

---

# An Algorithmic Framework for School Bus Routing with Stop Selection and On-Time Reliability

---

Monique Sciortino

A thesis submitted for the degree of  
*Doctor of Philosophy in Mathematics*

School of Mathematics

Cardiff University

28th January 2026



# Summary

The school bus routing problem (SBRP) is an NP-hard combinatorial optimisation problem with significant practical relevance. This research is driven by persistent operational challenges reported by guardians and school transport operators, including excessively early pick-up times, long and inconsistent travel durations, and occasional vehicle capacity shortages. These inefficiencies impact student well-being and service equity, revealing limitations in current planning methods. There is a pressing need for routing solutions that are not only cost-efficient but also scalable and reliable under daily operational uncertainty.

This dissertation addresses the SBRP in a realistic context, involving problem instances with sizes upwards of 1800 potential bus stops and 750 students. A mixed integer programming model is first developed to solve small-scale instances optimally and establish performance benchmarks. To handle the computational challenges of larger instances, a heuristic algorithmic framework is proposed, incorporating a bus stop selection component, multiple constructive and improvement heuristics, and a destroy-and-repair mechanism for effective solution space exploration.

The most advanced version of the framework extends to stochastic settings by modelling arc travel time uncertainty using shifted lognormal distributions, which reflect the skewed nature of real-world travel times. A correlation structure, based on a classification of arc pairs, is also introduced to capture spatial dependencies. Additionally, a percentile-based reliability metric is employed to account for the goal of on-time student arrival.

Computational experiments using real-world data, supported by Monte Carlo simulations, demonstrate that accounting for both travel time variability and spatial correlation is critical for generating reliable solutions that better meet service expectations. In contrast, approaches assuming deterministic or independent travel times may underestimate delays or even be practically unimplementable. By aligning algorithmic development with operational realities, this research contributes decision-support tools for enhancing reliability in school bus routing as well as in broader vehicle routing settings.



# Contents

Summary	i
List of Tables	vii
List of Figures	xiii
List of Acronyms and Initials	xxi
Dedication	xxv
Acknowledgements	xxvii
<b>1 Introduction</b>	<b>1</b>
1.1 SBRP Decomposition . . . . .	2
1.2 SBRP Classifications . . . . .	3
1.3 Research Aims . . . . .	4
1.4 Thesis Contributions . . . . .	4
1.5 Thesis Outline . . . . .	6
1.6 Academic Publications . . . . .	7
<b>2 Literature Review</b>	<b>9</b>
2.1 The VRP and its Variants . . . . .	9
2.2 SBRP Characteristics . . . . .	11
2.2.1 Number of Schools and Allowance of Mixed Loads . . . . .	11
2.2.2 Service Environment . . . . .	12
2.2.3 Morning Versus Afternoon Problem . . . . .	13
2.2.4 Fleet Mix . . . . .	13
2.2.5 Objectives . . . . .	14
2.2.6 Constraints . . . . .	16
2.3 Solution Methods for the SBRP . . . . .	18
2.3.1 Bus Stop Selection . . . . .	18
2.3.2 Route Generation . . . . .	23
2.3.3 School Bell Time Adjustment and Route Scheduling . . . . .	28

2.4	Summary . . . . .	29
<b>3</b>	<b>The Single-School Bus Routing Problem</b>	<b>31</b>
3.1	Introduction . . . . .	31
3.2	Modelling the S-SBRP . . . . .	32
3.2.1	Definitions and Assumptions . . . . .	32
3.2.2	Constraints . . . . .	34
3.2.3	Objectives . . . . .	38
3.2.4	Mathematical Formulation of the S-SBRP . . . . .	40
3.3	Computational Complexity of the S-SBRP . . . . .	47
3.4	Computational Results . . . . .	48
3.4.1	Problem Instances . . . . .	48
3.4.2	Analysis of Results . . . . .	51
3.5	Summary . . . . .	55
<b>4</b>	<b>A Heuristic Algorithm for the Single-School Bus Routing Problem</b>	<b>57</b>
4.1	Introduction . . . . .	57
4.2	Constructive Heuristics . . . . .	58
4.2.1	Related Works . . . . .	58
4.2.2	Implemented Constructive Heuristics . . . . .	63
4.3	Improvement Heuristics . . . . .	69
4.3.1	Related Works . . . . .	69
4.3.2	Implemented Intra-Route Operators . . . . .	71
4.3.3	Implemented Inter-Route Operators . . . . .	73
4.3.4	Application of Neighbourhood Operators . . . . .	75
4.4	Perturbation Mechanisms . . . . .	76
4.5	The Algorithmic Framework . . . . .	79
4.6	Computational Results . . . . .	80
4.7	Summary . . . . .	91
<b>5</b>	<b>Modelling Stochastic Travel Times</b>	<b>93</b>
5.1	Introduction . . . . .	93
5.2	Modelling a VRP with Stochastic Travel Times . . . . .	95
5.3	Travel Time Distributions . . . . .	99
5.4	Travel Time Reliability Metrics . . . . .	102
5.5	Proposed Framework . . . . .	106
5.5.1	Distributional Assumptions . . . . .	107
5.5.2	Problem and Algorithm Modifications . . . . .	112
5.6	Computational Results . . . . .	114

5.6.1	Baseline Configuration . . . . .	115
5.6.2	Other Configurations . . . . .	121
5.7	Summary . . . . .	125
<b>6</b>	<b>Modelling Travel Time Correlations</b>	<b>127</b>
6.1	Introduction . . . . .	127
6.2	Related Works . . . . .	130
6.2.1	Existence of Travel Time Correlations . . . . .	130
6.2.2	Modelling and Analysing the Impact of Travel Time Correlations . . . . .	133
6.3	Correlation Structure . . . . .	139
6.3.1	Classification of Arc Pairs . . . . .	139
6.3.2	Correlation Estimates . . . . .	141
6.3.3	Constructing the Correlation Matrix and Ensuring Positive Definiteness . . . . .	144
6.4	Computational Results . . . . .	148
6.5	Monte Carlo Simulations of Correlated SLN Travel Times . . . . .	155
6.5.1	Monte Carlo Sampling Methodology . . . . .	155
6.5.2	Performance Evaluation of Routing Solutions . . . . .	159
6.6	Summary . . . . .	164
<b>7</b>	<b>Conclusions and Future Research</b>	<b>167</b>
7.1	The Problem Investigated . . . . .	167
7.2	Summary of Findings . . . . .	168
7.2.1	First Research Aim . . . . .	168
7.2.2	Second Research Aim . . . . .	169
7.2.3	Third Research Aim . . . . .	171
7.2.4	Fourth Research Aim . . . . .	172
7.2.5	Fifth Research Aim . . . . .	173
7.3	Future Research . . . . .	174
7.3.1	Cost Function . . . . .	174
7.3.2	Bus Stop Selection Subproblem . . . . .	174
7.3.3	Constructive and Improvement Heuristics . . . . .	175
7.3.4	Hybrid Approaches . . . . .	176
7.3.5	Softer Approach to Travel Time Reliability . . . . .	176
7.3.6	Travel Time Distributions . . . . .	176
7.3.7	Spatial Correlations . . . . .	177
7.3.8	Dynamic and Multischool Settings . . . . .	178
7.4	Summary of Available Resources . . . . .	179
7.5	Closing Remarks . . . . .	179

<b>Appendices</b>	<b>211</b>
<b>A Visualizations of MIP Solutions</b>	<b>213</b>
<b>B Heuristic Algorithm Deterministic Scenario Results</b>	<b>219</b>
B.1 Total Journey Time Boxplots . . . . .	219
B.2 Independent-Samples Kruskal-Wallis Tests . . . . .	226
B.3 Average Results for the Other Algorithm Variants . . . . .	227
B.4 Visualizations of Best SRH-4 Solutions . . . . .	233
<b>C Heuristic Algorithm Independent Scenario Results</b>	<b>245</b>
C.1 Results for $\kappa = 0.95$ . . . . .	245
C.2 Computational Times and Feasibility Rates . . . . .	247
C.3 Visualizations of Best Solutions . . . . .	247
<b>D Heuristic Algorithm Correlated Scenario Results</b>	<b>259</b>
D.1 Computational Times and Feasibility Rates . . . . .	259
D.2 Visualizations of Best Solutions . . . . .	260

# List of Tables

2.1	Examples of objectives considered in the literature. . . . .	16
2.2	Examples of constraints considered in the literature. . . . .	17
3.1	Summary of the twenty real-world problem instances, listed in increasing order of $ V_1 $ . $S$ represents the total number of students, calculated as $\sum_{w \in V_2} s(w)$ . Distances $m_e$ and $m_w$ are given in km. . . . .	49
3.2	Descriptive statistics on the travelling times (rounded to the nearest second), namely the mean $\pm$ standard deviation, lower quartile (LQ), median and upper quartile (UQ). These are calculated across all arcs in $E_1$ . . . . .	49
3.3	MIP model statistics. . . . .	52
3.4	MIP solutions. $k$ denotes the number of routes in the solution, Cap. is short for Capacity, JT stands for Journey Time (minutes), and TJT stands for Total Journey Time (minutes). The routes are presented in terms of the indices of the vertices in $V_1$ , and the subscripts represent the number of boarding students. . . . .	53
4.1	Example of a subset of visited bus stops and their pairwise travel times (in minutes). The first column gives the number of boarding students. . . . .	64
4.2	Example of PNNH Initial Routes. Cap. is short for Capacity, DT stands for Dwell Time (minutes), TT stands for Travel Time (minutes), and JT stands for Journey Time (minutes). The routes are presented in terms of the indices of the vertices in $V_1$ , and the subscripts represent the number of boarding students. . . . .	65
4.3	Example of SRH Initial Routes. Cap. is short for Capacity, DT stands for Dwell Time (minutes), TT stands for Travel Time (minutes), and JT stands for Journey Time (minutes). The routes are presented in terms of the indices of the vertices in $V_1$ , and the subscripts represent the number of boarding students. . . . .	66

