{\text{col } 65}$ 10 0	42
0 16 18' 0	235	26	0 15 $R^{\text{del } 85}$ 0	34
0 17 18" 0	215	22	0 16 $R_{\text{col } 105}$ 11 $P^{\text{del } 45}$ 0	32
0 19 0	150	16	0 9 $N^{\text{del } 85}$ 0	38
0 10 24" 0	235	32	0 8 $N_{\text{col } 65}$ 6 0	34
0 11 24''' 0	195	22	0 5 $N_{\text{col } 20}$ $P^{\text{del } 20}$ 0	30
0 9 25' 0	235	38	0 12 $P_{\text{col } 65}$ 24 0	14
0 8 25" 0	215	28	0 17 $U^{\text{del } 85}$ 0	20
0 4 7' 0	235	28	0 18 U_{col} 19 0	20
0 3 7" 0	215	26	0 4 L^{del} 0	28
0 1 0	150	20	0 3 L_{col} 7 0	26
0 2 0	150	12	0 1 Y^{del} 0	20
0 5 0	150	30	0 2 Y_{col} 25 0	<u>16</u>
0 6 0	150	28		454
0 15 0	<u>150</u>	<u>34</u>		
	3750	488		

Table E.2: Improvement by coordination. Fleet size=12 (splitting) and 10 (other forms).

Demand(8)=Demand(21)=150

(i) All other demands=100

Route	Load	Distance
0 20 22 0	200	26
0 21 0	150	28
0 23 0	100	10
0 13 14 0	200	36
0 15 12 0	200	34
0 16 17 0	200	24
0 18 19 0	200	20
0 10 11 0	200	30
0 9 7 0	200	42
0 8 0	150	28
0 5 6 0	200	32
0 3 4 0	200	30
0 1 2 0	200	20
0 25 24 0	<u>200</u>	<u>10</u>
	2600	370

Table E.3: No coordination. Fleet size=9.

Splitting			Other forms	
Route	Load	Distance	Route	Distance
0 20 22 23' 0	235	26	0 20 22 $W^{\text{del } 35}$ 0	26
0 21 23" 0	215	28	0 21 $W_{\text{col } 35}$ 23 0	28
0 3 4 7' 0	235	30	0 9 $N^{\text{del } 35}$ 6 0	44
0 5 6 7" 0	235	36	0 8 $N_{\text{col } 35}$ 5 0	36
0 9 7''' 25 0	230	42	0 10 11 $P^{\text{del } 35}$ 0	30
0 8 24' 0	235	28	0 13 14 $P^{\text{del } 35}$ 0	36
0 13 14 24" 0	215	36	0 15 12 $P_{\text{col } 70}$ 7 0	42
0 15 12 0	200	34	0 16 17 0	24
0 16 17 0	200	24	0 18 19 0	20
0 1 2 0	200	20	0 4 3 0	30
0 18 19 0	200	20	0 1 2 0	20
0 10 11 0	<u>200</u>	<u>30</u>	0 25 24 0	<u>10</u>
	2600	354		346

Table E.4: Coordination. Fleet sizes: 9 (splitting), 8 (other forms).

(ii) All other demands=90

Route	Load	Distance
0 20 22 0	180	26
0 21 0	150	28
0 23 0	90	10
0 13 14 0	180	36
0 15 12 0	180	34
0 16 17 0	180	24
0 18 19 0	180	20
0 10 11 0	180	30
0 9 7 0	180	42
0 8 0	150	28
0 5 6 0	180	32
0 3 4 0	180	30
0 1 2 0	180	20
0 25 24 0	<u>180</u>	<u>10</u>
	2370	370

Table E.5: No coordination. Fleet size=9.

Splitting			Other forms	
Route	Load	Distance	Route	Distance
0 20 22 23' 0	235	26	0 20 22 $V^{del\ 55}$ 0	26
0 21 23" 0	185	28	0 21 V_{col} 23 0	28
0 3 4 7' 0	235	30	0 5 6 $Q^{del\ 55}$ 0	34
0 5 6' 9 0	235	48	0 8 Q_{col} 9 0	42
0 8 6" 0	185	34	0 10 11 0	30
0 10 11 12' 0	235	30	0 13 14 S_{col} 15 0	44
0 13 14 12" 0	215	36	0 16 17 $S^{del\ 55}$ 0	34
0 15 24 0	180	34	0 18 19 0	20
0 16 17 0	180	24	0 4 3 $Z^{del\ 55}$ 0	30
0 1 2 7" 0	215	24	0 1 2 0	20
0 18 19 0	180	20	0 24 Z_{col} 7 25 0	<u>16</u>
0 25 0	<u>90</u>	<u>8</u>		324
	2370	342		

Table E.6: Coordination. Fleet sizes: 9 (splitting), 8 (other forms).

(iii) All other demands=80

Route	Load	Distance
0 20 21 0	230	26
0 22 23 0	160	16
0 18 19 0	160	20
0 13 14 0	160	36
0 10 11 0	160	30
0 8 9 0	230	42
0 15 12 0	160	34
0 5 6 0	160	32
0 3 4 0	160	30
0 1 2 0	160	20
0 7 25 0	160	16
0 16 17 0	160	24
0 24 0	<u>80</u>	<u>6</u>
	2140	332

Table E.7: No coordination. Fleet size =7.

Splitting			Other forms	
Route	Load	Distance	Route	Distance
0 20 21 0	230	26	0 20 21 0	26
0 22 23 0	160	16	0 22 23 0	16
0 15 16 17' 0	235	32	0 17 16 $S^{\text{del } 75}$ 0	26
0 17" 18 19 0	165	26	0 12 15 $S_{\text{col } 75}$ $R^{\text{del } 150}$ 0	34
0 13 14 12' 0	235	36	0 13 14 $R_{\text{col } 150}$ 10 11 0	50
0 10 11 12" 0	165	30	0 8 9 $N^{\text{del } 5}$ 0	42
0 8 9 0	230	42	0 5 6 $N_{\text{col } 5}$ 7 0	36
0 5 6 7' 0	235	36	0 3 4 0	30
0 3 4 7" 0	165	30	0 1 2 0	20
0 1 2 0	160	20	0 24 25 0	10
0 24 25 0	<u>160</u>	<u>10</u>	0 18 19 0	<u>20</u>
	2140	304		310

Table E.8: Coordination. Fleet sizes: 8 (splitting), 7 (other forms).

(iv) All other demands=70

Route	Load	Distance
0 20 21 0	220	26
0 22 23 0	140	16
0 8 9 0	220	42
0 3 4 7 0	210	30
0 5 6 25 0	210	32
0 10 11 12 0	210	30
0 13 14 24 0	210	36
0 15 16 17 0	210	34
0 18 19 0	140	20
0 1 2 0	<u>140</u>	<u>20</u>
	1910	286

Table E.9: No coordination. Fleet size=7.

No improvement by splitting. Other forms of coordination yield a fleet size of 8.

Route	Distance
0 20 21 $X^{\text{del } 15}$ 0	26
0 22 23 $X^{\text{del } 25}$ 0	16
0 3 4 7 $T^{\text{del } 25}$ 0	30
0 8 9 $T^{\text{del } 15}$ 0	42
0 5 6 25 0	32
0 10 11 12 $T^{\text{del } 25}$ 0	30
0 24 13 14 $T_{\text{col } 65}$ 2 0	40
0 15 16 17 $X^{\text{del } 25}$ 0	34
0 18 19 $X_{\text{col } 65}$ 1 0	<u>38</u>
	288

Table E.10: Improvement by coordination other than split delivery. Fleet size=8.

(v) All other demands=60

Route	Load	Distance
0 20 21 0	210	26
0 22 23 0	120	16
0 8 9 0	210	42
0 3 4 7 0	180	30
0 5 6 25 0	180	32
0 10 11 12 0	180	30
0 13 14 24 0	180	36
0 15 16 17 0	180	34
0 18 19 0	120	20
0 1 2 0	<u>120</u>	<u>20</u>
	1680	286

Table E.11: No coordination. Fleet size=7.

No improvement by splitting. Other forms of coordination yield a fleet size of 7.

Route	Distance
0 20 21 $X^{\text{del } 25}$ 0	26
0 22 23 $X_{\text{col } 25}$ 1 2 0	34
0 8 9 $P^{\text{del } 25}$ 0	42
0 10 11 12 $P_{\text{col } 25}$ 5 6 0	50
0 3 4 7 $T^{\text{del } 55}$ 0	30
0 15 16 17 $T_{\text{col } 55}$ 25 0	36
0 13 14 24 0	36
0 18 19 0	<u>20</u>
	274

Table E.12: Improvement by coordination other than split delivery. Fleet size=7.

(vi) All other demands=50

Route	Load	Distance
0 20 21 0	200	26
0 22 23 1 2 0	200	34
0 3 4 7 25 0	200	30
0 5 6 24 0	150	32
0 8 9 0	200	42
0 10 11 13 14 0	200	50
0 16 17 18 19 0	200	30
0 15 12 0	<u>100</u>	<u>34</u>
	1450	278

Table E.13: No coordination. Fleet size=8.

Splitting			Other forms	
Route	Load	Distance	Route	Distance
0 20 21 0	200	26	0 20 21 $X^{\text{del } 35}$ 0	26
0 22 23 1 2 0	200	34	0 22 23 $X_{\text{col } 35}$ 1 2 7 $M^{\text{del } 20}$ 0	42
0 3 4 7 5 6' 0	235	50	0 3 4 5 6 $M_{\text{col } 20}$ 25 0	50
0 8 9 6" 0	215	48	0 8 9 $P^{\text{del } 35}$ 0	42
0 10 11 12 25 0	200	32	0 10 11 12 $P_{\text{col } 35}$ 13 14 0	50
0 13 14 24 0	150	36	0 24 15 16 17 0	34
0 15 16 17 0	150	34	0 18 19 0	<u>20</u>
0 18 19 0	<u>100</u>	<u>20</u>		264
	1450	280		

Table E.14: Improvement by coordination. Fleet size= 7 (splitting) and 6 (other forms).

Medium grid

maxVehCap=235, maxVehDist=50

The investigation involves a single depot located at either D1 or D2 in Figure E.2.




Figure E.2: Medium density grid.

Some coordination sites are co-located with nodes, and denoted by "(node)", as indicated in Figures E.3 and E.4.

Depot at D1




Figure E.3: Coordination site locations used in improvement by coordination for systems with depot at D1. The sites are generated within the improvement process and are not pre-selected.

The investigation includes systems of several nodes with demand 150 and all other nodes with equal demand:

- All demands=150
- Demand(13)=Demand(16)=Demand(18)=Demand(23)=150
- Demand(16)=Demand(19)=150
- Demand(13)=Demand(23)=150.

All demands=150

Route	Load	Distance
0 1 0	150	2
0 2 0	150	4
0 3 0	150	4
0 4 0	150	10
0 5 0	150	28
0 6 0	150	8
0 7 0	150	10
0 8 0	150	14
0 9 0	150	20
0 10 0	150	30
0 11 0	150	24
0 12 0	150	16
0 13 0	150	24
0 14 0	150	20
0 15 0	150	16
0 16 0	150	18
0 17 0	150	26
0 18 0	150	30
0 19 0	150	30
0 20 0	150	32
0 21 0	150	24
0 22 0	150	14
0 23 0	150	32
0 24 0	150	24
0 25 0	<u>150</u>	<u>20</u>
	3750	480

Table E.15: No coordination. Fleet size=10.

Splitting			Other forms	
Route	Load	Distance	Route	Distance
0 5 11' 0	235	30	0 5 E ^{del 85} 0	28
0 11" 13 0	215	28	0 4 E _{col 65} 11 0	10
0 4 12' 0	235	18	0 13 E _{col 20} G ^{del 105} 0	24
0 12" 17 0	215	26	0 12 G _{col 105} 17 K ^{del 40} 0	30
0 14' 19 0	235	32	0 14 K ^{del 85} 0	20
0 14" 18 0	215	32	0 15 K _{col 125} 18 L ^{del 60} 0	32
0 15' 20 0	235	34	0 19 L _{col 60} N ^{del 145} 0	34
0 15" 21 0	215	26	0 20 N _{col 145} 21 S ^{del 80} 0	40
0 22' 23 0	235	32	0 22 S _{col 80} 23 Q ^{del 15} 0	34
0 22" 24 0	215	24	0 24 Q _{col 15} (25) ^{del 100} 0	28
0 25 0	150	20	0 25 (25) _{col 100} 10 (8) ^{del 35} 0	34
0 10 8' 0	235	32	0 8 (8) ^{del 85} 0	14
0 8" 9 0	215	20	0 9 R _{col 120} 7 (7) ^{del 55} 0	22
0 16 6' 0	235	18	0 16 (7) _{col 55} (6) ^{del 140} 0	18
0 7 6" 0	215	10	0 6 (6) _{col 140} 3 0	8
0 3 0	150	4	0 1 2 0	<u>6</u>
0 2 0	150	4		382
0 1 0	<u>150</u>	<u>2</u>		
	3750	392		

Table E.16: Improvement by coordination. Fleet size=9 (splitting and other forms).