4.4	Example of RIH Initial Routes. Cap. is short for Capacity, DT stands for Dwell Time (minutes), TT stands for Travel Time (minutes), and JT stands for Journey Time (minutes). The routes are presented in terms of the indices of the vertices in $V_1$ , and the subscripts represent the number of boarding students. . . . .	68
4.5	Results are averaged across 25 runs. For each constructive heuristic, the best total journey time (TJT) average (in minutes) across all perturbation mechanisms is presented, with the mechanism number(s) in superscript. Bold results represent the overall best averages. . . . .	81
4.6	Mann-Whitney test results. The values under the SRH-3, SRH-4, RIH-3, and RIH-4 columns represent the mean ranks, $U$ represents the test statistic, while bold $p$ -values are significant at the 0.05 level.	82
4.7	PM 4 results for SRH. Results are averaged across 25 runs and presented as mean $\pm$ standard deviation. Time gives the computational time (seconds), FR stands for Feasibility Rate (percentage of iterations yielding a solution with all route journey times being at most 45 minutes), and TJT stands for Total Journey Time (minutes). Bold TJT values represent the best averages across all SRH perturbation mechanisms. . . . .	83
4.8	Sensitivity analysis assessing the impact of each neighbourhood operator on the algorithm's performance. Inc. represents the relative increase in the TJT average across 25 runs, while bold $p$ -values are significant at the 0.05 level. . . . .	85
4.9	Sensitivity analysis replacing SD with RVND, and doubling the number of iterations. Inc. (Dec.) represents the relative increase (decrease) in the TJT average across 25 runs. The values under the SD and RVND columns represent the mean ranks, $U$ represents the Mann-Whitney test statistic, while bold $p$ -values are significant at the 0.05 level. . . . .	86
4.10	Results for SRH-4 with RVND, and SRH-4 with RVND and double iterations. Results are averaged across 25 runs and presented as mean $\pm$ standard deviation. Time gives the computational time (seconds), FR stands for Feasibility Rate (percentage of iterations yielding a solution with all route journey times being at most 45 minutes), and TJT stands for Total Journey Time (minutes). . . .	88

4.11	Best SRH-4 (with RVND and double iterations) solutions. Cap. is short for Capacity, JT stands for Journey Time (minutes), and TJT stands for Total Journey Time (minutes). The routes are presented in terms of the indices of the vertices in $V_1$ , and the subscripts represent the number of boarding students. . . . .	89
5.1	Different percentile journey times in seconds (varying parameters $\alpha$ , $\beta$ , and $\kappa$ ) for a route containing a single arc with mean travel time 1000s and one boarding student. . . . .	114
5.2	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.2, \beta = 0.5$ . TMT and TPT stand for Total Mean Time and Total 99 <sup>th</sup> Percentile Time (minutes), presented as mean $\pm$ standard deviation. Comparative results for the deterministic scenario are also presented, with the column IR indicating the number of runs for which the deterministic solution is infeasible with respect to the on-time arrival chance constraints. A hyphen indicates infeasibility. . . . .	116
5.3	Mann-Whitney test results comparing the TMTs and TPTs of the deterministic and independent scenarios. The values under the Det. and Ind. columns represent the mean ranks, and $U$ represents the test statistic. . . . .	117
5.4	Best solutions for the independent scenario with parameters $\alpha = 0.2, \beta = 0.5$ . Cap. is short for Capacity, MT stands for Mean Time (minutes), and PT stands for Percentile Time (minutes). The routes are presented in terms of the indices of the vertices in $V_1$ , and the subscripts represent the number of boarding students. . . . .	117
5.5	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.2, \beta = 0.25$ . . . . .	121
5.6	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.2, \beta = 0.75$ . . . . .	122
5.7	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.1, \beta = 0.25$ . . . . .	122
5.8	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.1, \beta = 0.5$ . . . . .	123
5.9	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.1, \beta = 0.75$ . . . . .	123
5.10	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.3, \beta = 0.25$ . . . . .	124

5.11	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.3, \beta = 0.5$ . . . . .	124
5.12	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.3, \beta = 0.75$ . . . . .	125
6.1	Summary of assigned correlation estimates. . . . .	144
6.2	Proportions of arc pairs in each category (%): (a) first-order adjacent, (b) second-order adjacent, (c) opposing, (d) first-order nearby, (e) second-order nearby, (f) overlapping, and (g) unrelated. PDA stands for the number of pairs of distinct arcs (in millions). . . . .	149
6.3	Results averaged across 25 runs for the correlated scenario with parameters $\alpha = 0.2, \beta = 0.5$ . TMT and TPT stand for Total Mean Time and Total 99 <sup>th</sup> Percentile Time (minutes), presented as mean $\pm$ standard deviation. Comparative results for the deterministic and independent scenarios are also presented, with the column IR indicating the number of runs for which the solution is infeasible with respect to the on-time arrival chance constraints. An asterisk signifies an instance for which $\kappa$ was reduced to 0.95 due to infeasibility. . . . .	150
6.4	Mann-Whitney test results comparing the TMTs and TPTs of the deterministic and correlated scenarios. The Det. and Corr. columns report the mean ranks, $U$ represents the test statistic, while bold $p$ -values are less than 0.05. . . . .	151
6.5	Mann-Whitney test results comparing the TMTs and TPTs of the independent and correlated scenarios. The Ind. and Corr. columns report the mean ranks, $U$ represents the test statistic, while bold $p$ -values are less than 0.05. . . . .	151
6.6	Best solutions for the correlated scenario with parameters $\alpha = 0.2, \beta = 0.5$ . Cap. is short for Capacity, MT stands for Mean Time (minutes), and PT stands for Percentile Time (minutes). The routes are presented in terms of the indices of the vertices in $V_1$ , and the subscripts represent the number of boarding students. . . . .	152
6.7	Monte Carlo results for $I = 1,000$ samples, averaged across 25 solutions for parameters $\alpha = 0.2, \beta = 0.5$ and presented as mean $\pm$ standard deviation. The column FS indicates the number of samples out of 1,000 for which the solution is feasible (all routes with journey times not exceeding 45 minutes). TJT stands for Total Journey Time (minutes), averaged across the feasible samples. . . . .	160

6.8	Monte Carlo results for $I = 1,000$ samples of the shifted gamma, generalized Pareto, and Burr distributions, averaged across 25 solutions and presented as mean $\pm$ standard deviation. The values indicate the number of samples out of 1,000 for which the solution is feasible (all routes with journey times not exceeding 45 minutes).	163
B.1	Kruskal-Wallis test results for comparing the algorithm variants.	226
B.2	PM 3 results for SRH.	227
B.3	PM 1 and PM 2 results for SRH.	228
B.4	PM 1 and PM 2 results for PNNH.	229
B.5	PM 3 and PM 4 results for PNNH.	230
B.6	PM 1 and PM 2 results for RIH.	231
B.7	PM 3 and PM 4 results for RIH.	232
C.1	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.2, \beta = 0.5, \kappa = 0.95$ .	245
C.2	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.2, \beta = 0.75, \kappa = 0.95$ .	245
C.3	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.1, \beta = 0.5, \kappa = 0.95$ .	246
C.4	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.1, \beta = 0.75, \kappa = 0.95$ .	246
C.5	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.3, \beta = 0.5, \kappa = 0.95$ .	246
C.6	Results averaged across 25 runs for the independent scenario with parameters $\alpha = 0.3, \beta = 0.75, \kappa = 0.95$ .	246
C.7	Computational times and feasibility rates of the independent scenario's runs with $\alpha = 0.2, \beta = 0.5$ .	247
D.1	Computational times and feasibility rates of the correlated scenario's runs with $\alpha = 0.2, \beta = 0.5$ .	259



# List of Figures

3.1	An S-SBRP solution with $k = 2$ , $ V_1  = 26$ , $ V'_1  = 14$ and $ V_2  = 30$ .	36
3.2	A bipartition with $k = 5$ and $b = 8$ . The node labels represent the route loads and bus capacities. An example of a minimum-weight maximum-cardinality matching is displayed in red. Note that edge weights have been excluded for clarity. . . . .	40
3.3	A subtour demonstrating a violation of an MTZ constraint. Without loss of generality, setting $l_{7i} = 8$ , $l_{4i} = 14$ , and $l_{2i} = 24$ results in a violation of the MTZ constraint for arc $(2, 7)$ . . . . .	43
3.4	Victoria (left) and Cardiff (right) problem instances. The lime dot represents the school, the red dots represent the potential bus stops, and the yellow dots represent the student addresses. . . . .	50
3.5	Mgarr MIP solution with 4 routes (red, lime, blue, yellow), 190 students, average walk 6.06 minutes, and average journey time 13.50 minutes. The lime dot represents the school, the red dots represent the visited bus stops, and the yellow dots represent the student addresses. . . . .	54
5.1	Different probability density functions (varying parameters $\alpha$ and $\beta$ ) for an arc with mean travel time 1000s. . . . .	109
6.1	Road network diagram illustrating the propagation of traffic conditions – congestion on roads $a$ and $b$ can impact roads $c$ , $d$ , and $e$ . . . . .	127
6.2	Classification of arc pairs for correlation structure. . . . .	140
6.3	Example of a valid correlation matrix yielded by the scaled spectral method. . . . .	147
A.1	Mgarr MIP solution with 4 routes (red, lime, blue, yellow), 190 students, average walk 6.06 minutes, and average journey time 13.50 minutes. . . . .	213

A.2	Mellieħa MIP solution with 4 routes (red, lime, blue, yellow), 171 students, average walk 5.04 minutes, and average journey time 14.07 minutes. . . . .	214
A.3	Porthcawl MIP solution with 1 route (red), 66 students, average walk 14.26 minutes, and average journey time 26.87 minutes. . . . .	214
A.4	Qrendi MIP solution with 5 routes (red, lime, blue, yellow, magenta), 255 students, average walk 6.31 minutes, and average journey time 16.03 minutes. . . . .	215
A.5	Suffolk MIP solution with 5 routes (red, lime, blue, yellow, magenta), 209 students, average walk 7.68 minutes, and average journey time 22.85 minutes. . . . .	215
A.6	Senglea MIP solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 266 students, average walk 6.19 minutes, and average journey time 12.88 minutes. . . . .	216
A.7	Cardiff MIP solution with 3 routes (red, lime, blue), 156 students, average walk 11.04 minutes, and average journey time 19.08 minutes. . . . .	216
A.8	Milton Keynes MIP solution with 5 routes (red, lime, blue, yellow, magenta), 274 students, average walk 11.40 minutes, and average journey time 13.36 minutes. . . . .	217
B.1	Total journey time values for the Mġarr instance. . . . .	219
B.2	Total journey time values for the Mellieħa instance. . . . .	220
B.3	Total journey time values for the Porthcawl instance. . . . .	220
B.4	Total journey time values for the Qrendi instance. . . . .	220
B.5	Total journey time values for the Suffolk instance. . . . .	221
B.6	Total journey time values for the Senglea instance. . . . .	221
B.7	Total journey time values for the Victoria instance. . . . .	221
B.8	Total journey time values for the Pembroke instance. . . . .	222
B.9	Total journey time values for the Canberra instance. . . . .	222
B.10	Total journey time values for the Ħandaq instance. . . . .	222
B.11	Total journey time values for the Valletta instance. . . . .	223
B.12	Total journey time values for the Birkirkara instance. . . . .	223
B.13	Total journey time values for the Ħamrun instance. . . . .	223
B.14	Total journey time values for the Cardiff instance. . . . .	224
B.15	Total journey time values for the Milton Keynes instance. . . . .	224
B.16	Total journey time values for the Bridgend instance. . . . .	224
B.17	Total journey time values for the Edinburgh-2 instance. . . . .	225
B.18	Total journey time values for the Edinburgh-1 instance. . . . .	225
B.19	Total journey time values for the Adelaide instance. . . . .	225