Demand(13)=Demand(16)=Demand(18)=Demand(23)=150

(i) All other demands=100

Route	Load	Distance
0 5 11 0	200	32
0 13 0	150	24
0 4 12 0	200	18
0 14 17 0	200	28
0 18 0	150	30
0 15 19 0	200	32
0 20 21 0	200	32
0 23 0	150	32
0 24 25 0	200	26
0 22 10 0	200	32
0 16 0	150	18
0 8 9 0	200	20
0 6 7 0	200	10
0 3 2 0	200	6
0 1 0	<u>100</u>	<u>2</u>
	2700	344

Table E.17: No coordination. Fleet size=7.

Splitting			Other forms	
Route	Load	Distance	Route	Distance
0 5 11 4' 0	235	32	0 5 11 E ^{del 35} 0	32
0 4" 13 0	215	24	0 E ^{del 35} 12 17 0	28
0 12 17 0	200	26	0 4 E _{col 70} 13 0	24
0 15' 14 19 0	235	32	0 14 15 K ^{del 35} 0	20
0 15" 18 0	215	32	0 19 K _{col 35} 18 0	42
0 20 21 0	200	32	0 20 21 P ^{del 35} 0	32
0 22' 23 0	215	32	0 23 P _{col 35} 24 0	38
0 22" 24 25 0	235	26	0 22 25 0	20
0 10 8 0	200	34	0 8 9 C ^{del 35} 0	20
0 9 7 6' 0	235	20	0 16 C _{col 35} 7 0	18
0 16 6" 0	215	18	0 6 3 0	8
0 2 3 0	200	6	0 10 2 0	32
0 1 0	<u>100</u>	<u>2</u>	0 1 0	<u>2</u>
	2700	316		316

Table E.18: Improvement by coordination. Fleet size=8.

(ii) All other demands=90

Route	Load	Distance
0 5 11 0	180	32
0 13 0	150	24
0 4 12 0	180	18
0 14 15 0	180	20
0 17 19 0	180	38
0 18 0	150	30
0 20 21 0	180	32
0 23 0	150	32
0 24 25 0	180	20
0 22 10 0	180	32
0 8 9 0	180	20
0 16 0	150	18
0 6 7 0	180	10
0 3 2 0	180	6
0 1 0	<u>90</u>	<u>2</u>
	2490	334

Table E.19: No coordination. Fleet size=8.

Splitting			Other forms	
Route	Load	Distance	Route	Distance
0 5 11 12' 0	235	32	0 5 11 E ^{del 55} 0	32
0 13 12" 0	185	24	0 4 E _{col 55} 13 0	24
0 14 17 0	180	28	0 15 14 K ^{del 55} 0	20
0 4 1 0	180	10	0 19 K _{col 55} 18 0	36
0 15 19 0	180	32	0 12 17 0	26
0 18 0	150	30	0 20 21 P ^{del 55} 0	32
0 20 21 0	180	32	0 23 P _{col 55} 24 0	34
0 22' 23 0	235	32	0 22 25 0	20
0 22" 25 24 0	185	20	0 10 8 D ^{del 55} 0	32
0 10 8 0	180	32	0 3 D _{col 55} 16 0	18
0 9 7 6' 0	235	20	0 9 7 0	20
0 16 6" 0	185	18	0 6 2 0	8
0 2 3 0	<u>180</u>	<u>6</u>	0 1 0	<u>2</u>
	2490	316		304

Table E.20: Improvement by coordination. Fleet size=8 (splitting and other forms).

(iii) All other demands=80

Route	Load	Distance
0 5 11 0	160	32
0 4 13 0	230	24
0 12 17 0	160	26
0 14 15 0	160	20
0 18 19 0	230	36
0 20 21 0	160	32
0 23 24 0	230	34
0 22 25 0	160	20
0 10 8 0	160	34
0 16 9 0	230	24
0 7 6 0	160	10
0 3 2 0	160	6
0 1 0	<u>80</u>	<u>2</u>
	2280	300

Table E.21: No coordination. Fleet size=8.

Splitting			Other forms	
Route	Load	Distance	Route	Distance
0 5 11 4' 0	235	32	0 5 11 E ^{del 75} 0	32
0 4" 13 12 0	235	24	0 4 E _{col 75} 13 12 0	24
0 14' 18 17 0	235	34	0 17 18 K ^{del 5} 0	32
0 15 14" 19 0	235	32	0 15 14 K _{col 5} 19 0	32
0 20 21 0	160	32	0 T ^{del 75} 21 20 0	32
0 22' 23 24 0	235	34	0 22 T _{col 75} 23 24 0	36
0 22" 25 10 0	235	32	0 25 10 (8) ^{del 75} 0	32
0 2 3 0	160	6	0 8 (8) _{col 75} 9 7 0	20
0 8 9 7' 0	235	20	0 16 6 0	18
0 16 7" 6 0	235	18	0 3 2 0	6
0 1 0	<u>80</u>	<u>2</u>	0 1 0	<u>2</u>
	2280	266		266

Table E.22: Improvement by coordination. Fleet size=7.

(iv) All other demands=70

Route	Load	Distance
0 5 11 4 0	210	32
0 13 12 0	220	24
0 14 15 19 0	210	32
0 17 18 0	220	32
0 20 21 0	140	32
0 23 24 0	220	34
0 22 25 10 0	210	32
0 16 9 0	220	24
0 8 7 6 0	210	16
0 1 2 3 0	<u>210</u>	<u>8</u>
	2070	266

Table E.23: No coordination. Fleet size=7.

Splitting			Other forms	
Route	Load	Distance	Route	Distance
0 1' 5 11 4 0	235	32	0 (1) ^{del} 5 11 4 0	32
0 1" 13 12 0	235	24	0 (1) ^{del} 13 12 0	24
0 1''' 14 15 19 0	235	32	0 (1) ^{del} 14 15 19 0	32
0 1'''' 17 18 0	225	32	0 (1) _{col} 17 18 0	32
0 2 3 0	140	6	0 2 3 0	6
0 20 21 0	140	32	0 20 21 0	32
0 23 24 0	220	34	0 23 24 0	34
0 22 25 10 0	210	32	0 22 25 10 0	32
0 16 9 0	220	24	0 16 9 0	24
0 8 7 6 0	<u>210</u>	<u>16</u>	0 8 7 6 0	<u>16</u>
	2280	266		266

Table E.24: Improvement by coordination. Fleet size=7.

(v) All other demands=60

Route	Load	Distance
0 5 11 4 0	180	30
0 13 12 0	210	24
0 17 18 0	210	32
0 14 15 19 0	180	32
0 20 21 0	120	32
0 23 24 0	210	34
0 22 25 10 0	180	32
0 16 9 0	210	24
0 8 7 6 0	180	16
0 1 2 3 0	<u>180</u>	<u>8</u>
	1860	264

Table E.25: No coordination. Fleet size=7.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 1' 5 11 4 0	235	30	0 4 5 11 $E^{del 55}$ 0	30
0 1" 13 12 0	215	24	0 13 $E_{col 55}$ 12 17 $K^{del 20}$ 0	36
0 17 18 0	210	32	0 14 15 $K_{col 20}$ 20 21 $T^{del 15}$ 0	38
0 14 15 19 0	180	32	0 18 19 0	36
0 20 21 0	120	32	0 22 $T_{col 15}$ 25 10 8 $B^{del 10}$ 0	34
0 23 24 0	210	34	0 23 24 0	34
0 22 25 10 0	180	32	0 9 16 $B_{col 10}$ 7 0	26
0 16 9 0	210	24	0 6 3 2 0	10
0 8 7 6 0	180	16	0 1 0	<u>2</u>
0 2 3 0	<u>120</u>	<u>6</u>		246
	1860	262		

Table E.26: Improvement by coordination. Fleet size=7 (splitting and other forms).

(vi) All other demands=50

Route	Load	Distance
0 4 5 11 0	150	30
0 1 13 0	200	24
0 18 19 0	100	36
0 14 15 17 12 0	200	28
0 23 24 0	200	34
0 20 21 25 22 0	200	38
0 10 9 8 3 0	200	36
0 16 2 0	200	20
0 7 6 0	<u>200</u>	<u>12</u>
	1650	258

Table E.27: No coordination. Fleet size=7.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 4 5 11 12 1' 0	235	32	0 4 5 11 $F^{del\ 35}$ 12 0	32
0 13 17 1" 0	235	34	0 13 $F_{col\ 35}$ 17 14 $U^{del\ 20}$ 0	36
0 18 19 1'" 0	230	36	0 15 $U_{col\ 20}$ 19 18 0	38
0 14 15 20 21 22' 0	235	40	0 20 21 24 25 $(25)^{del\ 35}$ 0	40
0 22" 23 24 0	235	34	0 22 23 $(25)_{col\ 35}$ 10 0	48
0 22'" 25 10 9 8 0	230	38	0 9 8 7 6 $(3)^{del\ 35}$ 0	20
0 16 7 0	100	18	0 16 3 $(3)_{col\ 35}$ 2 0	20
0 6 3 2 0	<u>150</u>	<u>10</u>	0 1 0	<u>2</u>
	1650	242		236

Table E.28: Improvement by coordination. Fleet size=6 (splitting and other forms).

Demand(16)=Demand(19)=150

(i) All other demands=100

Route	Load	Distance
0 5 11 0	200	30
0 13 12 0	200	24
0 17 18 0	200	32
0 14 15 0	200	20
0 19 0	150	30
0 1 4 0	200	10
0 20 21 0	200	32
0 23 24 0	200	34
0 22 25 0	200	20
0 16 0	150	18
0 8 9 0	200	20
0 6 7 0	200	12
0 3 2 0	200	6
0 10 0	<u>100</u>	<u>30</u>
	2600	318

Table E.29: No coordination. Fleet size=7.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4' 0	235	30	0 5 11 C ^{del 35} 0	30
0 4" 13 12 0	235	24	0 13 12 C ^{del 35} 0	24
0 4''' 17 18 0	230	34	0 4 C _{col 70} 17 18 0	34
0 15' 19 0	235	32	0 14 15 M ^{del 35} 0	22
0 15" 14 1 0	215	20	0 19 M _{col 35} 21 S ^{del 20} 0	38
0 20 21 0	200	32	0 20 23 S ^{del 35} 0	42
0 23 24 0	200	34	0 25 24 S ^{del 35} 0	26
0 22 25 0	200	20	0 22 S _{col 90} 10 8 D ^{del 25} 0	34
0 16 3' 0	235	18	0 16 D ^{del 85} 0	18
0 8 9 3" 0	215	20	0 9 D _{col 100} 7 6 0	20
0 7 6 0	200	12	0 3 2 0	6
0 2 10 0	<u>200</u>	<u>32</u>	0 1 0	<u>2</u>
	2600	308		296

Table E.30: Improvement by coordination. Fleet size=8 (splitting) and 7 (other forms).

(ii) All other demands=90

Route	Load	Distance
0 16 0	150	18
0 19 0	150	30
0 5 11 0	180	30
0 4 13 0	180	24
0 1 12 0	180	16
0 15 14 0	180	20
0 17 18 0	180	32
0 21 20 0	180	32
0 23 24 0	180	34
0 22 25 0	180	20
0 10 0	90	30
0 8 9 0	180	20
0 6 7 0	180	12
0 2 3 0	<u>180</u>	<u>6</u>
	2370	324

Table E.31: No coordination. Fleet size=7.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4' 0	235	30	0 5 11 E ^{del 55} 0	30
0 4" 13 12 0	215	24	0 4 E _{col 55} 13 12 0	24
0 15' 19 0	235	30	0 15 14 U ^{del 55} 0	20
0 17 18 14' 0	235	34	0 17 18 U ^{del 55} 0	32
0 21 20 15" 14" 0	220	38	0 19 U _{col} 20 21 0	48
0 22' 23 24 0	235	34	0 R ^{del 55} 24 23 0	34
0 22" 25 10 0	215	32	0 22 R _{col 55} 25 10 0	32
0 8 9 7' 0	235	20	0 8 9 C ^{del 55} 0	20
0 16 7" 6' 0	235	18	0 16 C _{col 55} 7 (6) ^{del 50} 0	18
0 6" 3 2 0	220	10	0 6 (6) _{col 50} 3 2 0	10
0 1 0	<u>90</u>	<u>2</u>	0 1 0	<u>2</u>
	2370	272		270

Table E.32: Improvement by coordination. Fleet size=7.

(iii) All other demands=80

Route	Load	Distance
0 5 11 0	160	30
0 13 4 0	160	24
0 12 17 0	160	26
0 14 15 0	160	20
0 18 19 0	230	36
0 20 21 0	160	32
0 22 23 0	160	32
0 24 25 0	160	26
0 10 8 0	160	32
0 9 16 0	230	24
0 7 6 0	160	12
0 2 3 0	160	6
0 1 0	<u>80</u>	<u>2</u>
	2140	302

Table E.33: No coordination. Fleet size=8.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4' 0	235	30	0 5 11 E ^{del 75} 0	30
0 4" 13 12 0	165	24	0 4 E _{col 75} 13 12 0	24
0 17 14 15' 0	235	28	0 17 18 K ^{del 75} 0	32
0 15" 18 19 0	235	38	0 15 14 K _{col 75} 19 0	32
0 21 20 22' 0	235	34	0 T ^{del 75} 21 20 0	32
0 22" 23 24 25' 0	235	26	0 22 T _{col 75} 23 24 0	36
0 25" 10 8 3' 0	235	34	0 25 10 (8) 0	32
0 9 16 0	230	24	0 8 (8) _{col 75} 9 7 0	20
0 7 6 3" 0	175	12	0 16 6 0	18
0 1 2 0	<u>160</u>	<u>6</u>	0 3 2 0	6
	2140	256	0 1 0	<u>2</u>
				264

Table E.34: Improvement by coordination. Fleet size=7.

(iv) All other demands=70

Route	Load	Distance
0 5 11 13 0	210	34
0 4 12 17 0	210	28
0 1 14 15 0	210	20
0 18 19 0	220	36
0 20 21 22 0	210	34
0 23 24 25 0	210	36
0 10 8 3 0	210	32
0 9 16 0	220	24
0 7 6 2 0	<u>210</u>	<u>10</u>
	1910	254

Table E.35: No coordination. Fleet size=7.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 13 4' 0	235	34	0 5 11 13 (12) ^{del 25} 0	34
0 4" 12 17 14' 0	235	30	0 4 12 (12) _{col 25} 17 (14) ^{del 50} 0	28
0 14" 18 19 0	235	36	0 15 14 (14) _{col 50} 20 21 0	34
0 14'" 15 20 21 22' 0	235	40	0 18 19 0	36
0 22" 23 24 25 0	235	36	0 R ^{del 25} 23 24 25 0	36
0 22'" 10 8 3 0	235	32	0 22 R _{col 25} 10 8 D ^{del 50} 0	34
0 9 16 0	220	24	0 9 16 0	24
0 7 6 2 0	210	10	0 7 6 D _{col} 3 2 0	10
0 1 0	<u>70</u>	<u>2</u>	0 1 0	<u>2</u>
	1910	244		238

Table E.36: Improvement by coordination. Fleet size=7.

(v) All other demands=60

Route	Load	Distance
0 5 11 13 0	180	34
0 4 12 17 0	180	28
0 1 14 15 0	180	20
0 18 19 0	210	36
0 20 21 22 0	180	34
0 23 24 25 0	180	36
0 10 8 3 0	180	32
0 9 16 0	210	24
0 7 6 2 0	<u>180</u>	<u>10</u>
	1680	254

Table E.37: No coordination. Fleet size=7.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 13 4' 0	235	34	0 5 11 13 (12) ^{del 55} 0	34
0 4" 12 17 14 15' 0	235	30	0 4 12 (12) _{col 55} 17 14 (15) ^{del 50} 0	28
0 15" 18 19 0	220	38	0 15 (15) _{col 50} 18 19 M ^{del 15} 0	38
0 16 7 6' 0	235	18	0 1 M _{col 15} 21 20 23 R ^{del 20} 0	42
0 8 9 6" 3 0	215	20	0 22 R _{col 20} 24 25 10 (3) ^{del 15} 0	38
0 2 1 0	120	6	0 16 7 (3) ^{del 25} 0	18
0 20 21 23 22' 0	235	42	0 8 9 6 (3) _{col 40} 3 0	20
0 22" 24 25 10 0	<u>185</u>	<u>38</u>	0 1 2 0	<u>6</u>
	1680	226		224

Table E.38: Improvement by coordination. Fleet size=6.

(vi) All other demands=50

Route	Load	Distance
0 5 11 13 4 0	200	34
0 12 17 14 15 0	200	28
0 18 19 0	200	36
0 21 20 23 24 0	200	44
0 22 25 10 8 0	200	34
0 9 16 0	200	24
0 6 7 3 2 0	200	12
0 1 0	<u>50</u>	<u>2</u>
	1450	214

Table E.39: No coordination. Fleet size=6.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 13 4 1' 0	235	34	0 4 5 11 13 (12) ^{del 35} 0	34
0 1" 12 17 14 15 0	215	28	0 12 (12) _{col 35} 17 18 14 15	34
			U ^{del 20} 0	
0 18 19 0	200	36	0 3 1 U _{col 20} 19 0	34
0 21 20 23 24 22' 0	235	44	0 21 20 23 24 R ^{del 35} 0	44
0 22" 25 10 8 7 6' 0	235	36	0 22 R _{col 35} 25 10 8 3 0	34
0 9 16 6" 0	230	24	0 9 16 0	24
0 2 3 0	<u>100</u>	<u>6</u>	0 7 6 0	<u>10</u>
	1450	208		214

Table E.40: Improvement by coordination. Fleet size=6.

Demand(13)=Demand(23)=150

(i) All other demands=100

Route	Load	Distance
0 5 11 0	200	30
0 13 0	150	24
0 4 12 0	200	18
0 1 17 0	200	26
0 14 15 0	200	20
0 18 19 0	200	36
0 20 21 0	200	32
0 23 0	150	32
0 22 0	100	14
0 24 25 0	200	26
0 10 8 0	200	34
0 16 9 0	200	24
0 6 7 0	200	10
0 3 2 0	200	6
	2600	332

Table E.41: No coordination. Fleet size=8.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4' 0	235	30	0 5 11 (4) ^{del 35} 0	30
0 13 4" 0	215	24	0 4 (4) _{col 35} 13 0	24
0 12 17 0	200	26	0 12 17 0	26
0 14 15 0	200	20	0 15 14 0	20
0 18 19 0	200	36	0 18 19 0	36
0 20 21 0	200	32	0 20 21 (22) ^{del 35} 0	34
0 22' 23 0	235	32	0 22 (22) _{col 35} 23 0	32
0 22" 25 24 0	215	26	0 24 25 0	26
0 10 8 0	200	34	0 10 8 D ^{del 35} 0	34
0 16 9 0	200	24	0 16 9 D ^{del 35} 0	24
0 7 6 0	200	10	0 7 6 D _{col 70} 3 0	10
0 3 2 0	200	6	0 2 1 0	6
0 1 0	100	2		302
	2600	302		

Table E.42: Improvement by coordination. Fleet size=8.

(ii) All other demands=90

Route	Load	Distance
0 5 11 0	180	30
0 13 0	150	24
0 4 12 0	180	18
0 1 17 0	180	26
0 14 15 0	180	20
0 18 19 0	180	36
0 20 21 0	180	32
0 23 0	150	32
0 22 0	90	14
0 24 25 0	180	26
0 10 8 0	180	34
0 16 9 0	180	24
0 6 7 0	180	10
0 3 2 0	<u>180</u>	<u>6</u>
	2370	332

Table E.43: No coordination. Fleet size=8.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 1' 0	235	30	0 5 11 E ^{del 55} 0	30
0 13 4' 0	235	24	0 4 E _{col 55} 13 (12) ^{del 50} 0	24
0 4" 12 17 14' 0	235	30	0 12 (12) _{col 50} 17 14 (15) ^{del 15} 0	28
0 14" 18 15 0	220	34	0 18 19 (15) ^{del 55} 0	38
0 19 20 21' 0	235	38	0 15 Z _{col 70} 20 21 T ^{del 35} 0	34
0 21" 23 22' 0	235	38	0 22 T _{col} 23 0	34
0 22" 25 24 0	220	26	0 24 25 0	26
0 10 8 3' 0	235	34	0 10 8 D ^{del 55} 0	34
0 9 16 6' 0	235	24	0 16 9 D _{col 55} 3 0	24
0 7 6" 0	125	10	0 7 6 0	10
0 1" 2 3" 0	<u>160</u>	<u>8</u>	0 1 2 0	<u>6</u>
	2370	296		288

Table E.44: Improvement by coordination. Fleet size=8.

(iii) All other demands=80

Route	Load	Distance
0 5 11 0	160	30
0 13 4 0	230	24
0 12 17 0	160	26
0 15 14 0	160	20
0 18 19 0	160	36
0 20 21 0	160	32
0 22 23 0	230	32
0 24 25 0	160	26
0 10 8 0	160	32
0 16 9 0	160	24
0 6 7 0	160	12
0 3 2 0	160	6
0 1 0	<u>80</u>	<u>2</u>
	2140	302

Table E.45: No coordination. Fleet size=8.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4' 0	235	30	0 5 11 G ^{del 75} 0	30
0 4" 13 12 0	235	24	0 13 12 G _{col 75} 4 0	26
0 15' 17 18 0	235	34	0 17 14 (15) ^{del 75} 0	28
0 15" 14 19 1' 0	235	32	0 15 (15) _{col 75} 18 19 M ^{del 70} 0	38
0 20 21 22' 0	235	34	0 1 M _{col 70} 20 21 S ^{del 65} 0	32
0 22" 23 24 0	235	34	0 23 24 S ^{del 5} 0	34
0 25 10 8' 0	235	34	0 22 S _{col 70} 25 10 (8) ^{del 65} 0	36
0 8" 9 16 7' 0	235	24	0 8 (8) _{col 65} 9 16 D ^{del 60} 0	24
0 7" 6 3 1" 0	180	12	0 3 D _{col 60} 6 7 0	10
0 2 0	<u>80</u>	<u>4</u>	0 2 0	<u>4</u>
	2140	262		262

Table E.46: Improvement by coordination. Fleet size=7 (splitting) and 6 (other forms).

(iv) All other demands=70

Route	Load	Distance
0 5 11 4 0	210	30
0 13 12 0	220	24
0 17 18 19 0	210	38
0 14 15 1 0	210	20
0 20 21 22 0	210	34
0 23 24 0	220	34
0 25 10 3 0	210	32
0 8 9 16 0	210	24
0 7 6 2 0	<u>210</u>	<u>12</u>
	1910	248

Table E.47: No coordination. Fleet size=6.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4 0	210	30	0 $J^{\text{del } 25}$ 5 11 4 0	30
0 13 12 0	220	24	0 13 12 0	24
0 14' 17 18 19 0	235	40	0 $U^{\text{del } 25}$ 15 14 17 0	28
0 14" 15 20 21' 0	235	38	0 18 19 $U_{\text{col } 25}$ $J_{\text{col } 25}$ 1 0	36
0 21" 23 22' 0	235	38	0 20 21 22 0	34
0 22" 24 25 10 0	215	38	0 23 24 0	34
0 8 9 16 0	210	24	0 25 10 8 0	34
0 7 6 3 0	210	10	0 9 16 7 0	24
0 1 2 0	<u>140</u>	<u>6</u>	0 6 3 2 0	<u>10</u>
	1910	248		254

Table E.48: Improvement by coordination. Fleet size=6 (splitting) and no improvement (other forms).

(v) All other demands=60

Route	Load	Distance
0 5 11 4 0	180	30
0 13 12 0	210	24
0 17 18 19 0	180	38
0 14 15 1 0	180	20
0 20 21 0	120	32
0 23 24 0	210	34
0 22 25 10 0	180	32
0 3 8 9 0	180	20
0 16 7 6 0	180	18
0 2 0	<u>60</u>	<u>4</u>
	1680	252

Table E.49: No coordination. Fleet size=6.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 13' 0	235	30	0 4 5 11 $G^{del 55}$ 0	30
0 4 13" 12 17 1' 0	235	24	0 13 $G_{col 55}$ 12 17 $K^{del 20}$ 0	38
0 15' 14 18 19 1" 0	235	38	0 15 14 $K_{col 20}$ 18 19 0	38
0 1'" 15" 20 21 22 0	225	36	0 20 21 $R^{del 55}$ 22 0	34
0 23 24 25' 0	235	36	0 23 24 $R_{col 55}$ 25 0	38
0 25" 10 8 3 0	215	34	0 10 9 8 $D^{del 55}$ 0	36
0 9 16 7 0	180	24	0 3 $D_{col 55}$ 6 7 16 0	20
0 6 2 0	<u>120</u>	<u>10</u>	0 2 0	<u>4</u>
	1680	232		238

Table E.50: Improvement by coordination. Fleet size=6.