B.20	Total journey time values for the Brisbane instance. . . . .	226
B.21	Mġarr best SRH-4 solution with 4 routes (red, lime, blue, yellow), 190 students, average walk 6.05 minutes, and average journey time 13.53 minutes. . . . .	233
B.22	Mellieħa best SRH-4 solution with 4 routes (red, lime, blue, yellow), 171 students, average walk 5.03 minutes, and average journey time 14.07 minutes. . . . .	234
B.23	Porthcawl best SRH-4 solution with 1 route (red), 66 students, average walk 14.26 minutes, and average journey time 26.87 minutes.	234
B.24	Qrendi best SRH-4 solution with 5 routes (red, lime, blue, yellow, magenta), 255 students, average walk 6.60 minutes, and average journey time 15.15 minutes. . . . .	235
B.25	Suffolk best SRH-4 solution with 3 routes (red, lime, blue), 209 students, average walk 7.33 minutes, and average journey time 38.24 minutes. . . . .	235
B.26	Senglea best SRH-4 solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 266 students, average walk 6.72 minutes, and average journey time 11.89 minutes. . . . .	236
B.27	Victoria best SRH-4 solution with 4 routes (red, lime, blue, yellow), 171 students, average walk 7.42 minutes, and average journey time 23.15 minutes. . . . .	236
B.28	Pembroke best SRH-4 solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 335 students, average walk 6.86 minutes, and average journey time 14.65 minutes. . . . .	237
B.29	Canberra best SRH-4 solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 499 students, average walk 7.19 minutes, and average journey time 25.45 minutes. . . . .	237
B.30	Ħandaq best SRH-4 solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 285 students, average walk 6.97 minutes, and average journey time 15.94 minutes. . . . .	238
B.31	Valletta best SRH-4 solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 268 students, average walk 6.82 minutes, and average journey time 17.06 minutes. . . . .	238
B.32	Birkirkara best SRH-4 solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 306 students, average walk 7.17 minutes, and average journey time 15.65 minutes. . . . .	239
B.33	Ħamrun best SRH-4 solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 321 students, average walk 6.94 minutes, and average journey time 13.91 minutes. . . . .	239

B.34 Cardiff best SRH-4 solution with 2 routes (red, lime), 156 students, average walk 10.12 minutes, and average journey time 33.56 minutes. 240

B.35 Milton Keynes best SRH-4 solution with 4 routes (red, lime, blue, yellow), 274 students, average walk 11.41 minutes, and average journey time 14.76 minutes. . . . . 240

B.36 Bridgend best SRH-4 solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 381 students, average walk 11.75 minutes, and average journey time 28.32 minutes. . . . . 241

B.37 Edinburgh-2 best SRH-4 solution with 4 routes (red, lime, blue, yellow), 320 students, average walk 11.10 minutes, and average journey time 14.78 minutes. . . . . 241

B.38 Edinburgh-1 best SRH-4 solution with 9 routes (red, lime, blue, yellow, magenta, cyan, black, brown, green), 680 students, average walk 10.28 minutes, and average journey time 15.64 minutes. . . . 242

B.39 Adelaide best SRH-4 solution with 8 routes (red, lime, blue, yellow, magenta, cyan, black, brown), 565 students, average walk 12.01 minutes, and average journey time 15.80 minutes. . . . . 242

B.40 Brisbane best SRH-4 solution with 10 routes (red, lime, blue, yellow, magenta, cyan, black, brown, green, navy), 757 students, average walk 11.85 minutes, and average journey time 20.70 minutes. 243

C.1 Mġarr best independent solution with 4 routes (red, lime, blue, yellow), 190 students, average walk 6.02 minutes, average journey time 13.63 minutes, and average percentile time 24.49 minutes. . . 248

C.2 Mellieħa best independent solution with 4 routes (red, lime, blue, yellow), 171 students, average walk 4.98 minutes, average journey time 14.19 minutes, and average percentile time 24.94 minutes. . . 248

C.3 Porthcawl best independent solution with 2 routes (red, lime), 66 students, average walk 12.38 minutes, average journey time 12.14 minutes, and average percentile time 26.39 minutes. . . . . 249

C.4 Qrendi best independent solution with 5 routes (red, lime, blue, yellow, magenta), 255 students, average walk 6.67 minutes, average journey time 15.15 minutes, and average percentile time 27.21 minutes. . . . . 249

C.5 Suffolk best independent solution with 4 routes (red, lime, blue, yellow), 209 students, average walk 8.05 minutes, average journey time 28.86 minutes, and average percentile time 38.31 minutes. . . 250

C.6	Senglea best independent solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 266 students, average walk 7.09 minutes, average journey time 12.18 minutes, and average percentile time 23.11 minutes. . . . .	250
C.7	Victoria best independent solution with 4 routes (red, lime, blue, yellow), 171 students, average walk 7.32 minutes, average journey time 23.26 minutes, and average percentile time 38.21 minutes. . .	251
C.8	Pembroke best independent solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 335 students, average walk 6.78 minutes, average journey time 14.94 minutes, and average percentile time 27.30 minutes. . . . .	251
C.9	Canberra best independent solution with 8 routes (red, lime, blue, yellow, magenta, cyan, black, brown), 499 students, average walk 6.97 minutes, average journey time 23.60 minutes, and average percentile time 32.26 minutes. . . . .	252
C.10	Ħandaq best independent solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 285 students, average walk 6.80 minutes, average journey time 16.19 minutes, and average percentile time 28.45 minutes. . . . .	252
C.11	Valetta best independent solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 268 students, average walk 6.67 minutes, average journey time 17.67 minutes, and average percentile time 34.06 minutes. . . . .	253
C.12	Birkirkara best independent solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 306 students, average walk 6.54 minutes, average journey time 15.44 minutes, and average percentile time 26.84 minutes. . . . .	253
C.13	Ħamrun best independent solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 321 students, average walk 6.66 minutes, average journey time 14.14 minutes, and average percentile time 25.04 minutes. . . . .	254
C.14	Cardiff best independent solution with 4 routes (red, lime, blue, yellow), 156 students, average walk 10.06 minutes, average journey time 16.98 minutes, and average percentile time 37.64 minutes. . .	254
C.15	Milton Keynes best independent solution with 4 routes (red, lime, blue, yellow), 274 students, average walk 11.44 minutes, average journey time 14.76 minutes, and average percentile time 27.97 minutes. . . . .	255

C.16	Bridgend best independent solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 381 students, average walk 11.82 minutes, average journey time 25.39 minutes, and average percentile time 35.40 minutes. . . . .	255
C.17	Edinburgh-2 best independent solution with 4 routes (red, lime, blue, yellow), 320 students, average walk 10.81 minutes, average journey time 15.03 minutes, and average percentile time 27.52 minutes. . . . .	256
C.18	Edinburgh-1 best independent solution with 9 routes (red, lime, blue, yellow, magenta, cyan, black, brown, green), 680 students, average walk 10.28 minutes, average journey time 15.62 minutes, and average percentile time 28.51 minutes. . . . .	256
C.19	Adelaide best independent solution with 8 routes (red, lime, blue, yellow, magenta, cyan, black, brown), 565 students, average walk 12.12 minutes, average journey time 15.90 minutes, and average percentile time 30.32 minutes. . . . .	257
C.20	Brisbane best independent solution with 10 routes (red, lime, blue, yellow, magenta, cyan, black, brown, green, navy), 757 students, average walk 11.85 minutes, average journey time 20.91 minutes, and average percentile time 27.97 minutes. . . . .	257
D.1	Mġarr best correlated solution with 4 routes (red, lime, blue, yellow), 190 students, average walk 6.01 minutes, average journey time 13.68 minutes, and average percentile time 24.82 minutes. . .	260
D.2	Mellieħa best correlated solution with 4 routes (red, lime, blue, yellow), 171 students, average walk 5.04 minutes, average journey time 14.21 minutes, and average percentile time 25.70 minutes. . .	261
D.3	Porthcawl best correlated solution with 2 routes (red, lime), 66 students, average walk 12.38 minutes, average journey time 12.14 minutes, and average percentile time 26.31 minutes. . . . .	261
D.4	Qrendi best correlated solution with 5 routes (red, lime, blue, yellow, magenta), 255 students, average walk 6.70 minutes, average journey time 15.18 minutes, and average percentile time 27.42 minutes. . . . .	262
D.5	Suffolk best correlated solution with 4 routes (red, lime, blue, yellow), 209 students, average walk 7.69 minutes, average journey time 29.17 minutes, and average percentile time 39.21 minutes. . .	262

D.6	Senglea best correlated solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 266 students, average walk 6.49 minutes, average journey time 12.36 minutes, and average percentile time 22.46 minutes. . . . .	263
D.7	Victoria best correlated solution with 5 routes (red, lime, blue, yellow, magenta), 171 students, average walk 7.29 minutes, average journey time 18.51 minutes, and average percentile time 31.94 minutes. . . . .	263
D.8	Pembroke best correlated solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 335 students, average walk 7.21 minutes, average journey time 14.96 minutes, and average percentile time 27.99 minutes. . . . .	264
D.9	Canberra best correlated solution with 8 routes (red, lime, blue, yellow, magenta, cyan, black, brown), 499 students, average walk 7.21 minutes, average journey time 23.66 minutes, and average percentile time 32.44 minutes. . . . .	264
D.10	Ħandaq best correlated solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 285 students, average walk 6.79 minutes, average journey time 16.06 minutes, and average percentile time 28.58 minutes. . . . .	265
D.11	Valetta best correlated solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 268 students, average walk 6.84 minutes, average journey time 18.07 minutes, and average percentile time 34.15 minutes. . . . .	265
D.12	Birkirkara best correlated solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 306 students, average walk 6.62 minutes, average journey time 15.62 minutes, and average percentile time 27.82 minutes. . . . .	266
D.13	Ħamrun best correlated solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 321 students, average walk 6.99 minutes, average journey time 14.16 minutes, and average percentile time 25.44 minutes. . . . .	266
D.14	Cardiff best correlated solution with 4 routes (red, lime, blue, yellow), 156 students, average walk 9.65 minutes, average journey time 17.41 minutes, and average percentile time 38.57 minutes. . .	267
D.15	Milton Keynes best correlated solution with 4 routes (red, lime, blue, yellow), 274 students, average walk 11.54 minutes, average journey time 14.76 minutes, and average percentile time 28.31 minutes. . . . .	267

D.16 Bridgend best correlated solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 381 students, average walk 11.83 minutes, average journey time 25.43 minutes, and average percentile time 35.52 minutes. . . . . 268

D.17 Edinburgh-2 best correlated solution with 4 routes (red, lime, blue, yellow), 320 students, average walk 11.10 minutes, average journey time 14.78 minutes, and average percentile time 27.69 minutes. . . 268