(vi) All other demands=50

Route	Load	Distance
0 4 5 11 12 0	200	32
0 13 1 0	200	24
0 14 17 18 19 0	200	40
0 15 20 21 22 0	200	36
0 23 0	150	32
0 24 25 10 0	150	26
0 7 16 9 8 0	200	24
0 6 3 2 0	<u>150</u>	<u>10</u>
	1450	224

Table E.51: No coordination. Fleet size=6.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 4 5 11 12 1' 0	235	32	0 4 5 11 $F^{\text{del } 35}$ 12 0	34
0 1" 13 3 0	215	26	0 13 $F_{\text{col } 35}$ 17 14 $(15)^{\text{del } 20}$ 0	36
0 15' 14 17 18 19 0	215	40	0 15 $(15)_{\text{col } 20}$ 18 19 20 21 0	46
0 15" 20 21 25 22 0	235	40	0 22 23 $Q^{\text{del } 35}$ 0	32
0 23 24 0	200	34	0 24 $Q_{\text{col } 35}$ 25 10 9 8 $D^{\text{del } 20}$ 0	46
0 16 9 7 6 0	200	24	0 16 7 6 $D_{\text{col } 20}$ 3 2 0	20
0 10 8 2 0	<u>150</u>	<u>34</u>	0 1 0	<u>2</u>
	1450	230		216

Table E.52: Improvement by coordination. Fleet size=6.

Depot at D2




Figure E.4: Coordination site locations used in improvement by coordination for systems with depot at D2. The sites are generated within the improvement process and are not pre-selected.

The investigation includes systems of several nodes with demand 150 and all other nodes with equal demand:

- All demands=150
- Demand(13)=Demand(16)=Demand(18)=Demand(23)=150
- Demand(16)=Demand(19)=150
- Demand(13)=Demand(23)=150.

All demands=150

Route	Load	Distance
0 1 0	150	16
0 2 0	150	14
0 3 0	150	10
0 4 0	150	20
0 5 0	150	38
0 6 0	150	6
0 7 0	150	4
0 8 0	150	8
0 9 0	150	6
0 10 0	150	28
0 11 0	150	34
0 12 0	150	26
0 13 0	150	34
0 14 0	150	30
0 15 0	150	30
0 16 0	150	4
0 17 0	150	36
0 18 0	150	40
0 19 0	150	40
0 20 0	150	46
0 21 0	150	38
0 22 0	150	28
0 23 0	150	46
0 24 0	150	38
0 25 0	<u>150</u>	<u>34</u>
	3750	654

Table E.53: No coordination. Fleet size=16.

Splitting			Other forms	
Route	Load	Distance	Route	Distance
0 5 4' 0	235	38	0 5 H ^{del 85} 0	38
0 11 4" 0	215	34	0 11 H _{col 65} 13 0	38
0 12' 13 0	235	34	0 12 H ^{del 45} (3) ^{del 40} 0	28
0 12" 17 0	215	36	0 4 H _{col 65} 17 0	38
0 14' 18 0	235	42	0 18 M ^{del 85} 0	44
0 14" 19 0	215	42	0 19 M _{col 85} 20 0	50
0 15' 20 0	235	48	0 14 N ^{del 85} 0	32
0 15" 21 0	215	40	0 15 N _{col 85} 21 Q ^{del 20} 0	40
0 22' 23 0	235	46	0 23 Q ^{del 65} 0	46
0 22" 24 0	215	38	0 22 Q _{col} 24 (7) ^{del 20} 0	40
0 25 0	150	34	0 25 A ^{del 85} 0	34
0 9' 10 0	235	28	0 10 A _{col} 8 0	28
0 9" 8 0	215	10	0 1 (3) ^{del 85} 0	16
0 3' 1 0	235	16	0 2 (3) _{col 125} 3 V ^{del 60} 0	14
0 3" 2 0	215	14	0 6 (7) _{col 80} 7 0	6
0 7' 6 0	235	6	0 9 0	6
0 16 7" 0	<u>215</u>	<u>8</u>	0 16 0	<u>4</u>
	3750	514		502

Table E.54: Improvement by coordination. Fleet size= 12.

Demand(13)=Demand(16)=Demand(18)=Demand(23)=150

(i) All other demands=100

Route	Load	Distance
0 5 11 0	200	40
0 13 0	150	34
0 12 17 0	200	36
0 14 15 0	200	32
0 1 19 0	200	42
0 18 0	150	40
0 3 4 0	200	20
0 23 0	150	46
0 20 21 0	200	46
0 24 25 0	200	40
0 2 22 0	200	28
0 6 7 0	200	6
0 10 8 0	200	28
0 9 0	100	6
0 16 0	<u>150</u>	<u>4</u>
	2700	448

Table E.55: No coordination. Fleet size=11.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 3' 0	235	40	0 5 11 H ^{del 35} 0	40
0 13 4' 0	235	34	0 13 H _{col 35} 4 (3) ^{del 20} 0	34
0 4" 12 17 3" 0	235	38	0 17 12 (3) ^{del 35} 0	36
0 3''' 14 15 0	235	32	0 1 19 (3) ^{del 35} 0	42
0 1 19 3''' 0	210	42	0 14 15 (3) ^{del 35} 0	32
0 18 2' 0	235	44	0 18 (3) ^{del 85} 0	40
0 20 21 22' 0	235	46	0 20 21 p ^{del 35} 0	46
0 23 22" 0	215	46	0 23 p ^{del 85} 0	46
0 24 25 2" 0	215	40	0 24 25 Y _{col 120} 22 (3) _{col 45} 2	46
			(3) _{col 115} 3 0	
0 6 7 0	200	6	0 6 7 0	6
0 10 8 0	200	28	0 10 8 0	28
0 9 0	100	6	0 9 0	6
0 16 0	<u>150</u>	<u>4</u>	0 16 0	<u>4</u>
	2700	406		406

Table E.56: Improvement by coordination. Fleet size=10.

(ii) All other demands=90

Route	Load	Distance
0 5 11 0	180	40
0 13 0	150	34
0 4 12 0	180	28
0 17 19 0	180	48
0 18 0	150	40
0 20 21 0	180	46
0 22 24 0	180	38
0 23 0	150	46
0 25 2 0	180	34
0 14 15 0	180	32
0 1 3 0	180	16
0 6 7 0	180	6
0 9 0	90	6
0 16 0	150	4
0 10 8 0	<u>180</u>	<u>28</u>
	2490	446

Table E.57: No coordination. Fleet size=11.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4' 0	235	40	0 5 11 H ^{del 55} 0	40
0 13 4" 3' 0	235	34	0 4 H _{col 55} J ^{del 50} 13 0	34
0 3" 12 17 6' 0	235	36	0 12 J _{col 50} 17 K ^{del 15} 3 0	38
0 1' 15 14 0	235	32	0 18 K ^{del 20} (1) ^{del 65} 0	42
0 6" 18 0	225	40	0 19 K _{col 35} 14 15 0	44
0 7 19 1" 0	215	42	0 20 21 R ^{del 55} 0	48
0 20 21 2' 0	235	46	0 23 R _{col 55} 24 (1) _{col 65} 1 0	50
0 23 22' 0	235	46	0 22 25 0	34
0 24 25 22" 2" 0	220	40	0 10 8 0	28
0 10 8 0	180	28	0 16 0	4
0 9 0	90	6	0 9 7 0	10
0 16 0	<u>150</u>	<u>4</u>	0 2 6 0	<u>14</u>
	2490	394		386

Table E.58: Improvement by coordination. Fleet size=10 (splitting) and 9 (other forms).

(iii) All other demands=80

Route	Load	Distance
0 5 11 0	160	40
0 13 12 0	230	34
0 1 4 0	160	20
0 17 18 0	230	42
0 14 15 0	160	32
0 3 19 0	160	40
0 20 21 0	160	46
0 23 24 0	230	48
0 22 25 0	160	34
0 2 6 0	160	14
0 8 10 0	160	28
0 16 9 0	230	10
0 7 0	<u>80</u>	<u>4</u>
	2280	392

Table E.59: No coordination. Fleet size=10.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4' 0	235	40	0 5 11 H ^{del 75} 0	40
0 4" 13 12 0	235	34	0 4 H _{col 75} 13 12 0	34
0 17 18 14' 0	235	44	0 17 18 K ^{del 5} 0	42
0 15 14" 19 0	235	44	0 15 14 K _{col 5} 19 0	44
0 20 21 2' 0	235	46	0 20 21 U ^{del 75} 0	46
0 2" 23 24 0	235	48	0 23 24 U _{col 75} 1 0	50
0 22 25 6' 0	235	34	0 22 25 (3) ^{del 75} 0	34
0 1 3 6" 7' 0	235	16	0 6 2 (3) _{col 75} 3 0	10
0 10 8 7" 0	170	30	0 7 0	4
0 16 9 0	<u>230</u>	<u>10</u>	0 8 10 0	28
	2280	346	0 16 9 0	<u>10</u>
				342

Table E.60: Improvement by coordination. Fleet size=8.

(iv) All other demands=70

Route	Load	Distance
0 5 11 4 0	210	40
0 13 12 0	220	34
0 17 18 0	220	42
0 14 15 19 0	210	48
0 23 24 0	220	48
0 2 20 21 0	210	46
0 22 25 1 0	210	36
0 9 10 8 0	210	28
0 3 6 7 0	210	10
0 16 0	<u>150</u>	<u>4</u>
	2070	336

Table E.61: No coordination. Fleet size=8.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4 3' 0	235	40	0 5 11 4 (3) ^{del 25} 0	40
0 13 12 3" 0	235	34	0 13 12 (3) ^{del 15} 0	34
0 17 18 3''' 0	235	42	0 17 18 (3) ^{del 15} 0	42
0 14 15 19 3'''' 0	225	48	0 14 15 19 (3) _{col 45} 3 0	48
0 23 24 0	220	48	0 23 24 (3) ^{del 15} 0	48
0 2 21 20 0	210	46	0 2 21 20 (3) ^{del 25} 0	46
0 22 25 1 0	210	36	0 22 25 1 (3) _{col 50} 6 0	36
0 9 10 8 0	210	28	0 9 10 8 0	28
0 6 7 0	140	6	0 7 16 0	<u>8</u>
0 16 0	<u>150</u>	<u>4</u>		330
	2070	332		

Table E.62: Improvement by coordination. Fleet size=8.

(v) All other demands=60

Route	Load	Distance
0 5 11 4 0	180	40
0 13 12 0	210	34
0 17 18 0	210	42
0 14 15 19 0	180	48
0 23 24 0	210	48
0 2 20 21 0	180	46
0 22 25 1 0	180	36
0 9 10 8 0	180	28
0 3 6 7 0	180	10
0 16 0	<u>150</u>	<u>4</u>
	1860	336

Table E.63: No coordination. Fleet size=8.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4 3' 0	235	40	0 5 11 J ^{del 55} 12 0	42
0 13 12 3" 0	215	34	0 4 13 J _{col 55} 17 K ^{del 20} 0	44
0 17 18 6' 0	235	42	0 18 19 K _{col 20} 14 0	48
0 14 15 19 6" 0	215	48	0 3 1 15 0	30
0 23 24 7' 0	235	48	0 20 21 Q ^{del 55} 2 0	46
0 2 21 20 7" 0	215	46	0 23 24 Q ^{del 25} 0	48
0 22 25 1 0	180	36	0 22 25 Q _{col 80} 6 7 0	36
0 9 10 8 0	180	28	0 9 10 8 0	28
0 16 0	<u>150</u>	<u>4</u>	0 16 0	<u>4</u>
	1860	326		326

Table E.64: Improvement by coordination. Fleet size=8.

(vi) All other demands=50

Route	Load	Distance
0 5 11 4 3 0	200	40
0 12 13 0	200	34
0 1 15 14 17 0	200	40
0 18 19 0	200	46
0 20 21 2 6 0	200	46
0 23 24 0	200	48
0 22 25 10 8 0	200	38
0 16 9 0	200	10
0 7 0	<u>50</u>	<u>4</u>
	1650	306

Table E.65: No coordination. Fleet size=7.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4 3 7' 0	235	40	0 5 11 4 J ^{del 35} 3 0	42
0 12 13 7" 0	215	34	0 13 J _{col 35} 12 17 K ^{del 20} 0	34
0 1 15 14 17 6' 0	235	40	0 18 19 K _{col 20} 14 0	46
0 18 19 6" 0	215	46	0 1 15 0	30
0 20 21 22 2 0	200	48	0 20 21 S ^{del 35} 22 2 0	48
0 23 24 0	200	48	0 23 24 S _{col 35} 6 7 0	48
0 9 8 10 25 0	200	40	0 25 10 8 9 0	40
0 16 0	<u>150</u>	<u>4</u>	0 16 0	<u>4</u>
	1650	300		292

Table E.66: Improvement by coordination. Fleet size=7.