D.18 Edinburgh-1 best correlated solution with 9 routes (red, lime, blue, yellow, magenta, cyan, black, brown, green), 680 students, average walk 10.24 minutes, average journey time 15.62 minutes, and average percentile time 29.24 minutes. . . . . 269

D.19 Adelaide best correlated solution with 8 routes (red, lime, blue, yellow, magenta, cyan, black, brown), 565 students, average walk 12.23 minutes, average journey time 16.04 minutes, and average percentile time 28.82 minutes. . . . . 269

D.20 Brisbane best correlated solution with 10 routes (red, lime, blue, yellow, magenta, cyan, black, brown, green, navy), 757 students, average walk 12.02 minutes, average journey time 20.53 minutes, and average percentile time 27.58 minutes. . . . . 270

# List of Acronyms and Initials

<b>Abbreviation</b>	<b>Full Form</b>
3DM	3-Dimensional Matching
3-SAT	3-Satisfiability
ACO	Ant Colony Optimization
ALNS	Adaptive Large Neighbourhood Search
ARL	Allocation-Routing-Location
B&B	Branch-and-Bound
B&C	Branch-and-Cut
B&P	Branch-and-Price
BTI	Buffer Time Index
C	Vehicle Capacity
CAP	Capacity
CCP	Chance-Constrained Program
CDF	Cumulative Distribution Function
CF	Characteristic Function
CG	Column Generation
COL	Chance of Being Late
COO	Chance of Overcrowding
CoV	Coefficient of Variation
CU	Capacity Utilization
CVRP	Capacitated Vehicle Routing Problem
CWS	Clarke-Wright Savings
DBSCAN	Density-Based Spatial Clustering of Applications with Noise
EPT	Earliest Pick-up Time
FHWA	Federal Highway Administration
FOC	Frequency of Congestion
FR	Feasibility Rate
GA	Genetic Algorithm
GAP	Generalized Assignment Problem
GE	General Education
GENI	Generalized Insertion Procedure
GIS	Geographic Information System

GPS	Global Positioning System
GRASP	Greedy Randomized Adaptive Search Procedure
HCM	Highway Capacity Manual
HFVRP	Heterogeneous Fleet Vehicle Routing Problem
ILP	Integer Linear Programming
ILS	Iterated Local Search
IR	Infeasible Runs
JT	Journey Time
LAR	Location-Allocation-Routing
LB	Load-Balancing or Ride Time-Balancing
LRA	Location-Routing-Allocation
LQ	Lower Quartile
MGF	Moment-Generating Function
MI	Misery Index
MILP	Mixed Integer Linear Programming
MIP	Mixed Integer Programming
MRL	Maximum Route Length
MRT	Maximum Riding Time
MSN	Minimum Route Load
MSR	Maximum Stops per Route
MT	Mean Time
MTZ	Miller-Tucker-Zemlin
MWT	Maximum Walking Distance or Time
N	Number of Routes
NN	Nearest Neighbour
NP	Nondeterministic Polynomial
NS	Number of Stops
NT	Number of Transfers
OVRP	Open Vehicle Routing Problem
PDA	Pairs of Distinct Arcs
PDF	Probability Density Function
PM	Perturbation Mechanism
PNNH	Parallel Nearest Neighbour Heuristic
PT	Percentile Time
PTC	Percentage of Travel Under Congestion
PTI	Planning Time Index
RCL	Restricted Candidate List
RIH	Regret Insertion Heuristic
RO	Robust Optimization

RVND	Random Variable Neighbourhood Descent
S-SBRP	Single-School Bus Routing Problem
SA	Simulated Annealing
SAT	Satisfiability
SBRP	School Bus Routing Problem
SBS	Number of Stops Shared by Multiple Routes
SD	Steepest Descent
SDVRP	Split Delivery Vehicle Routing Problem
SE	Special Education
SF	Safety Factor
SLN	Shifted Lognormal
SNA	Sweep Nearest Algorithm
SP	Stochastic Programming
SPP	Shortest Path Problem
SPR	Stochastic Program with Recourse
SRH	Sequential Random Heuristic
STD	Standard Deviation
SVRP	Stochastic Vehicle Routing Problem
SWD	Total Student Walking Distance
TBD	Total Bus Travel Time or Distance
TC	Total Cost
TCQSM	Transit Capacity and Quality of Service Manual
TJT	Total Journey Time
TL	Time Loss
TMT	Total Mean Time
TPT	Total Percentile Time
TRC	Trip Compatibility
TS	Tabu Search
TSD	Total Student Riding Time or Distance
TSP	Travelling Salesman Problem
TTR	Travel Time Reliability
TW	Time Window
UQ	Upper Quartile
VND	Variable Neighbourhood Descent
VRP	Vehicle Routing Problem
VRPB	Vehicle Routing Problem with Backhauls
VRPPD	Vehicle Routing Problem with Pickup and Delivery
VRPSPD	Vehicle Routing Problem with Simultaneous Pickup and Delivery
VRPTW	Vehicle Routing Problem with Time Windows



To my son Kai, my heart's greatest joy.

And to Auntie Maria and Nanna Mary –  
whose love continues to guide me from heaven.



# Acknowledgements

First and foremost, I thank God for granting me the strength and wisdom to undertake and complete this work. Without His guidance and blessings, this milestone would not have been possible.

I would like to express my deepest gratitude to my supervisors, Prof. Rhyd Lewis and Prof. Jonathan Thompson, for their guidance, support, and insightful feedback throughout the course of this research. Their mentorship has been invaluable, and I am sincerely thank them for their time and expertise.

To Mum and Dad, thank you for instilling in me the values of perseverance, humility, and hard work. You have been my foundation and my constant source of strength.

To my husband Rodrick, thank you for walking this journey by my side, cheering me on through every challenge, and offering endless encouragement when it was most needed. Your love and belief in me carried me forward.

To my son Kai, thank you for filling my heart with love and purpose. You inspire me to dream bigger and to be the best version of myself.

A heartfelt thank you also goes to my in-laws and Auntie Carmen for the immense support they gave by caring for Kai while I was focused on my studies. Your help made all the difference.

I would also like to acknowledge the financial support received from the University of Malta Scholarship for Part-Time Study, and the Tertiary Education Scholarship Scheme (TESS). Their generous contributions were instrumental in enabling me to pursue and complete my PhD.

Finally, to all those who supported me throughout this journey – family, friends, and colleagues – thank you for your encouragement and kindness. Special thanks to Mr. Kristian Gixti for introducing me to the field of parallel computing and for his generous assistance throughout the final phase of the experiments.



# Chapter 1

## Introduction

The *school bus routing problem* (SBRP) is a combinatorial optimization problem in transportation that focuses on designing efficient, safe, and reliable routes for transporting students to and from school. In many countries, school bus transportation forms part of the government's administrative mechanism and is funded through taxation. Students who live beyond a certain distance from their school are eligible for free or subsidized transportation. By optimizing school bus routes, governments can better utilize resources such as fuel and drivers' time, reduce traffic congestion, promote environmental sustainability, ensure students arrive at school on time, alleviate stress for both students and parents, and enhance safety by minimizing travel time and avoiding congested or hazardous areas. While optimizing these routes offers numerous benefits, it also presents challenges. Schools, in partnership with governments, must continuously adjust bus routes to accommodate yearly changes in student populations. Factors such as vehicle and driver availability, student safety concerns, traffic patterns, school start and end times, as well as road and weather conditions, must all be considered when designing bus routes.

Transportation authorities usually need to work closely with bus fleet owners and operators to create a cohesive system for coordinating routes, optimizing resources, and exploring sustainable initiatives such as electric buses. Fuelled by technological advancements, operators are increasingly implementing routing and real-time tracking software to facilitate route planning, monitor drivers and students on buses, respond quickly in case of emergencies (e.g., vehicle breakdowns), and enhance communication with parents (e.g., alerting them when their child's bus is delayed, or approaching the designated stop). By tracking the location of school buses, operators can also conduct route compliance checks against initial travel time estimates to identify areas for improvement. Such analysis and adjustments are crucial, especially since school bus fleets often operate on limited budgets. These financial challenges can be alleviated by reducing the costs

associated with implementing and running the service through careful planning and execution. With a well-organized transportation system, schools can remain focused on students' well-being and their overarching goal of providing quality education.

## 1.1 SBRP Decomposition

First investigated by Newton and Thomas (1969), the SBRP falls into a larger class of problems known as *vehicle routing problems* (VRPs). These problems involve designing optimal delivery or collection routes from one or more depots to a finite set of geographically scattered customers, subject to various constraints (Laporte et al., 1988). A common objective in VRPs is to minimize the total operating costs of the vehicle fleet. Typical constraints include maximum capacity restrictions on vehicles and maximum time or distance restrictions on routes.

The SBRP presents a challenging task from both logistical and financial perspectives. Typically, the SBRP involves compiling a list of home addresses for all students requiring school transportation, identifying potential bus stops accessible to the students, determining the stops to be visited by the buses, assigning students to their respective bus stops, and designing routes that optimize operational efficiency without compromising safety or service quality. These objectives are often conflicting, as improvements in service quality can increase costs, and vice-versa.

Desrosiers et al. (1981) decompose the SBRP into five subproblems. These subproblems are described briefly below and will be discussed in more detail in Chapter 2. The first subproblem, *data preparation*, generates a network that includes the school(s), students' home addresses, potential bus stop locations, and bus depots. Also considered are the number of students requiring transportation at each address, the destination school of each student, the types and number of buses available, and their capacities. The second subproblem, *bus stop selection*, concerns the selection of bus stops to be visited from the set of all potential bus stops in the relevant area, and the assignment of students to these stops. In this subproblem, local policies, such as students' maximum walking distance, are also considered.

The third subproblem, *route generation*, is the core of the SBRP and involves constructing the optimal set of bus routes, subject to several constraints. One important restriction is the maximum riding time for students, i.e., the maximum time students should spend on the bus. The last two subproblems, *school bell time adjustment* and *route scheduling*, are specific to multischool routing problems. They concern the optimization of transportation logistics across multiple schools

simultaneously, addressing the increased complexity of serving several student populations and school schedules. The school bell time adjustment subproblem optimizes the school bell times (i.e., the opening and closing times) to allow buses to service multiple schools. Meanwhile, the route scheduling subproblem assigns sequences of routes to buses based on the school bell times.

## 1.2 SBRP Classifications

Any VRP, and in particular an SBRP, can be *static* or *dynamic*, and *deterministic* or *stochastic*. All four combinations exist and by classification, we refer to one of the following combinations:

- static and deterministic;
- static and stochastic;
- dynamic and deterministic;
- dynamic and stochastic.

If the required inputs in the data preparation subproblem are available before the routes are constructed and do not change thereafter, the SBRP is classified as static. Generally, if the problem involves creating a set of preplanned routes that are not reoptimized and are based on inputs that do not change in real time, it is considered static. In contrast, if the inputs (such as the number of students requiring transportation) evolve and are updated concurrently with the construction of the bus routes, the problem is characterized as dynamic.

A static SBRP is deterministic if all of its inputs are assumed to be known with certainty and remain constant. In this case, the problem operates under fixed conditions. For example, a school district must plan bus routes for the upcoming academic year. The routes will remain fixed because the demand (the number of students and their addresses) is known in advance and will not change throughout the year. Additionally, the travel times between bus stops are known with certainty and are assumed to remain unaffected by traffic conditions, weather, accidents, or other factors.

A static SBRP is stochastic if there is uncertainty in some of its inputs, and this uncertainty is only revealed during the execution of the routes. For example, a school district must plan bus routes for the upcoming academic year. The routes will be preplanned and remain unchanged (static), but there are uncertainties about some parameters, such as student absenteeism and traffic conditions. These uncertainties can be modelled using probability distributions, making the problem stochastic. The number of students present for transportation on a given day is uncertain and may vary. Historical data could provide the probability of a student being absent on any given day. Similarly, the travel time between geographic

locations can fluctuate and can be estimated using probability distributions based on historical traffic data.

An SBRP is dynamic and deterministic when its inputs are fixed, but new information or changes may arise during the execution phase that require real-time route adjustments. For example, a bus driver may encounter a road closure due to construction or an accident. The driver receives this information in real-time via global positioning system (GPS) tracking and dynamically adjusts the route to avoid the closed road. Additionally, a last-minute adjustment may be needed if a parent calls with an emergency request to change a student's pick-up or drop-off location.

Lastly, an SBRP is dynamic and stochastic if probabilistic information is available about the inputs that change over time, and the routes are updated as these inputs are revealed during execution. For example, if the system receives absenteeism information while a bus is already en route, the driver may be instructed to skip certain stops.

### 1.3 Research Aims

This thesis focuses on the static and single-school variant of the SBRP (S-SBRP), addressing both deterministic and stochastic versions of the problem. The main aims of the research in this thesis are as follows:

- To evaluate a new mixed integer programming formulation for the S-SBRP that accounts for several realistic operational constraints.
- To evaluate and compare the performance of various bespoke constructive and improvement (local search) heuristics for the S-SBRP.
- To investigate the impact of different mechanisms for varying bus stop selection on the quality of the resulting solutions.
- To explore the effects of incorporating stochastic travel times, as opposed to deterministic ones, within a heuristic algorithmic framework.
- To assess whether incorporating correlations between travel times leads to better solutions, and to evaluate the additional computational complexity introduced by this consideration.

### 1.4 Thesis Contributions

The contributions of this thesis are as follows.

- In Chapter 3, we propose and evaluate a mixed integer programming formulation for the S-SBRP. This formulation incorporates bus stop selection, maximum student walking distances, bus dwell times, maximum bus riding

times, and multiple visits to bus stops. To our knowledge, no other formulation that includes all of these features exists in the literature. The closest formulations are those by Riera-Ledesma and Salazar-González (2013) (which does not address bus dwell times or multiple visits to bus stops); Schittekat et al. (2013) (which does not address bus dwell times, maximum bus riding time, or multiple visits to bus stops); and Sales et al. (2018) (which does not address bus dwell times or maximum bus riding time). It should be noted that presenting this formulation does not imply that the problem can always be solved using exact solution methods, given its underlying computational complexity. Rather, the purpose of the formulation is to clarify the constraints and help the reader to compare it with other formulations in the literature.

- In Chapter 4, we describe an iterated local search algorithmic framework for the S-SBRP. We propose three constructive heuristics – Parallel Nearest Neighbour Heuristic (PNNH), Sequential Random Heuristic (SRH), and Regret Insertion Heuristic (RIH) – with the first and third being original contributions. The RIH heuristic is an extended version of a procedure proposed by Pacheco and Marti (2006), which, unlike the original, can handle multiple visits to bus stops. Additionally, we consider six improvement heuristics and design four variants of an operator that alters the selection of bus stops, and repairs a given set of routes based on the newly selected stops (Perturbation Mechanism (PM) 1 to 4). The PM 2 variant is featured in the work of Lewis and Smith-Miles (2018), while the other variants are original contributions. Computational results show that SRH combined with PM 4, or RIH combined with PM 4, generally produce superior solutions. Further sensitivity analysis reveals that the generalized Or-opt and cross-exchange improvement heuristics are the most impactful, and that random variable neighbourhood descent performs marginally better than steepest descent.
- In Chapter 5, we propose an extension of the algorithmic framework from the previous chapter that incorporates stochastic and independent travel times. We assume that each travel time follows a shifted lognormal distribution, with an adjustable free-flow component and coefficient of variation. Our extension accommodates on-time arrival constraints for the routes, which can be adjusted by varying the travel time reliability level. Computational results highlight the following advantages of modelling stochastic travel times: identifying infeasibility cases at very high travel time reliability levels, recognizing the need for additional buses to guarantee on-time arrival, and improving travel time reliability while maintaining reasonable bus capacity utilization and mean journey times.

- In Chapter 6, we further extend our framework by allowing stochastic travel times to exhibit spatial correlation. We introduce a novel classification of pairs of arcs, which we use to model varying levels of spatial correlation. To our knowledge, one category in this classification has not yet been explored in the literature. It concerns pairs of arcs that overlap across certain road segments, giving rise to inherent interactions resulting from shared traffic flow along these segments. Spatial correlations capture cascading effects, where traffic flow disruptions in one area lead to delays in other areas. Computational results show that this interdependence increases the variability of travel times, directly impacting travel time reliability.
- In Chapter 6, we also present a step-by-step guide for generating Monte Carlo simulations of correlated shifted lognormal travel times. This is also applicable to other distributions. The proposed procedure helps to model and analyse real-world transportation networks under uncertainty. By randomly simulating a large number of travel times, a distribution of possible outcomes is generated. This approach enables transportation planners to evaluate the impact of traffic disruptions, estimate travel times more accurately, and gain better insights into the likelihood of on-time arrival.

Building on the work of Lewis et al. (2018), this study emphasizes practical applicability in both its problem formulation and data design. The formulation is based on real-world specifications provided by Welsh county councils, ensuring it captures the operational constraints and challenges faced by actual transportation authorities. Consultations with representatives from the Maltese government further informed our understanding of school transport issues and policy-driven considerations. Additionally, the data used in this study was generated from real-world schools and addresses using web mapping services, enabling accurate representation of geographic features and infrastructure. Consequently, walking and driving times and distances are derived from realistic spatial data, ensuring that travel patterns closely reflect real-world conditions. This contrasts with some studies in the literature that rely on artificial or very small-scale datasets, which do not accurately capture real-world travel dynamics.

## 1.5 Thesis Outline

This thesis is dedicated to the development of heuristic algorithmic frameworks for the S-SBRP. The next chapter provides an overview of existing literature on the different subproblems of the SBRP, with a primary focus on the first three subproblems. In Chapter 3, we first define and mathematically formulate the S-SBRP. We then present the real-world problem instances addressed in this

research, followed by results obtained from mixed integer programming. Finally, we prove the NP-hardness of the problem.

In Chapter 4, we shift our focus to heuristic approaches and develop an iterative heuristic algorithm that incorporates bus stop selection, route construction and improvement, and a solution destroy-and-repair operator. We analyse and compare the performance of three constructive heuristics, six improvement heuristics, and four variants of the destroy-and-repair operator under deterministic travel times. To make our framework more realistic and applicable in practice, we then extend it to incorporate stochastic travel times. We begin by assuming that travel times are independent in Chapter 5, and later drop this assumption to allow correlations between travel times in Chapter 6. Chapter 6 also presents the results of Monte Carlo simulations conducted to assess the feasibility and quality of the solutions presented in Chapters 4 to 6. Finally, Chapter 7 provides concluding remarks and discusses the outcomes and possible directions for future work.

## 1.6 Academic Publications

The following academic articles have been produced from the research presented in this thesis:

- Sciortino, M., Lewis, R., and Thompson, J. (2021). A heuristic algorithm for school bus routing with bus stop selection. In Zarges, C, and Verel, S., editors, *Evolutionary Computation in Combinatorial Optimization: 21st European Conference, EvoCOP 2021, Held as Part of EvoStar 2021, Virtual Event, April 7–9, 2021, Proceedings*, pages 202-218. Springer, Cham, Switzerland.
- Sciortino, M., Lewis, R. and Thompson, J. (2023). A school bus routing heuristic algorithm allowing heterogeneous fleets and bus stop selection. *SN Computer Science*, 4:74.

These articles correspond to content presented in Chapters 3 and 4 of this thesis. For both articles, I served as the principal author, responsible for developing the underlying methodology, conducting the computational experiments, and preparing the manuscripts. My supervisors contributed by reviewing the drafts and providing feedback on their structure and clarity. A third article, related to the work presented in Chapters 5 and 6, is in preparation and will be submitted to a journal within the next few months.



# Chapter 2

## Literature Review

This chapter provides a review of the current literature relevant to the field of school bus routing. We begin with a brief overview of the VRP and some of its main variants in Section 2.1. The focus then shifts to the SBRP in Section 2.2, where we summarize the key characteristics commonly addressed in the literature, including problem assumptions, objectives, and constraints. Finally, in Section 2.3, we provide an overview of the solution methods proposed to tackle the SBRP subproblems.

### 2.1 The VRP and its Variants

The VRP is a generalization of the *travelling salesperson problem* (TSP), which seeks the shortest possible route that visits every city in a given set and returns to the starting point. The VRP was originally proposed by Dantzig and Ramser (1959), who initially referred to it as the *truck dispatching problem*. They modelled it as an extension of the TSP, incorporating additional complexities such as multiple routes and customer demands to be fulfilled by homogeneous vehicles with limited carrying capacity. This formulation represents the classical variant of the VRP, also known as the *capacitated VRP* (CVRP). Dantzig and Ramser (1959) considered a real-world application minimizing the total mileage travelled by a fleet of gasoline delivery trucks between one central hub and several service stations. In their work, they not only proposed the first mathematical programming formulation of the problem but also introduced the first heuristic approach for addressing it. Five years after the initial introduction of the VRP, Clarke and Wright (1964) improved upon Dantzig and Ramser (1959)'s approach by developing the so-called savings heuristic. Building on these two pioneering studies, the VRP has since attracted increasing attention in the Operations Research community, with numerous variants being studied over the years. To commemorate the fiftieth anniversary of the VRP, Laporte (2009) published a review highlighting

the key contributions in the problem's history.

In a taxonomic review of 277 papers published between 2009 and 2015, Braekers et al. (2016) identified five main problem characteristics that can be used to categorize different variants of the VRP. These were based on the taxonomy developed by Eksioglu et al. (2009), who reviewed 1,021 VRP articles published between 1959 and 2008. The five characteristics are as follows:

1. *Type of study*: for example, the applied solution method;
2. *Scenario characteristics*: for example, the time horizon, whether demand splitting is allowed, the inclusion of time windows, and whether the problem involves both pickups and deliveries;
3. *Problem physical characteristics*: for example, the number of depots, the number of vehicles, whether vehicles have homogeneous or heterogeneous capacities, and the formulation of the objective function (time-based, distance-based, vehicle-dependent, costs related to lateness, etc.);
4. *Information characteristics*: for example, the evolution of information (static, dynamic, etc.) and the quality of information (deterministic, stochastic, unknown, etc.);
5. *Data characteristics*: whether the data used, if any, is synthetic or real-world.