Demand(16)=Demand(19)=150

(i) All other demands=100

Route	Load	Distance
0 5 11 0	200	40
0 13 4 0	200	34
0 12 17 0	200	36
0 1 18 0	200	42
0 14 15 0	200	32
0 19 0	150	40
0 20 21 0	200	46
0 23 24 0	200	48
0 22 25 0	200	34
0 10 8 0	200	28
0 16 0	150	4
0 2 3 0	200	14
0 6 7 0	200	6
0 9 0	<u>100</u>	<u>6</u>
	2600	410

Table E.67: No coordination. Fleet size=10.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 3' 0	235	40	0 5 11 (3) ^{del 35} 0	40
0 4 13 3" 0	235	34	0 4 13 (3) ^{del 35} 0	34
0 3''' 12 17 7' 0	235	36	0 12 17 (3) _{col 65} 3 0	36
0 6' 1 18 0	235	42	0 14 15 (3) ^{del 35} 0	32
0 6" 14 15 0	235	32	0 1 18 (3) ^{del 35} 0	42
0 7''' 6''' 19 0	235	40	0 19 (3) _{col 15} 6 0	40
0 2' 20 21 0	235	46	0 20 21 T ^{del 35} 0	46
0 2" 23 24 0	235	48	0 22 25 T ^{del 35} 0	34
0 2''' 22 25 7''' 0	235	34	0 23 24 T _{col 70} 2 (3) _{col 60} 0	48
0 10 8 7''' 0	235	30	0 9 7 0	10
0 16 0	150	4	0 10 8 0	28
0 9 0	<u>100</u>	<u>6</u>	0 16 0	<u>4</u>
	2600	392		394

Table E.68: Improvement by coordination. Fleet size=10.

(ii) All other demands=90

Route	Load	Distance
0 5 11 0	180	40
0 13 4 0	180	34
0 12 17 0	180	36
0 14 15 0	180	32
0 19 0	150	40
0 1 18 0	180	42
0 20 21 0	180	46
0 23 24 0	180	48
0 22 25 0	180	34
0 10 8 0	180	28
0 2 3 0	180	14
0 6 7 0	180	6
0 9 0	90	6
0 16 0	<u>150</u>	<u>4</u>
	2370	410

Table E.69: No coordination. Fleet size=10.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4' 0	235	40	0 5 11 H ^{del 55} 0	40
0 4" 13 12 3' 0	235	34	0 4 13 H _{col 55} 12 E ^{del 20} 0	36
0 3" 14 15 0	235	32	0 14 15 E ^{del 55} 0	32
0 1' 17 18 0	235	42	0 17 18 E _{col 35} 1 0	44
0 1" 19 6' 0	235	42	0 19 E _{col 40} 2 3 0	44
0 20 21 2' 6" 0	235	46	0 20 21 Q ^{del 55} 0	46
0 23 24 2" 0	235	48	0 23 24 Q _{col 55} 22 0	50
0 22 25 2'" 3'" 0	215	34	0 25 10 0	38
0 10 8 0	180	28	0 8 9 0	10
0 9 7 0	180	10	0 16 0	4
0 16 0	<u>150</u>	<u>4</u>	0 6 7 0	<u>6</u>
	2370	360		350

Table E.70: Improvement by coordination. Fleet size=9 (splitting) and 8 (other forms).

(iii) All other demands=80

Route	Load	Distance
0 5 11 0	160	40
0 13 4 0	160	34
0 12 17 0	160	36
0 14 15 0	160	32
0 18 19 0	230	46
0 20 21 0	160	46
0 23 24 0	160	48
0 22 25 0	160	34
0 10 8 0	160	28
0 1 2 0	160	18
0 3 6 0	160	10
0 16 9 0	230	10
0 7 0	<u>80</u>	<u>4</u>
	2140	386

Table E.71: No coordination. Fleet size=9.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4' 0	235	40	0 5 11 $G^{\text{del } 75}$ 0	40
0 4" 13 12 3' 0	235	34	0 12 17 $G^{\text{del } 75}$ 0	36
0 3" 17 1 6' 0	235	38	0 4 13 $G_{\text{col}} 15 14 D^{\text{del } 65}$ 0	48
0 18 19 0	230	46	0 18 19 $D^{\text{del } 5}$ 0	46
0 20 21 2' 0	235	46	0 20 21 $Q^{\text{del } 75}$ 0	46
0 23 24 2" 6" 7' 0	235	48	0 23 24 $Q^{\text{del } 75}$ 0	48
0 22 25 7" 0	185	34	0 22 25 $Q_{\text{col } 150} 2 D_{\text{col } 70} 1 3 0$	34
0 14 15 0	160	32	0 10 8 0	28
0 10 8 0	160	28	0 16 9 0	10
0 16 9 0	<u>230</u>	<u>10</u>	0 7 6 0	<u>6</u>
	2140	356		342

Table E.72: Improvement by coordination. Fleet size=8.

(iv) All other demands=70

Route	Load	Distance
0 5 11 13 0	210	44
0 18 19 0	220	46
0 4 12 17 0	210	38
0 1 14 15 0	210	32
0 23 24 25 0	210	50
0 20 21 22 0	210	48
0 10 8 7 0	210	30
0 6 3 2 0	210	14
0 16 9 0	<u>220</u>	<u>10</u>
	1910	312

Table E.73: No coordination. Fleet size=7.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 13 3' 0	235	44	0 5 11 13 (3) ^{del 25} 0	44
0 4 12 17 3" 0	235	46	0 4 12 17 (3) ^{del 25} 0	38
0 18 19 2' 0	235	50	0 18 19 U ^{del 25} 0	46
0 1 14 15 2" 3'" 0	235	34	0 1 14 15 U ^{del 25} 0	34
0 23 24 25 2' 0	235	50	0 3 (3) _{col 50} 23 24 25 0	50
0 20 21 22 2" 0	235	48	0 2 U _{col 50} 22 21 20 0	50
0 9 10 8 0	210	28	0 9 8 10 0	28
0 7 6 0	140	6	0 7 6 0	6
0 16 0	<u>150</u>	<u>4</u>	0 16 0	<u>4</u>
	1910	310		300

Table E.74: Improvement by coordination. Fleet size=7.

(v) All other demands=60

Route	Load	Distance
0 5 11 13 0	180	44
0 18 19 0	210	46
0 4 12 17 0	180	38
0 1 14 15 0	180	32
0 23 24 25 0	180	50
0 20 21 22 0	180	48
0 10 8 7 0	180	30
0 6 3 2 0	180	14
0 16 9 0	<u>210</u>	<u>10</u>
	1680	312

Table E.75: No coordination. Fleet size=7.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 4' 5 11 13 0	235	44	0 5 11 13 $G^{del 55}$ 0	44
0 4" 12 17 3 6' 0	235	38	0 17 12 $G_{col 55}$ 4 3 $C^{del 50}$ 0	38
0 1 14 15 6" 7' 0	235	32	0 1 14 15 $C_{col 5}$ 7 0	32
0 18 19 7" 0	225	46	0 18 19 $C_{col 35}$ 6 0	46
0 2' 22 21 20 0	235	48	0 22 $S^{del 55}$ 21 20 0	48
0 2" 23 24 25 0	185	50	0 23 24 $S_{col 55}$ 25 2 0	50
0 8 10 9 0	180	28	0 8 10 9 0	28
0 16 0	<u>150</u>	<u>4</u>	0 16 0	<u>4</u>
	1680	290		290

Table E.76: Improvement by coordination. Fleet size=7.

(vi) All other demands=50

Route	Load	Distance
0 5 11 13 4 0	200	44
0 12 17 18 1 0	200	44
0 19 14 0	200	42
0 15 20 21 2 0	200	48
0 22 25 24 23 0	200	50
0 10 8 3 6 0	200	34
0 16 9 0	200	10
0 7 0	<u>50</u>	<u>4</u>
	1450	276

Table E.77: No coordination. Fleet size=6.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 13 4 3' 0	235	44	0 5 11 13 4 $V^{\text{del } 35}$ 0	44
0 12 17 18 1 3" 6' 0	235	44	0 1 $V_{\text{col } 35}$ 12 17 18 $L^{\text{del } 70}$ 0	44
0 19 14 6" 0	230	42	0 19 $L_{\text{col } 70}$ 14 15 0	44
0 15 20 21 2 7' 0	235	48	0 3 2 21 20 $P^{\text{del } 35}$ 0	46
0 22 23 24 25 7" 0	215	50	0 22 23 24 25 $P_{\text{col } 35}$ 6 0	50
0 16 0	150	4	0 9 10 8 0	28
0 9 10 8 0	<u>150</u>	<u>28</u>	0 16 7 0	<u>6</u>
	1450	260		262

Table E.78: Improvement by coordination. Fleet size=6.

Demand(13)=Demand(23)=150

(i) All other demands=100

Route	Load	Distance
0 5 11 0	200	40
0 13 0	150	34
0 4 12 0	200	28
0 18 17 0	200	42
0 1 19 0	200	40
0 14 15 0	200	32
0 20 21 0	200	46
0 23 0	150	46
0 24 25 0	200	40
0 2 22 0	200	28
0 10 8 0	200	28
0 3 6 0	200	10
0 16 9 0	200	10
0 7 0	<u>100</u>	<u>4</u>
	2600	428

Table E.79: No coordination. Fleet size=9.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4' 0	235	40	0 5 11 H ^{del 35} 0	40
0 4" 13 6' 0	235	34	0 13 H _{col 35} 4 0	34
0 1' 12 17 0	235	38	0 E ^{del 35} 12 17 0	36
0 1" 14 15 0	235	32	0 E ^{del 35} 14 15 0	32
0 1''' 18 19 0	230	46	0 18 19 E _{col 70} 1 0	48
0 22' 21 20 0	235	48	0 P ^{del 35} 20 21 0	46
0 22" 23 6" 0	235	46	0 P ^{del 55} 23 0	46
0 6''' 24 25 0	235	40	0 22 P _{col 90} 24 25 0	42
0 6''' 3 2 0	225	14	0 10 8 B ^{del 35} 0	28
0 10 8 0	200	28	0 2 3 B ^{del 35} 0	14
0 16 9 0	200	10	0 7 6 B _{col 70} 9 0	14
0 7 0	<u>100</u>	<u>4</u>	0 16 0	<u>4</u>
	2600	380		384

Table E.80: Improvement by coordination. Fleet size=8.

(ii) All other demands=90

Route	Load	Distance
0 5 11 0	180	40
0 13 0	150	34
0 4 12 0	180	28
0 18 17 0	180	42
0 1 19 0	180	40
0 14 15 0	180	32
0 20 21 0	180	46
0 23 0	150	46
0 24 25 0	180	40
0 2 22 0	180	28
0 10 8 0	180	28
0 3 6 0	180	10
0 16 9 0	180	10
0 7 0	<u>90</u>	<u>4</u>
	2370	428

Table E.81: No coordination. Fleet size=9.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 3' 0	235	40	0 5 11 H ^{del 55} 0	40
0 13 4' 0	235	34	0 13 H _{col 55} 12 V ^{del 50} 0	36
0 4" 12 17 1' 0	235	38	0 17 18 K ^{del 55} 0	42
0 18 19 1" 3" 0	235	46	0 3 V _{col 50} 19 6 0	40
0 14 15 3'" 2' 0	235	34	0 15 14 K _{col 55} 4 0	36
0 22' 21 20 0	235	48	0 20 21 R ^{del 50} 0	48
0 22" 23 2" 0	235	46	0 23 R _{col 50} 24 S ^{del 45} 0	50
0 2'" 24 25 0	185	40	0 2 22 S _{col 45} 25 0	34
0 10 8 0	180	28	0 10 8 0	28
0 16 9 0	180	10	0 16 9 0	10
0 7 6 0	<u>180</u>	<u>10</u>	0 7 0	<u>4</u>
	2370	374		368

Table E.82: Improvement by coordination. Fleet size=9.

(iii) All other demands=80

Route	Load	Distance
0 5 11 0	160	40
0 13 4 0	230	34
0 12 17 0	160	36
0 14 15 0	160	32
0 18 19 0	160	46
0 20 21 0	160	46
0 23 24 0	230	48
0 22 25 0	160	32
0 10 8 0	160	28
0 16 9 0	160	10
0 1 2 0	160	18
0 3 6 0	160	10
0 7 0	<u>80</u>	<u>4</u>
	2140	386

Table E.83: No coordination. Fleet size=8.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4' 0	235	40	0 5 11 H ^{del 75} 0	40
0 4" 13 12 0	235	34	0 13 12 H _{col 75} 4 0	36
0 17 14' 15 0	235	40	0 17 18 L ^{del 75} 0	42
0 14" 18 19 1' 0	235	48	0 19 L _{col 75} 14 15 (1) ^{del 70} 0	44
0 20 21 22' 0	235	48	0 1 (1) _{col 70} 21 20 P ^{del 65} 0	46
0 22" 23 24 0	235	48	0 24 25 P _{col 65} 22 0	40
0 1" 2 25 0	170	38	0 23 2 0	46
0 10 8 7' 0	235	30	0 9 8 (8) ^{del 75} 0	10
0 16 9 0	160	10	0 10 (8) _{col 75} 6 7 0	32
0 7" 6 3 0	<u>165</u>	<u>10</u>	0 16 0	<u>4</u>
	2140	346		340

Table E.84: Improvement by coordination. Fleet size=8.