The classical VRP model has been extended in various ways by incorporating additional real-world features. For example, in the *open VRP* (OVRP), the routes do not start and end at a depot, as in the classical VRP, but either start or end at a depot. Varying the vehicles' capacities leads to the *heterogeneous fleet VRP* (HFVRP), while allowing split deliveries – where customer demand can be divided and delivered by multiple vehicles – results in the *split delivery VRP* (SDVRP). Another well-known extension is the *VRP with time windows* (VRPTW), which assumes that deliveries must occur within specific time intervals, varying from one customer to another. There are also variants of the VRP involving both pickups and deliveries, such as the *VRP with pickup and delivery* (VRPPD), where goods are picked up and delivered in the same route, and the *VRP with backhauls* (VRPB), which involves a mix of customers requiring deliveries (linehauls) and others requiring pickups (backhauls).

As we shall see in Chapter 3, in this thesis we address an open SDVRP with capacity and time constraints. In relation to this research, Braekers et al. (2016) noted that heterogeneous vehicles, split deliveries, and stochastic travel times are considered in 16.51%, 6.12%, and 2.75% of the surveyed articles, respectively. Moreover, objective functions involving routing costs (either time-based or distance-based) appear in 92.35% of the articles, while vehicle-dependent costs are included in 38.23%. While most studies use artificial data, only 6.12% of the

articles utilize real-life data.

In another taxonomic review of 299 articles published between 2009 and 2017, Elshaer and Awad (2020) observed that the HFVRP, SDVRP, OVRP, and stochastic travel VRP are considered in 13.41%, 5.43%, 4.71%, and 3.26% of the articles, respectively. In a more recent review by Konstantakopoulos et al. (2022), which examined 263 articles published between 2010 and 2020, the corresponding percentages for the first three variants are 17.11%, 3.80%, and 4.56%, respectively.

Since we now turn our focus to the SBRP, which can be considered as a specialized variant of the VRP, the reader is referred to the following recent reviews for more comprehensive information on existing and emerging VRP variants: Vidal et al. (2020), Tan and Yeh (2021) (covering 88 articles published between 2019 and 2021), Zhang et al. (2022), and Hao et al. (2024) (covering articles published between 2017 and 2023).

## 2.2 SBRP Characteristics

The SBRP is characterized by numerous structural, operational, and contextual variations that influence how it is modelled and addressed. To better understand the complexity and diversity of the problem, this section reviews key classification dimensions including the number of schools involved, whether mixed loads are allowed, the nature of the service environment, differences between morning and afternoon routing, the composition of the bus fleet, the objectives pursued in the routing process, and the various constraints that must be respected.

### 2.2.1 Number of Schools and Allowance of Mixed Loads

An SBRP may involve either a single school or multiple schools. Most studies in the literature focus on the single-school case, including the original study by Newton and Thomas (1969). Other examples include Dulac et al. (1980), Bowerman et al. (1995), Li and Fu (2002), Bektaş and Elmastaş (2007), Schittekat et al. (2013), Kinable et al. (2014), Siqueira et al. (2016), Sales et al. (2018), Dang et al. (2019), and Hou et al. (2022b). In multischool scenarios, students from different schools can share the same bus to improve vehicle utilization. According to Spada et al. (2005), there are two primary approaches to route generation in this setting: school-based and home-based.

The school-based approach does not permit students from different schools to be transported together on the same bus simultaneously (i.e., no *mixed loads*). Instead, it generates separate route sets for each school, which are then assigned

to the vehicle fleet based on the schools' bell times. In contrast, the home-based approach allows mixed loads. This latter approach was developed by Braca et al. (1997), who proposed a simple insertion rule that adds one stop to a route at a time. If the corresponding school for the stop is not already included in the route, the best insertion point for the school is also determined. Notable examples of SBRP works involving multiple schools include those of Angel et al. (1972), Bodin and Berman (1979), Desrosiers et al. (1981), Caceres et al. (2017), and Bertsimas et al. (2019) for configurations without mixed loads, and Spada et al. (2005), Park et al. (2012), Campbell et al. (2015), Calvete et al. (2020), and Ozmen and Sahin (2021) for those that permit mixed loads.

Ellegood et al. (2020) noted that more than half of the multischool studies allow mixed loads. The advantages of mixed loads include greater flexibility, potential cost savings, and the fact that bus stops are typically visited only once (rather than being served by multiple buses or skipped due to students attending different schools than those already on board). However, mixed loads also present some drawbacks, such as additional travel between schools, increased student waiting times, and the need for greater coordination. For a discussion of the various mixed load transportation policies explored in the SBRP literature, the reader is referred to Ellegood et al. (2020).

### 2.2.2 Service Environment

The structure of an SBRP also varies depending on whether the service environment is urban or rural. Urban areas are characterized by dense street networks and more developed pedestrian infrastructure, which can facilitate safer conditions for students walking between their homes and designated bus stops. However, actual safety levels may vary depending on factors such as traffic volume, neighbourhood crime rates, and the availability of pedestrian crossings. Furthermore, the high student density in urban settings means that many students live within a reasonable walking distance of a shared bus stop. In contrast, in rural areas, buses often pick up and drop off students directly at their homes. As a result, the bus stop selection subproblem is typically not considered in rural settings.

Chen et al. (2015) highlighted several ways in which the rural school environment differs, including lower population density, more stops per route, fewer students per stop, longer travel distances, scarcity or absence of one-way streets, a higher likelihood of buses being parked overnight at drivers' homes, and fewer alternative routes. These distinctions often cause rural routes to be constrained by a maximum riding time before the bus capacity is fully utilized. On the other

hand, urban routes typically reach full bus capacity before the maximum riding time is exceeded (Bowerman et al., 1995). Ripplinger (2005) observed that, due to the small problem size, manually generated solutions can be optimal for rural areas. Rural SBRPs have also been addressed in the works of Thangiah and Nygard (1992), Corberán et al. (2002), Spada et al. (2005), Faraj et al. (2014), Souza Lima et al. (2016), Sun et al. (2018), and Miranda et al. (2021).

### **2.2.3 Morning Versus Afternoon Problem**

In school bus routing, the morning problem involves planning transportation from students' homes to school. Braca et al. (1997) state that this problem is generally more challenging than the afternoon problem, which concerns transportation from school to students' homes. They attribute this to typically heavier traffic in the morning and to school start times being less dispersed than dismissal times. The latter is particularly relevant in multischool configurations. As noted by Park and Kim (2010), for these reasons, the afternoon problem has received less attention in the literature. However, the afternoon problem can often be derived from the morning problem with just minor modifications. For example, morning routes may be reversed (a "last on, first off" policy, resulting in shorter afternoon routes) or replicated (a "first on, first off" policy, promoting more balanced riding times) (Bodin and Berman, 1979). That said, morning and afternoon routes may not always mirror each other exactly (Ellegood et al., 2020). For instance, the number of students requiring transportation can differ between morning and afternoon, as some parents may be available to drop off or pick up their children. Additionally, students may participate in before- or after-school programmes, such as breakfast clubs, sports, or arts and crafts.

### **2.2.4 Fleet Mix**

The composition of a school bus fleet is often determined by institutional policies related to student safety, comfort, and regulatory standards. A central consideration is the level of crowding permitted on buses, which is typically established by individual schools or governing authorities. These policies may include whether standing passengers are allowed and often vary depending on the age of the students, with more restrictive rules generally applied to younger children.

The term 'fleet mix' refers to whether the buses in operation have uniform (homogeneous) or differing (heterogeneous) capacities. A homogeneous fleet is appropriate, for example, in rural areas where buses typically do not operate at full capacity and routes are more often constrained by maximum riding time limits. A heterogeneous fleet can be modelled by assigning a distinct capacity to

each bus (Kim et al., 2012) or by assuming the existence of different bus types, with all buses of the same type having equal capacity (Chen et al., 2015). While heterogeneous fleets require more complex planning and may necessitate a larger number of buses to transport the same number of students, they offer greater flexibility in terms of scheduling and stop selection (Mohd Azmi et al., 2024).

Park et al. (2012) noted that only 3 of 13 single-school publications and 2 of 13 multischool publications considered a heterogeneous fleet. In the review by Ellegood et al. (2020), the corresponding figures are 3 out of 36 and 13 out of 28, respectively. Mohd Azmi et al. (2024) observed that the number of studies between 2018 and 2022 modelling homogeneous and heterogeneous fleets is roughly equivalent, while Díaz-Ramírez et al. (2024) reported that 63%, 24%, and 13% of a total of 49 papers (published after 2010) address homogeneous, heterogeneous, and unspecified fleets, respectively. Examples of studies that consider heterogeneous fleets include Thangiah and Nygard (1992), Li and Fu (2002), Spada et al. (2005), Siqueira et al. (2016), Souza Lima et al. (2017), Miranda et al. (2018), Mokhtari and Ghezavati (2018), Sales et al. (2018), Li and Chow (2021), and Paul et al. (2023).

## 2.2.5 Objectives

Savas (1978) proposed three performance measures for public services: *efficiency*, *effectiveness*, and *equity*. These were later adopted by Bowerman et al. (1995) to assess the quality of school bus service provision. We begin by outlining these measures and then present examples from the literature that apply them.

### 2.2.5.1 Efficiency

In the context of public services, efficiency refers to the ratio of service level to service cost. In cases where the service level is fixed, efficiency can be assessed based solely on the cost of service provision (Bowerman et al., 1995). Given the substantial global investment in school transportation, managing costs effectively is critical. For instance, the Maltese government allocated €63 million to provide free school transport to over 33,000 students during the 2023–2024 scholastic year (Vassallo, 2024).

A key strategy for improving efficiency is minimizing the number of routes ( $N$ ), as the major cost components include bus acquisition, maintenance, and drivers' salaries. As Bowerman et al. (1995) noted, "Since the capital cost is significantly larger per bus than the incremental cost over the year, the number of routes generated should be held to a minimum." Reducing the number of routes also helps mitigate issues such as driver or vehicle shortages.

Another important consideration is minimizing operational costs, particularly fuel expenses, by keeping routes as short as possible. However, this objective may conflict with the goal of minimizing the number of routes. This is because fewer routes will mean having to use larger buses and longer routes. This trade-off was evident in Malta during the 2023–2024 scholastic year when school transport operators attributed early pick-up times to the government’s ban on smaller vehicles (Carabott, 2023). Operational costs can be reduced by minimizing the total bus travel time or distance (TBD).