(iv) All other demands=70

Route	Load	Distance
0 5 11 4 0	210	40
0 13 12 0	220	34
0 17 18 19 0	210	48
0 1 14 15 0	210	32
0 23 24 0	220	48
0 20 21 2 0	210	46
0 3 22 25 0	210	34
0 16 6 7 0	210	10
0 9 8 10 0	<u>210</u>	<u>30</u>
	1910	322

Table E.85: No coordination. Fleet size=8.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 3' 5 11 4 0	235	40	0 (3) ^{del 25} 5 11 4 0	40
0 3" 13 12 0	235	34	0 (3) ^{del 15} 13 12 0	34
0 3'" 17 18 19 0	235	48	0 (3) ^{del 25} 17 18 19 0	48
0 3''' 14 15 1 2' 0	220	32	0 1 14 15 (3) _{col 45} 3 0	32
0 2" 23 24 0	235	48	0 23 24 p ^{del 15} 0	48
0 6 2'" 21 20 0	235	46	0 20 21 2 p ^{del 25} 0	46
0 2''' 22 25 10 0	235	38	0 6 (3) _{col 20} p _{col 40} 22 25 10 0	38
0 9 8 7 0	210	12	0 9 8 7 0	12
0 16 0	<u>70</u>	<u>4</u>	0 16 0	<u>4</u>
	1910	302		302

Table E.86: Improvement by coordination. Fleet size=7.

(v) All other demands=60

Route	Load	Distance
0 5 11 4 0	180	40
0 13 12 0	210	34
0 17 18 19 0	180	48
0 1 14 15 0	180	32
0 23 24 0	210	48
0 20 21 2 0	180	46
0 9 8 10 0	180	30
0 3 22 25 0	180	34
0 16 7 6 0	<u>180</u>	<u>10</u>
	1680	322

Table E.87: No coordination. Fleet size=8.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 1' 4 5 11 0	235	42	0 E ^{del 55} 4 5 11 0	40
0 1" 13 12 0	215	34	0 13 12 E _{col 35} 1 0	36
0 14' 19 18 17 0	235	50	0 17 18 19 (14) ^{del 55} 0	50
0 14" 15 20 21 3' 0	235	50	0 14 (14) _{col 55} 15 20 21 P ^{del 50} 0	50
0 23 24 3" 0	220	48	0 23 24 P _{col 50} 3 0	48
0 22 25 2 0	180	34	0 2 22 25 0	34
0 16 7 6 0	180	10	0 16 7 6 0	10
0 9 8 10 0	<u>180</u>	<u>30</u>	0 9 8 10 0	<u>30</u>
	1680	298		298

Table E.88: Improvement by coordination. Fleet size=7.

(vi) All other demands=50

Route	Load	Distance
0 5 11 4 3 0	200	40
0 13 12 0	200	34
0 17 18 15 14 0	200	46
0 19 20 21 2 0	200	50
0 23 24 0	200	48
0 22 25 10 9 0	200	38
0 8 1 6 7 0	200	20
0 16 0	<u>50</u>	<u>4</u>
	1450	280

Table E.89: No coordination. Fleet size=7.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 5 11 4 3 1' 0	235	40	0 4 5 11 J ^{del 35} 12 0	42
0 13 12 1" 6' 0	235	34	0 13 J _{col 35} 1 3 (6) ^{del 20} 0	34
0 17 18 19 14 6" 7' 0	235	50	0 17 18 19 14 (6) ^{del 35} 0	50
0 15 20 21 22 7" 0	235	50	0 15 20 21 22 (6) _{col 55} 6 7 0	50
0 23 24 2' 0	225	48	0 23 24 U ^{del 35} 0	48
0 7'" 2" 25 10 9 16 0	235	42	0 9 8 10 25 U _{col 35} 2 0	40
0 8 0	<u>50</u>	<u>10</u>	0 16 0	<u>4</u>
	1450	274		268

Table E.90: Improvement by coordination. Fleet size=6.

Dense grid

maxVehCap=235, maxVehDist=50




Figure E.5: Dense grid layout.

The investigation includes systems of several nodes with demand 150 and all other nodes with equal demand:

- All demands=150
- Demand(9)=Demand(22)=150.




Figure E.6: Coordination site locations on the dense grid. The sites are generated within the improvement process and are not pre-selected.

All demands=150

Route	Load	Distance
0 1 0	150	10
0 2 0	150	16
0 3 0	150	20
0 4 0	150	20
0 5 0	150	14
0 6 0	150	12
0 7 0	150	20
0 8 0	150	16
0 9 0	150	26
0 10 0	150	32
0 11 0	150	36
0 12 0	150	32
0 13 0	150	14
0 14 0	150	16
0 15 0	150	12
0 16 0	150	16
0 17 0	150	30
0 18 0	150	16
0 19 0	150	12
0 20 0	150	12
0 21 0	150	10
0 22 0	150	16
0 23 0	150	24
0 24 0	150	20
0 25 0	<u>150</u>	<u>10</u>
	3750	462

Table E.91: No coordination. Fleet size=10.

Splitting			Other forms	
Route	Load	Distance	Route	Distance
0 12 11' 0	235	38	0 10 T ^{del 85} 0	32
0 11" 10 0	215	38	0 11 T _{col 65} 14 0	42
0 2 3' 0	235	20	0 12 T ^{del 85} 0	32
0 3" 4 0	215	22	0 15 T _{col 65} 17 0	32
0 5 6' 0	235	16	0 16 T _{col 40} Q ^{del 40} 0	20
0 6" 7 0	215	22	0 18 Z ^{del 85} 0	16
0 8 13' 0	235	20	0 19 Z _{col 85} 20 0	14
0 9 13" 0	215	30	0 22 Y ^{del} 0	16
0 14 15' 0	235	18	0 21 Y _{col} 23 0	24
0 15" 16 0	215	16	0 24 G ^{del} 0	20
0 17 18' 0	235	30	0 25 G _{col} 1 0	14
0 18" 19 0	215	16	0 2 J ^{del} 0	18
0 20 21' 0	235	12	0 3 J _{col} 4 0	22
0 22 21" 0	215	16	0 5 L ^{del} 0	14
0 23 25' 0	235	24	0 7 L _{col} 6 Q ^{del 20} 0	22
0 24 25" 0	215	20	0 8 Q ^{del 20} 0	16
0 1 0	<u>150</u>	<u>10</u>	0 9 Q _{col 80} 13 0	<u>34</u>
	3750	368		388

Table E.92: Improvement by coordination. Fleet size= 8.

Demand(9)=Demand(22)=150

(i) All other demands=100

Route	Load	Distance
0 9 0	150	26
0 22 0	150	16
0 11 12 0	200	38
0 10 8 0	200	32
0 13 14 0	200	20
0 15 16 0	200	16
0 17 18 0	200	30
0 19 20 0	200	14
0 21 23 0	200	24
0 24 25 0	200	20
0 1 2 0	200	16
0 3 4 0	200	22
0 5 7 0	200	22
0 6 0	<u>100</u>	<u>12</u>
	2600	308

Table E.93: No coordination. Fleet size=7.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 9 13' 0	235	28	0 9 $N^{\text{del } 85}$ 0	28
0 11 12 13" 6' 0	235	42	0 11 $N_{\text{col } 85}$ 10 $T^{\text{del } 20}$ 0	40
0 10 14 15' 0	235	38	0 12 14 $T^{\text{del } 35}$ 0	32
0 17 16 15" 0	235	30	0 17 16 $T^{\text{del } 35}$ 0	30
0 15" 18 19 0	230	12	0 13 $T_{\text{col } 90}$ 15 $X^{\text{del } 25}$ 0	24
0 8 7 6" 0	235	26	0 18 19 $X^{\text{del } 35}$ 0	16
0 6" 5 1 0	235	18	0 20 21 $X^{\text{del } 35}$ 0	14
0 6" 4 3 0	210	24	0 22 $X_{\text{col } 95}$ 23 $E^{\text{del } 80}$ 0	24
0 21' 22 0	235	16	0 24 $E_{\text{col } 80}$ 25 1 0	24
0 20 23 21" 0	215	26	0 $K^{\text{del } 35}$ 6 8 0	18
0 25 24 0	200	20	0 $K^{\text{del } 35}$ 7 5 0	22
0 2 0	<u>100</u>	<u>16</u>	0 2 3 $K_{\text{col } 70}$ 4 0	<u>30</u>
	2600	296		302

Table E.94: Improvement by coordination. Fleet size= 7.

(ii) All other demands=90

Route	Load	Distance
0 9 0	150	26
0 22 0	150	16
0 11 12 0	180	38
0 10 8 0	180	32
0 13 14 0	180	20
0 15 16 0	180	16
0 17 18 0	180	30
0 19 20 0	180	14
0 21 23 0	180	24
0 24 25 0	180	20
0 1 2 0	180	16
0 3 4 0	180	22
0 5 7 0	180	22
0 6 0	<u>90</u>	<u>12</u>
	2370	308

Table E.95: No coordination. Fleet size=7.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 9 8' 0	235	26	0 9 N ^{del 85} 0	28
0 13 10 8" 6' 0	235	14	0 10 N _{col 85} 11 12 S ^{del 50} 0	40
0 11 12 15' 0	235	40	0 15 S _{col 35} 14 13 0	22
0 14 16 15" 0	215	20	0 16 0	16
0 17 18 19' 0	235	30	0 17 18 Z ^{del 55} 0	30
0 19" 20 21 0	215	16	0 19 20 Z _{col 40} 21 0	18
0 22 25' 0	235	20	0 22 Z _{col 15} 25 0	20
0 23 24 25" 0	185	26	0 23 24 G ^{del} 0	26
0 3 4 0	180	22	0 D ^{del} 5 1 0	18
0 6" 7 5 0	190	24	0 3 4 G _{col} 2 0	32
0 1 2 0	<u>180</u>	<u>16</u>	0 8 7 D _{col} 6 0	<u>30</u>
	2370	254		280

Table E.96: Improvement by coordination. Fleet size= 6.

(iii) All other demands=80

Route	Load	Distance
0 1 2 0	160	16
0 3 4 0	160	22
0 5 6 0	160	16
0 7 8 0	160	24
0 9 10 0	230	32
0 11 12 0	160	38
0 13 14 0	160	20
0 15 16 0	160	16
0 17 18 0	160	30
0 19 20 0	160	14
0 21 22 0	230	16
0 23 24 0	160	26
0 25 0	<u>80</u>	<u>10</u>
	2140	280

Table E.97: No coordination. Fleet size=6.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 1' 2 3 0	235	20	0 1 2 $H^{\text{del } 75}$ 0	16
0 1" 4 5 6' 0	235	24	0 3 4 $H^{\text{del } 75}$ 0	22
0 6" 7 8 13' 0	235	30	0 5 6 $H_{\text{col } 150}$ 7 8 0 $P^{\text{del } 65}$ 0	32
0 9 10 13" 0	235	32	0 13 $P_{\text{col } 65}$ 11 12 $S^{\text{del } 60}$ 0	42
0 13'" 11 12 14' 0	235	38	0 9 10 0	32
0 14" 16 15 0	175	20	0 14 $S_{\text{col } 60}$ 15 16 $V^{\text{del } 55}$ 0	22
0 17 18 19' 0	235	30	0 17 18 $V_{\text{col } 55}$ 19 0	34
0 19" 20 21 0	165	16	0 20 21 0	14
0 22 23 0	230	24	0 22 23 0	24
0 24 25 0	<u>160</u>	<u>20</u>	0 24 25 0	<u>20</u>
	2140	254		258

Table E.98: Improvement by coordination. Fleet size=6 (splitting) and 7 (other forms).

(iv) All other demands=70

Route	Load	Distance
0 9 8 0	220	26
0 22 23 0	220	24
0 2 3 4 0	210	22
0 5 6 7 0	210	24
0 10 11 12 0	210	40
0 13 14 15 0	210	22
0 16 17 18 0	210	30
0 19 20 21 0	210	16
0 24 25 1 0	<u>210</u>	<u>24</u>
	1910	228

Table E.99: No coordination. Fleet size=6.

No advantage in splitting.

Other forms	
Route	Dist
0 2 3 4 (1) ^{del 25} 0	22
0 6 7 5 (1) ^{del 25} 0	26
0 23 24 25 (1) _{col 50} 1 0	30
0 21 22 0	16
0 8 9 0	26
0 U ^{del 25} 10 11 12 0	40
0 U ^{del 25} 17 18 19 0	30
0 13 14 U _{col 50} 15 16 0	28
0 20 0	<u>12</u>
	230

Table E.100: Improvement by coordination other than splitting.
Fleet size= 6.

(v) All other demands=60

Route	Load	Distance
0 13 9 0	210	30
0 10 11 12 0	180	40
0 8 7 6 0	180	26
0 14 15 16 0	180	20
0 17 18 19 0	180	30
0 2 3 4 0	180	22
0 22 20 0	210	18
0 5 1 25 0	180	20
0 21 23 24 0	<u>180</u>	<u>26</u>
	1680	232

Table E.101: No coordination. Fleet size=5.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 8 9 13' 0	235	30	0 8 9 N ^{del 25} 0	28
0 13" 10 11 12 14' 0	235	40	0 13 N _{col 25} 10 11 12 0	40
0 14" 15 16 18 19' 0	235	26	0 14 15 16 V ^{del} 0	22
0 17 19" 20 21 0	225	34	0 17 18 19 V _{col} 20 0	36
0 22 23 0	210	24	0 22 21 0	16
0 24 25 0	120	20	0 23 24 25 F ^{del} 0	26
0 1' 2 3 4 0	235	22	0 4 3 2 F _{col} 1 0	26
0 1" 5 7 6 0	<u>185</u>	<u>26</u>	0 5 7 6 0	<u>24</u>
	1680	222		218

Table E.102: Improvement by coordination. Fleet size= 5.

(vi) All other demands=50

Route	Load	Distance
0 1 2 3 4 0	200	22
0 6 5 7 8 0	200	28
0 13 9 0	200	28
0 10 11 12 14 0	200	40
0 15 16 0	100	16
0 17 18 19 20 0	200	32
0 21 22 0	200	16
0 23 24 25 0	<u>150</u>	<u>26</u>
	1450	208

Table E.103: No coordination. Fleet size=5.

Splitting			Other forms	
Route	Load	Dist	Route	Dist
0 25 2 3 4 1' 0	235	24	0 8 $M^{del 35}$ 7 5 6 0	28
0 1" 5 7 8 6 0	215	30	0 13 $M_{col 35}$ 9 10 $T^{del 70}$ 0	36
0 13 9 0	200	26	0 11 12 14 $T_{col 20}$ 16 15 0	42
0 10 11 12 14 15' 0	235	42	0 17 18 19 20 0	32
0 15" 16 17 18 19 0	215	30	0 22 21 0	16
0 20 22 0	200	18	0 23 24 25 0	26
0 24 23 21 0	<u>150</u>	<u>26</u>	0 1 2 3 4 0	<u>22</u>
	1450	196		202

Table E.104: Improvement by coordination. Fleet size= 5.

APPENDIX F

DATA FOR EXAMPLES IN CHAPTER 6

Standard layout




Figure F.1: Example structure.

maxVehDist=50
maxVehCap=200

Node locations and demands

The depot is located at [12,12].

Unless stated otherwise, the following node locations and demands apply.

Node	Location		Demand
1	19	4	43
2	17	15	166
3	16	8	12
4	18	14	28
5	3	8	114
6	18	18	119
7	8	17	118
8	2	4	3
9	5	3	32
10	8	18	17
11	9	8	18
12	16	10	63
13	4	13	58
14	13	7	200
15	14	8	50

Table F.1: Node locations and demands.

BIBLIOGRAPHY

- [AF93] S. Anily and A. Federgruen. Two-echelon distribution systems with vehicle routing costs and central inventories. *Operations Research*, Vol. 41, No. 1, January-February 1993, 37-47
- [AG91] K. Altinkemer and B. Gavish. Parallel savings based heuristics for the delivery problem. *Operations Research*, Vol. 39, No. 3, May - June 1991, 456-469
- [AMS89] Y. Agarwal, K. Mathur and H. Salkin. A set partitioning-based exact algorithm for the vehicle routing problem. manuscript, 1989
- [B81] L. Bodin and B. Golden. Classification in vehicle routing and scheduling. *Networks* 11, 97-108, 1981
- [B83] L. Bodin, B. Golden, A. Assad and M. Ball. Routing and scheduling of vehicles and crews. *Computers and Operations Research* 10:2, 62-212, 1983
- [B92] D. Bertsimas. A vehicle routing problem with stochastic demand. *Operations Research* 40, 574-585, 1992
- [B93] M.D. Byrne. *Aspects of the Vehicle Routing Problem with Pickup and Delivery*. PhD thesis, Massey University, 1993
- [B95] D. Bertsimas, P. Chervi and M. Peterson. Computational approaches to stochastic vehicle routing problems. *Transportation Science* Vol. 29, No. 4, November 1995, 342-352

- [B97] J.R. Birge. Stochastic programming computation and applications. *INFORMS Journal on Computing* Vol. 9, No. 2, Spring 1997, 111-133
- [B00] J.M. Belenguer, M.C. Martinez and E. Mota. A lower bound for the split delivery vehicle routing problem. *Operations Research* Vol. 48, No. 5, September-October 2000, 801-810
- [B01] I.R. Beale. *Subset selection routing: Modelling and practicalities*. PhD thesis, Massey University, 2001
- [BCP95] D. Bertsimas, P. Chervi and M. Peterson. Computational approaches to stochastic vehicle routing problems. *Transportation Science*, Vol. 29, No. 4, 1995, 342-352
- [BDF83] W. Bell, L. Dalberto, M. Fisher, A. Greenfield, R. Jaikumar, P. Kedia, R. Mack and P. Prutzman. Improving the distribution of industrial gases with an on-line computerized routing and scheduling optimizer. *Interfaces* 4-23, 1983
- [BFS93] H.M. Baker, L.S. Franz, and J.R. Sweigart, Coordinated transportation systems: An alternative approach to traditional independent systems. *European Journal of Operational Research* 66, 1993, 341-352
- [BH93] M.E. Berge and C.A. Hopperstad. Demand driven dispatch: A method for dynamic aircraft capacity assignment, models and algorithms. *Operations Research* Vol. 41, 1993, 153-168
- [BHJD98] J.F. Bard, L. Huang, P. Jaillet, and M. Dror. A decomposition approach to the inventory routing problem with satellite facilities. *Transportation Science* Vol. 32, No. 2, May 1998, 189-203
- [BHM96] A. Bachem, W. Hochstättler and M. Malich, The simulated trading heuristic for solving vehicle routing problems. *Discrete Applied Mathematics* 65 (1996) 47-72.

- [BJO90] D. Bertsimas, P. Jaillet and A.R. Odoni. A priori optimization. *Operations Research* Vol. 38, 1019-1033
- [BL93] M.O. Ball and F.L. Lin. A reliability model applied to emergency service vehicle location. *Operations Research* Vol. 41, No. 1, January-February 1993, 18-23
- [BS96] D.J. Bertsimas and D. Simchi-Levi. A new generation of vehicle routing research: robust algorithms, addressing uncertainty. *Operations Research* Vol. 44, No. 2, March-April 1996, 286-304
- [BV91] D.J. Bertsimas and G. van Ryzin. A stochastic and dynamic vehicle routing problem in the Euclidean plane. *Operations Research* Vol. 39, 1991, 601-615.
- [BV93] D.J. Bertsimas and G. van Ryzin. Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles. *Operations Research* Vol. 41, 1993, 60-76.
- [C76] N. Christofides. *Worst-case analysis of a new heuristic for the travelling salesman problem*. Management Sciences Research Report No. 388, Carnegie-Mellon University, Feb. 1976
- [C85] N. Christofides. Vehicle Routing, in *The Traveling Salesman Problem*. E. Lawler, J.K. Lenstra, A.H.G. Rinnooy Kan and D.B. Shmoys (eds.), John Wiley and Sons, New York, 1985, 431-448
- [C90] J.F. Campbell. Locating transportation terminals to serve an expanding demand. *Transportation Research* 24 (B) 3, 1990, 173-192
- [C93] M. Carey and A. Srinivasan. Externalities, average and marginal costs, and tolls on congested networks with time-varying flows. *Operations Research* Vol. 41, No. 1, January-February 1993, 217-231

- [C98] R.K. Cheung. A two-stage stochastic network model and solution methods for the dynamic empty container allocation problem. *Transportation Science*, Vol. 32, No. 2, May 1998, 142-162
- [CBW89] T. Chien, A. Balakrishnan and R. Wong. An integrated inventory allocation and vehicle routing problem. *Transportation Science*, 1989, 67-76.
- [CCKS97] A. Campbell, L. Clarke, A. Kleywegt, and M. Savelsbergh. The Inventory Routing Problem. manuscript, 1997
- [CJR81] F. Cullen, J. Jarvis and H. Ratcliff. Set partitioning based heuristics for interactive routing. *Networks*, 1981, 125-143.
- [CKR98] M. Charikar, S. Khuller and B. Raghavachari. Algorithms for capacitated vehicle routing. *1998 ACM Symposium on Theory of Computing* 349-358
- [CMT79] N. Christofides, A. Mingozzi and P. Toth. The vehicle routing problem, in N. Christofides, A. Mingozzi, P. Toth and C. Sandi, eds., *Combinatorial Optimization* (Wiley, New York, 1979).
- [CW64] G. Clarke and J. Wright. Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research* 12:4, 1964, 568-581
- [D78] C.F. Daganzo. An approximate analytic model of many-to-many demand responsive transportation systems. *Transportation Research* 12, 1978, 325-333
- [D95] R.B. Dial. Autonomous dial-a-ride transit introductory overview, *Transportation Research* 3C, 1995, 261-275
- [D99] J. Desrosiers, F. Soumis, and N. Bélanger. Fleet assignment and routing with schedule synchronization constraints. *European Journal of Operational Research* Vol. 119, No. 1, 1999, 75-90

- [D00] M. Dror, M. Haouari, and J. Chaouachi. Generalized spanning trees. *European Journal of Operations Research* Vol. 120, 2000, 583-592
- [DHM90] M.S. Daskin, W.J. Hopp and B. Medina. Forecast horizons and dynamic facility location planning. *Annals of Operations Research* Vol. 40, 1990, 125-151
- [DLS90] M. Desrochers, J.K. Lenstra and M.W.P. Savelsbergh. A classification scheme for vehicle routing and scheduling problems. *European Journal of Operations Research* Vol. 46, 1990, 322-332
- [DLT89] M. Dror, G. Laporte and P. Trudeau. Vehicle routing with stochastic demands: Properties and solution frameworks. *Transportation Science* Vol. 23, 1989, 166-176
- [DLT94] M. Dror, G. Laporte and P. Trudeau. Vehicle routing with split deliveries. *Discrete Applied Mathematics* Vol. 50, 1994, 239-254
- [DP93] M. Dror and W. Powell, Stochastic and Dynamic Models in Transportation. *Operations Research* Vol. 41, No. 1, January-February 1993, 11-14
- [DR59] G.B. Dantzig and J.H. Ramser. The truck dispatching problem. *Management Science* Vol. 6, No. 1, October 1959, 80-91
- [DS01] P.M. Dearing and R. Segars. An equivalence result for planar location problems with rectilinear distance and barriers. manuscript, 2001
- [DT86] M. Dror and P. Trudeau. 1986. Stochastic vehicle routing with modified savings algorithm. *European Journal of Operational Research*. Vol. 23, 228-235
- [DT89] M. Dror and P. Trudeau. Savings by split delivery routing. *Transportation Science* 23(2): 141-145 1989

- [DT90] M. Dror and P. Trudeau. Split delivery routing. *Naval Research Logistics* Vol. 37, 1990, 383-402
- [DW95] Z. Drezner and G.O. Wesolowsky. Location on a one-way rectilinear grid. *Journal of the Operational Research Society* 46, 1995, 735-746
- [E97] F. Ercal and H.C. Lee, Time efficient maze routing algorithms on reconfigurable mesh architectures. *Journal of Parallel and Distributed Computing*, 44(2):133-140, August 1997
- [F92] P.W. Frizzell and J.W. Giffin. The bounded split delivery vehicle routing problem with grid network distances. *Asia-Pacific Journal of Operational Research*, Vol. 9, 1992, 101-116
- [F93] P.W. Frizzell. *Aspects of the Vehicle Routing and Scheduling Problem*. PhD thesis, Massey University, 1993
- [F95] P.W. Frizzell and J.W. Giffin. The split delivery vehicle scheduling problem with time windows and grid network distances. *Computers and Operations Research*, Vol. 22, No. 6, 1995, 655-667
- [F99] M. Fleischer and S.H. Jacobson. Information theory and the finite-time behavior of the simulated annealing algorithm: Experimental results. *INFORMS Journal on Computing*. Vol. 11, No. 1, Winter 1999, 35-43
- [FBSTW93] T.L. Friesz, D. Bernstein, T. E. Smith, R.L Tobin and B.W. Wie. A variational inequality formulation of the dynamic network user equilibrium problem. *Operations Research* Vol. 41, No. 1, January-February 1993, 179-191
- [FJ81] M. Fisher and R. Jaikumar. A generalized assignment heuristic for vehicle routing. *Networks*, 11:109-124, 1981