### **2.2.5.2 Effectiveness**

Effectiveness is evaluated based on how successfully the demand is met, i.e., whether the service is accessible to all eligible students and satisfactory. This evaluation becomes particularly relevant in real-world situations where service shortfalls may arise. For instance, at the beginning of the 2018–2019 scholastic year, nearly 1,000 students in Malta lacked transport arrangements due to insufficient bus capacity (Sansone, 2018). To assess how satisfactory the service is, one can analyse indicators such as total student riding time or distance (TSD) and total student walking distance (SWD). Another useful measure of effectiveness is students’ time loss (TL) (Spada et al., 2005), defined as the difference between the travel time of a direct trip from a student’s home to school and the time spent waiting for and riding the bus.

### **2.2.5.3 Equity**

Equity evaluates how fairly or impartially the service is provided to the eligible students. A solution that demonstrates high efficiency in terms of cost and time may still be deemed unacceptable by key stakeholders, such as education departments, if it lacks fairness. Load-balancing and ride time-balancing (LB) are two common equity criteria. The former aims to distribute students evenly across routes serving the same school, ensuring that each bus carries a similar number of students. The latter, on the other hand, seeks to minimize significant disparities in the duration of bus routes for a single school. Another approach to enhancing equity is the implementation of a “first on, first off” policy, which ensures that students on the same route experience similar travel times. Equity has received limited attention in the SBRP literature.

### **2.2.5.4 Other Objectives and Some Examples**

Most SBRP studies focus on efficiency criteria, such as minimizing the number of routes (N), total bus travel time or distance (TBD), or total cost (TC). Less

common objectives include minimizing the maximum route length (MRL), the number of stops (NS), or the number of stops shared by multiple routes (SBS); maximizing capacity utilization (CU) or trip compatibility (TRC); and incorporating penalties based on the safety factor (SF) of each traversed road segment or the maximum allowable number of transfers a student may make during their trip (NT). It is worth noting that several studies (35% in Díaz-Ramírez et al. (2024)) address multiple objectives, either through a hierarchical approach – where objectives are prioritized in a strict order – or through an aggregated approach, in which all objectives are combined into a single function using assigned weights. Examples of objectives considered in the literature are presented in Table 2.1.

Table 2.1: Examples of objectives considered in the literature.

Reference	Objectives
Bennett and Gazis (1972)	TBD, TSD
Thangiah and Nygard (1992)	N, TBD, TSD
Bowerman et al. (1995)	N, SWD, LB
Corberán et al. (2002)	N, MRL
Li and Fu (2002)	N, TSD, TBD, LB
Spada et al. (2005)	TL
Martínez and Viegas (2011)	SWD, TC (N, TBD)
Riera-Ledesma and Salazar-González (2012)	TC (N, SWD)
Minocha and Tripathi (2014)	CU
Bögl et al. (2015)	TBD, NT, SWD, NS, SBS
Chalkia et al. (2016)	TBD, SF
Shafahi et al. (2018b)	TC (CU, TRC, TBD)
Bertsimas et al. (2019)	SWD, N, TSD
Hou et al. (2022b)	LB, N, TBD
Paul et al. (2023)	TBD, N

### 2.2.6 Constraints

Like the more general VRP, the SBRP has been studied under a variety of constraints, including:

- Vehicle capacity ( $C$ ): the maximum number of students allowed on a bus;
- Maximum riding time (MRT): the maximum duration a student is allowed to spend on the bus during a single trip;
- School time window (TW): the acceptable time frame for bus arrivals at a school;
- Maximum walking distance or time (MWT): the maximum distance or time

- a student may walk from home to their assigned bus stop;
- Earliest pick-up time (EPT): the earliest time at which a student can be picked up in the morning;
  - Minimum route load (MSN): the minimum number of students required to create a route;
  - Number of routes (N): the total or maximum number of routes allowed in the solution;
  - Maximum route length (MRL): the maximum total distance that a bus is allowed to travel in a single route;
  - Maximum stops per route (MSR): the maximum number of stops allowed on a single route;
  - Chance of overcrowding (COO): the likelihood that buses exceed their capacity;
  - Chance of being late (COL): the likelihood that students arrive late at school.

The first four constraints are the most common, while the others occur less frequently. Most SBRP studies incorporate combinations of multiple constraints, as illustrated in Table 2.2.

Table 2.2: Examples of constraints considered in the literature.

Reference	Constraints
Dulac et al. (1980)	C, MRT, MWT
Braca et al. (1997)	C, MRT, TW, EPT, MSN
Spada et al. (2005)	C, TW
Bektaş and Elmastaş (2007)	C, MRT
Park et al. (2012)	C, MRT, TW
Riera-Ledesma and Salazar-González (2013)	C, MRT, MWT, MSR, MSN
Faraj et al. (2014)	C, MRT, MWT
Chen et al. (2015)	C, TW, N
Siqueira et al. (2016)	C, TW
Caceres et al. (2017)	C, MRT, COO, COL
Mokhtari and Ghezavati (2018)	C, MRT, TW
Dang et al. (2019)	C, MRT
Miranda et al. (2021)	C, MRT, TW, MWT
Calvete et al. (2022)	C, N
Xue et al. (2023)	C, MRL

## 2.3 Solution Methods for the SBRP

Having reviewed the different characteristics of the SBRP, this section discusses studies that have addressed its key subproblems: bus stop selection, route generation, school bell time adjustment, and route scheduling. Although these subproblems are closely interrelated, they are often tackled independently and sequentially due to the complexity of the overall problem.

### 2.3.1 Bus Stop Selection

Few studies have focused specifically on bus stop selection, and those that do primarily rely on heuristic algorithms. For example, only 7 out of 29 studies in Park et al. (2012), 17 out of 64 in Ellegood et al. (2020), 12 out of 33 in Mohd Azmi et al. (2024), and 17 out of 49 in Díaz-Ramírez et al. (2024) explicitly address this subproblem. In many other cases, bus stops are assumed to be predetermined. To better understand existing approaches, we divide the following review into two categories: studies that use mathematical models for bus stop selection and those that rely on heuristic methods.

#### 2.3.1.1 Mathematical Models

The bus stop selection subproblem can be formulated as an assignment problem, with the objective of minimizing criteria such as the number of stops or the total student walking distance. Dulac et al. (1980) initially focused on minimizing the total walking distance but later considered minimizing the number of stops, as this “seems to determine better routes with a smaller number of stops”. Martínez and Viegas (2011) proposed a mathematical model that minimizes students’ walking distances to bus stops while accounting for the capacity limits of each stop (determined by the bus capacity) and the maximum allowable walking time for each student.

Bögl et al. (2015) tested three modelling formulations for bus stop selection and student assignment, minimizing the distance to students’ destinations, the number of stops serving students with different destinations, and the total number of stops. They found that the latter two formulations use fewer stops, resulting in lower total travel times, albeit at the cost of increased time loss (as students may have to walk longer distances) and a higher average number of transfers. Bertsimas et al. (2019) applied an integer optimization model with constraints on maximum walking distance. The objective was to minimize the total number of stops, thereby simplifying bus routing and reducing pickup time by accounting for actions such as braking, opening and closing doors, conducting safety checks,

and waiting for students to board. However, as emphasized by Galdi and Thebpanya (2016), reducing the number of bus stops may compromise walking safety, particularly for students in areas without sidewalks and with heavier traffic. It may also cause behavioural or safety issues due to the higher concentration of students at fewer stops.

Li and Chow (2021) examined different bus stop selection strategies for three school types: general education (GE), special education (SE), and mixed schools. For SE schools, each student’s residence is treated as a fixed stop. For GE schools, the objective is to minimize the number of stops required to cover all students (an approach also used by Guo and Samaranayake (2022)). For mixed schools, which serve both GE and SE students, the bus stop selection process involves three steps. First, a capacitated maximum flow model assigns as many GE students as possible to SE stops within a 0.5-mile radius (referred to as “near-SE” students). Second, new stops are introduced to cover the remaining GE students. Finally, in Step 3, the allocation is refined by minimizing the walking distances of “near-SE” students, considering both SE stops and the new stops introduced in Step 2.

### 2.3.1.2 Heuristic Approaches

In the work of Faraj et al. (2014), students are divided into “groupable” and “non-groupable”. A student is groupable if they have at least one other student within maximum walking distance. Non-groupable students are treated as isolated cases, and are assigned a bus stop directly at their residence. For groupable students, stop locations are determined at students’ residences using a heuristic for the dominant set problem. This involves assigning a priority to each node in a graph, where the nodes represent student residences and the edges represent the road network. Each node’s priority is directly proportional to the square of its degree and inversely proportional to the sum of the degrees of its neighbouring nodes. Nodes are then added to the dominant set in order of decreasing priority, aiming to minimize the number of stops. Sarubbi et al. (2016) improved upon the work of Faraj et al. (2014) through a five-step procedure: generating evenly spaced points along the road network, building a topology, projecting each student onto a vertex, calculating the walking distance from each student to every vertex, and keeping only those within the allowed walking distance. Bus stops are then selected using a pseudo-random greedy algorithm.

Galdi and Thebpanya (2016) proposed a geographic information system (GIS)-based heuristic aimed at reducing the number of bus stops without compromising student safety or accessibility. Bus stops were evaluated based on their proximity to other stops. Stops located close to one another, or those that did not significantly reduce walking distances for students, were identified as candidates

for elimination. The stops marked for elimination by the heuristic were then re-evaluated by school officials, with each school level (elementary, middle, high) considered independently due to varying road safety classifications at different educational levels. The works of Sarubbi et al. (2016) (single school) and Galdi and Thebpanya (2016) (multiple schools) are among the few that consider the bus stop selection subproblem as a stand-alone problem. Generally, bus stop selection is frequently addressed in conjunction with route generation through one of three general strategies: *location-allocation-routing* (LAR), *allocation-routing-location* (ARL), and *location-routing-allocation* (LRA).

### 2.3.1.2.1 Location-Allocation-Routing (LAR)

In the LAR strategy, a set of bus stops is first selected, students are assigned to these stops, and routes are created based on the demand at each stop. This approach, used in the earlier studies by Bodin and Berman (1979), Dulac et al. (1980), and Desrosiers et al. (1981), may lead to suboptimal solutions since bus stop selection and student assignment are performed without considering capacity constraints. For example, a solution may include more routes than necessary if multiple stops have many assigned students (Bowerman et al., 1995). To address this limitation, Riera-Ledesma and Salazar-González (2012) proposed an integrated bi-objective model that minimizes both total bus route distance and student walking distance. However, their method may yield solutions where buses are underutilized. To mitigate this, they later introduced a constraint ensuring a minimum number of students per route (Riera-Ledesma and Salazar-González, 2013).