- [FZ84] A. Federgruen and P. Zipkin. A combined vehicle routing and inventory allocation problem. *Operations Research* Vol. 32, 1984, 1019-1037.
- [G99] K.R. Gue. The effects of trailer scheduling on the layout of freight terminals. *Transportation Science*, Vol. 33, No. 4, November 1999, 419-428
- [GA88] B.L. Golden and A. Assad (eds.). *Vehicle Routing: Methods and Studies*, North-Holland, Amsterdam, 1988
- [G89] F. Glover. Tabu Search Part 1. *ORSA Journal on Computing* Vol. 1, No. 3, Summer 1989, 190-206
- [G89a] D. E. Goldberg. *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley Publishing Company, Inc., Reading, Massachusetts, 1989
- [G96] M. Gendreau, F. Guertin, J.Y. Potvin, and E. Taillard. Tabu search for real-time vehicle routing and dispatching. Technical Report CRT-96-47, Centre de recherche sur les transports, Universite de Montreal, Montreal, Canada, 1996
- [GB80] B. Golden, L. Bodin, T. Doyle and W.J. Stewart. Approximate traveling salesman algorithms. *Operations Research* 28, No. 3, Part II, May-June 1980, 694-711
- [GJ79] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the theory of NP-completeness*. Freeman, San Francisco, 1979
- [GLS95] M. Gendreau, G. Laporte and R. Seguin. An exact algorithm for the vehicle routing problem with stochastic demands and customers. *Transportation Science*, Vol. 29, No. 2, May 1995, 143-156

- [GM74] B.E. Gillett and L.R. Miller. A heuristic for the vehicle-dispatch problem. *Operations Research* 22, 1974, 340-349
- [GMN77] B. Golden, T. Magnanti, and H. Nyugen. Implementing vehicle routing algorithms. *Networks*, 7:113-148, 1977
- [GV99] N. Gans & G. van Ryzin. Dynamic vehicle dispatching: optimal heavy traffic performance and practical insights. *Operations Research*, Vol. 47, No. 5, September-October 1999, 675-692
- [H00] A. Hertz, G. Laporte, and M. Mittaz. A tabu search heuristic for the capacitated arc routing problem. *Operations Research*, Vol. 48, No. 1, January-February 2000, 129-135
- [H01] R.W. Hall. The study of transportation is paved with science. *OR/MS Today*, August 2000, 22-25
- [HM00] S.G. Henderson and A.J. Mason. BartSim: A tool for analysing and improving ambulance performance in Auckland, New Zealand. *Proceedings of the Operational Research Society of New Zealand 35th Annual Conference 2000*, 57-64
- [HM01] A. Hertz and M. Mittaz. A variable neighbourhood descent algorithm of the undirected capacitated arc routing problem. *Transportation Science*, Vol. 53, No. 4, November 2001, 425-434
- [HP97] R.W. Hall and J.G. Partyka. On the road to efficiency. *OR/MS Today*, June 1997, 38-41
- [HR85] M. Haimovich and A. Rinooy Kan. Bounds and heuristics for capacitated routing problems. *Mathematics of Operations Research*, 10:527-542, 1985
- [IGP00] S. Ichoua, M. Gendreau and J-Y. Potvin. Diversion issues in real-time vehicle dispatching. *Transportation Science*, Vol. 34 No. 4, November 2000, pp 426-438

- [KB95] G. Kontoravdis and J.F. Bard. A greedy random adaptive search procedure for the vehicle routing problem with time windows", ORSA Journal on Computing Vol. 7, No. 1, 1995, 10-23
- [KGV83] S. Kirkpatrick, S.D.J. Gelatt and M.P. Vecchi. Optimization by simulated annealing. *Science*, 220, No. 4598, 1983, 671-680
- [KNS97] A.J. Kleywegt, V.S. Nori, and M.W.P. Savelsbergh. Dynamic and stochastic minimization. manuscript, May 1997
- [LMR96] K. Lund, O.B.G. Madsen and J.M. Rygaard. *Vehicle routing problems with varying degrees of dynamism*. Technical Report IMM-REP-1996-1, 1996, Institut for Matematisk Modellering
- [LN87] G. Laporte and Y. Nobert. Exact algorithms for the vehicle routing problem. *Annals of Discrete Mathematics*, 31:147-184, 1987
- [L92] G. Laporte. Vehicle routing problem: An overview. *European Journal of Operational Research*, Vol. 59, 1992, 345-358
- [LH01] F.J. Lerch and D.E. Harter. Cognitive support for real-time dynamic decision making. *Information Systems Research*, Vol. 12, No. 1, March 2001, 63-82
- [LK73] S. Lin and B.W. Kernighan. An effective heuristic algorithm for the travelling salesman problem. *Operations Research*, Vol. 21, 1973, 498-516
- [LLM91] M. Labbé, G. Laporte and H. Mercure. Capacitated vehicle routing on trees. *Operations Research*, Vol. 39, No. 4, July-August 1991, 616-622
- [LMR96] K. Lund, O.B.G. Madsen and J.M. Rygaard. *Vehicle routing problems with varying degress of dynamism*, Technical Report IMM-REP-1996-1, 1996

- [M93] A.S. Minkoff. A Markov decision model and decomposition heuristic for dynamic vehicle dispatching. *Operations Research*, Vol. 41, No. 1 January-February 1993.
- [M97] P.A. Mullaseril, M. Dror and J. Leung. Split-delivery routing heuristic in livestock feed distribution. *Journal of the Operational Research Society*, Vol. 48, 1997, 107-116
- [NJK90] K.E. Nygard, P. Juell and N. Kadaba. Neural networks for selecting vehicle routing heuristics. *ORSA Journal on Computing*, Vol. 2, No.4, Fall 1990 353-364
- [NZE57] New Zealand Explosives Act 1957
- [NZRC] New Zealand Road Code
- [NZS88] NZ Standard 5433:1988 "Code of Practice for the Transport of Hazardous Substances on Land"
- [NZTAA95] New Zealand Transport Amendment Act 1995
- [O96] I.H. Osman and G. Laporte. Metaheuristics: A bibliography. *Annals of Operations Research*, Vol. 63, 1996, 513-623
- [O99] S.H. Owen. *Scenario planning approaches to facility location: Models and solution methods*. PhD dissertation, Northwestern University Evanston Illinois USA, 1999
- [P00] W.L. Perry. Knowledge and combat outcomes. *Military Operations Research*, Vol. 5, No. 1, 2000, 29-39
- [P88] W. Powell. A comparative review of alternative algorithms for the dynamic vehicle allocation problem. In *Vehicle routing: Methods and Studies* B. Golden and A. Assad (eds.) North-Holland, Amsterdam, 1998, 249-291.

- [P93] J-Y. Potvin. The travelling salesman problem: A neural network perspective. *ORSA Journal on Computing*, Vol. 5, No. 4, Fall 1993, 328-348.
- [Ps88] H.N. Psaraftis. Dynamic vehicle routing problems. In *Vehicle routing: Methods and studies* B. Golden and A. Assad (eds.) North-Holland, Amsterdam, 1988, 223-248
- [Ps95] H.N. Psaraftis. Dynamic vehicle routing: Status and prospects, *Annals of Operations Research* Vol. 61, 1995, 143-164
- [PSC00] W.B. Powell, W. Snow, and R.K. Cheung. Adaptive labeling algorithms for the dynamic assignment problem. *Transportation Science*, Vol. 34, No. 1, February 2000, 50-66
- [PSN88] W. Powell, Y. Sheffi, K.S. Nickerson, K. Butterbaugh, and S. Atherton, Maximising profits for North American Van Lines truckload division: A new framework for pricing operations, *Interfaces*, 18, 1988, 21-41
- [PT93] H.N. Psaraftis and J.N. Tsitsiklis. Dynamic shortest paths in acyclic networks with Markovian arc costs. *Operations Research*, Vol. 41, No.1, 1993, 91-101
- [R82] Ole M. Raft. A modular algorithm for an extended vehicle scheduling problem, *European Journal of Operational Research*, 11, 1982, 67-76
- [R94] Gerhard Reinelt. *The Traveling Salesman - Computational Solutions for TSP Applications*. Springer-Verlag. 1994

- [R97] C.M. Rivers. The en route re-distribution of payload. *Proceedings of the fourth conference of the Association of Asian-Pacific Operational Research Societies, within IFORS, 1997*
- [R98] C.M. Rivers. Coordination in vehicle routing. *Proceedings of the Operational Research Society of New Zealand, 33rd Annual Conference 1998, 314-323*
- [R99] C.M. Rivers. Dynamic coordination in vehicle routing. *Proceedings of the Operational Research Society of New Zealand, 34th Annual Conference 1999, 107-116*
- [RBL96] J. Renaud, F.F. Boctor and G. Laporte. An improved petal heuristic for the vehicle routing problem, *Journal of the Operational Research Society, Vol. 47, 1996, 329-336*
- [RM86] R.A. Russell and R.B. Morrel. Routing special-education school buses. *Interfaces 16:5 September-October 1986, 56-64*
- [RMJ94] A.C. Regan, H.S. Mahmassani and P. Jaillet. Improving efficiency of commercial vehicle operations using real-time information: Potential uses and assignment strategies. *Transportation Research Rec. 1493, 1994, 188-198*
- [RMJ95] A.C. Regan, H.S. Mahmassani and P. Jaillet. Dynamic decision making for commercial fleet operations using real-time information. *Transportation Research Rec. 1537, 1995, 91-97*
- [RT95] Y. Rochat and É. Taillard. Probabilistic diversification and intensification in local search for vehicle routing. *Journal of Heuristics, Vol. 1, 1995, 147-167*

- [RW96] R.E. Rosenthal and W.J. Walsh. Optimizing flight operations for an aircraft carrier in transit. *Operations Research* Vol. 44, No. 2, March-April 1996, 305-312
- [S01] N. Secomandi. A rollout policy for the vehicle routing problem with stochastic demands. *Operations Research*, Vol. 49, No. 5, September-October 2001, 796-802
- [S78] D.M. Stein. Scheduling dial-a-ride transportation systems. *Transportation Research*, 12, 1978, 232-249
- [S80] D.A. Schilling. Dynamic location modelling for public-sector facilities: A multicriteria approach. *Decision Sciences*, Vol. 11, 1980, 714-724
- [S99] K.A. Smith. Neural networks for combinatorial optimization: A review of more than a decade of research. *INFORMS Journal on Computing*, Vol. 11, Winter 1999, No. 1, 15-34
- [SD88] M.M. Solomon and J Desrosiers. Time window constrained routing and scheduling problems. *Transportation Science*, Vol. 22, No. 1, February 1988, 1-13
- [SG83] W. Stewart Jr and B. Golden. Stochastic vehicle routing: a comprehensive approach. *European Journal of Operational Research*, Vol. 14, 1983, 371-385.
- [SP95] Y. Shen and J.Y. Potvin. A computer assistant for vehicle dispatching with learning capabilities, *Annals of Operations Research*, 61, 1995, 189-211
- [SS98] M. Savelsbergh and M. Sol. DRIVE: Dynamic routing of independent vehicles. *Operations Research*, Vol. 46, No. 4, 1998, 474-490

- [ST98] G. Sierksma and G.A. Tijssen. Routing helicopters for crew exchanges on off-shore locations. *Annals of Operations Research*, 76, 1998, 261-286
- [T89] P. Trudeau, J.M. Rousseau, J.A. Ferland, and J. Choquette. An operations research approach for the planning and operating of an ambulance service. *INFOR* 27, 1989, 95-113
- [T97] É. Taillard, P. Badeau, M. Gendreau, F. Guertin and J-Y. Potvin. A tabu search heuristic for the vehicle routing problem with soft time windows. *Transportation Science*, Vol. 31, No. 2, May 1997, 170-186
- [TE98] S. Tsubakitani and J.R. Evans. An empirical study of a new metaheuristic for the travelling salesman problem. *European Journal of Operations Research* 104, 1998, 113-128
- [TLG96] É.D. Taillard, G. Laporte and M. Gendreau. Vehicle routeing with multiple use of vehicles. *Journal of the Operational Research Society* Vol. 47, 1996, 1065-1070
- [USHM90] The Hazardous Material Transportation Uniform Safety Act of 1990-HMTUSA
- [W86] S. Wiedenbeck. The feasibility of load-splitting in vehicle routing problems. Master's thesis, The University of Arizona, November 1986.
- [W98] O. Wendt. Cooperative simulated annealing. Internet: www.is-frankfurt.de/cosa/ 1998
- [WC77] N.H.M. Wilson and N.H. Colvin. Computer control of the Rochester dial-a-ride system. Report R77-31, Dept of Civil Engineering, MIT, Cambridge, MA, 1977

- [YHM90] A. Yamani, T.J. Hodgson and L.A. Martin-Vega. Single aircraft mid-air refueling using spherical distances. *Operations Research* Vol. 38, No. 5, September-October 1990, 792-800