Huo et al. (2014) grouped student addresses into circular areas with a 500-metre radius, using the centre of each circle as a bus stop. Sales et al. (2018) assigned students to their nearest bus stop within a walking distance threshold. Kotoula et al. (2017) applied  $k$ -means clustering to partition students into four demand clusters, each forming a route. The clustering process minimized inter-stop distances and ensured each bus crossed the road network once per cluster. Following clustering, a multi-phase routing algorithm was employed, starting by selecting the student farthest from the school. Additional stops were added while ensuring compliance with trip duration and bus capacity constraints. Once all students in a cluster were assigned or constraints were reached, the route was further optimized using a genetic algorithm. Clustering was iteratively refined based on route efficiency and student coverage.

Guo et al. (2018) used  $k$ -means clustering to reduce the problem size by grouping students based on proximity. Each cluster represents a shared bus stop within the maximum walking distance or a direct pickup location near the cluster

centroid. To further simplify the problem, they employed a heuristic-based edge compression technique to eliminate unlikely ride-sharing node combinations by focusing only on spatially proximate nodes. Ren et al. (2019) developed an iterative clustering method combining  $k$ -means and DBSCAN (density-based spatial clustering of applications with noise) to determine bus stop locations and student assignments, subject to walking accessibility constraints. The method minimizes the number of stops by incrementally increasing the number of centroids ( $k$ ) until all constraints are satisfied. When student density is sufficient, initial centroids are derived from DBSCAN; for sparse distributions,  $k$  is calculated based on the school district's area and the maximum walking distance. Centroids are then uniformly placed across the district.

Ochoa-Zezzatti et al. (2020) selected stops within 500-metre zones of high student density and assigned students to their nearest stop. Köksal Ahmed et al. (2021) also followed the LAR strategy, emphasizing the importance of determining students' positions relative to road segments before routing to assess whether students need to cross roads. Their school-centred method used a radar-like scanning pattern to assign buses, dividing pick-up points into sectors where geographically close students share a bus. Solutions can be improved by adjusting the "phasing" or the initial angle of the radar scan.

A recent study by Calvete et al. (2023) introduced a bi-level decision process that accounts for student preferences in bus stop selection. At the upper level, stops are selected and cost-minimizing routes are constructed, while at the lower level, students choose their preferred stops. Due to capacity constraints, some student preferences may not be feasible. Any stop that would be selected by more students than a bus can accommodate is excluded, and infeasible scenarios arise when students only have access to "banned" stops. Parvasi et al. (2017) adopted a similar bi-level framework, in which a transportation firm at the upper level selects stops and routes to maximize profit, while students at the lower level decide whether to use the system or seek alternative transport, aiming to minimize their costs. The interaction between the levels reflects a trade-off for the firm: activating more stops may attract more students and increase revenue, but it also raises service costs.

Calvete et al. (2020) proposed an adaptation of the LAR approach. They first assigned students with only one feasible stop due to walking constraints. Remaining stops were then selected either via a mixed integer linear programming (MILP) model minimizing routing costs, or by randomly selecting a set that covers all students. If this set led to an infeasible assignment under capacity constraints, it was adjusted. A partial allocation of students was then used to design initial routes with low routing costs. If feasible, allocation and routing were completed;

otherwise, more students were included in the initial allocation.

### **2.3.1.2.2 Allocation-Routing-Location (ARL)**

The ARL strategy addresses the limitations of the LAR strategy by first clustering students while considering vehicle capacity constraints. For each cluster, routing is then performed using students' home locations (or representative points) rather than predefined bus stops. Finally, bus stops are located along these preliminary routes by aggregating nearby students while satisfying constraints such as maximum walking distance and the maximum number of students per stop. This approach was applied in earlier studies by Chapleau et al. (1985) and Bowerman et al. (1995); however, to our knowledge, it has not appeared in recent publications. Although the ARL strategy enables load balancing during the allocation phase, it cannot ensure balanced route lengths, as routes are constructed only after student allocation and depend on the spatial dispersion of students within each cluster (Bowerman et al., 1995).

### **2.3.1.2.3 Location-Routing-Allocation (LRA)**

To our knowledge, the third strategy, LRA, was first used by Schittekat et al. (2013). It begins by identifying the set of feasible bus stops for each student and constructing routes that minimize total travel distance using only these feasible stops (the master problem). In the next step, students are assigned to the selected stops while ensuring that bus capacity constraints are satisfied. This assignment is handled in a subproblem, which evaluates whether a feasible student-to-stop allocation exists for the solution generated by the master problem.

An alternative strategy was later proposed by Calvete et al. (2022), who claimed it yielded solutions with shorter total walking distances than the LRA method. They initialized  $b_0$  routes, each beginning with a single pickup point, by iteratively selecting the pickup point farthest from the depot and the previously selected points. Next, they considered the set of remaining pickup points and student locations. Elements from this combined set were selected in random order. If the selected element was a pickup point, it was inserted into the route that incurred the lowest insertion cost. If it was a student location, the student was assigned to the nearest available pickup point in terms of walking distance, subject to bus capacity constraints. In their work, the authors emphasized that jointly selecting pickup points and student locations, rather than handling them in separate stages, led to solutions with shorter total walking distances.

## 2.3.2 Route Generation

Generating school bus routes is a complex task influenced by factors such as the geographic distribution of students, time constraints, traffic conditions, road network connectivity, and vehicle capacity. Overcoming this challenge often requires the application of advanced optimization techniques. Among the most prominent are exact methods that provide precise, model-driven solutions. On the other hand, heuristic algorithms offer practical, often faster approaches that deliver good-enough solutions within reasonable time frames, especially for large-scale problems. Metaheuristic algorithms build on heuristics by incorporating adaptive mechanisms to explore solution spaces more effectively. Finally, hybrid approaches combine the strengths of multiple methods to enhance performance and flexibility. The following sections examine key studies, highlighting the application of exact methods and metaheuristics in school bus routing. Works featuring heuristics will be reviewed in Chapter 4.

### 2.3.2.1 Exact Methods

Although several researchers have described the routing subproblem using integer programming models, they do not solve these models using exact solution methods. This section focuses exclusively on selected works that employ exact solution techniques. Such techniques – including branch-and-bound, cutting planes, and column generation – guarantee optimality but often suffer from prohibitive computational times as problem size increases. As a result, their applicability is generally limited to relatively small instances, typically involving fewer than 50 nodes, according to Díaz-Ramírez et al. (2024). Two early S-SBRP studies that developed MILP models and solved them using commercial optimization solvers include those by Schittekat et al. (2006) (10 stops, 50 students) and Bektaş and Elmastaş (2007) (29 stops, 519 students).

Martínez and Viegas (2011) proposed a two-phase MILP approach to the bus stop selection and route generation subproblems. A capacitated  $p$ -median model minimizing the total walking distance is implemented in the first phase. The second phase employs a standard vehicle flow formulation to generate routes minimizing the total cost. Sun et al. (2018) developed a robust optimization model accounting for stochastic and time-dependent travel times and minimizing the maximum total cost. A small-scale test instance was then designed to evaluate the model’s effectiveness, with exact solutions obtained using the CPLEX solver.

More recently, Hülágü and Celikoglu (2022) formulated a flow-based MILP model to compare homogeneous and heterogeneous fleets. Their results show that the use of a heterogeneous fleet can reduce overall routing costs by up to 30%

while achieving near-full utilization of bus capacities. Guo and Samaranayake (2022) introduced a decomposition framework for solving the S-SBRP based on a shareability network that captures feasible groupings of students. They reformulated the routing subproblem as a weighted set covering problem, reducing dimensionality and solving it via the Gurobi solver. To enhance tractability, the framework incorporates two preprocessing strategies: node compression, which uses an integer linear programming (ILP) model to choose stops within students' walking limits and thus shrinks the network's node count, and heuristic edge pruning, which sparsifies the network by eliminating less promising connections.

The branch-and-bound (B&B) method decomposes the original problem into a hierarchy of subproblems, each representing a subset of the solution space. By evaluating bounds on the optimal solution within each subproblem, the method discards those subproblems that cannot improve upon the current best-known solution, thereby narrowing the search. Gavish and Shlifer (1979) addressed the route generation subproblem for a single-school scenario by formulating a non-linear mixed integer programming (MIP) model. To obtain upper bounds, they solved a series of assignment problems and used a B&B approach to determine optimal solutions. White (1982) enhanced the algorithm developed by Gavish and Shlifer (1979) by substituting the assignment problem with a maximum matching problem. Similarly, Kumar and Jain (2015) formulated the route generation subproblem using an assignment-based model and solved it with B&B, while Kim et al. (2012) applied B&B to both the route generation and route scheduling subproblems.

Like B&B, column generation (CG) is also often effective for solving integer programming problems with discrete and prohibitively large solution spaces. Instead of considering all possible variables (or columns) from the outset, which can be computationally overwhelming, CG begins with a smaller, manageable subset of variables and iteratively introduces new variables that have the potential to improve the solution. These promising variables are identified by solving a specialized subproblem known as the pricing problem. In the context of the S-SBRP, CG has been applied by López Santana and Romero Carvajal (2015) who integrated it with a stop clustering approach. Caceres et al. (2019) also employed a CG procedure coupled with a greedy heuristic to route SE students from different schools. CG can be combined with B&B to form the branch-and-price (B&P) method, which was used by Riera-Ledesma and Salazar-González (2013) on a set partitioning formulation. A B&P framework based on set covering was also proposed by Kinable et al. (2014), who solved several of their test instances to optimality – including some with up to 40 stops and 800 students – and established strong bounds for the remaining ones.

Methods that use cutting planes begin by solving a relaxed version of an integer programming problem, where integer constraints are temporarily ignored. If the solution is not integer-feasible, additional constraints known as cutting planes, or simply cuts, are added to exclude the current fractional solution without eliminating any feasible integer solutions. These cuts iteratively refine the feasible region, guiding the solution toward optimality. To improve efficiency and reduce computational time, cutting planes are often used in combination with other exact techniques. When integrated with B&B, the approach is known as branch-and-cut (B&C), which has been implemented by Riera-Ledesma and Salazar-González (2012).

### 2.3.2.2 Metaheuristic Approaches

To effectively address the computational challenges associated with the route generation subproblem, various metaheuristic approaches have been employed, owing to their ability to produce high-quality solutions within reasonable time frames. We begin by outlining the methodological foundations of metaheuristics and subsequently examine their application to the routing subproblem.

Metaheuristics are high-level search strategies designed to explore a solution space  $\mathcal{S}$  efficiently. Let  $f: \mathcal{S} \rightarrow \mathbb{R}$  be the objective function that assigns a numerical value to each candidate solution in  $\mathcal{S}$ , typically to be minimized. Rather than exhaustively searching  $\mathcal{S}$ , metaheuristics iteratively explore promising regions of  $\mathcal{S}$  by applying a search operator  $\mathcal{N}: \mathcal{S} \rightarrow 2^{\mathcal{S}}$  (or  $\mathcal{N}: 2^{\mathcal{S}} \rightarrow 2^{\mathcal{S}}$ ), which maps a current solution (or a subset of solutions) to a subset of *neighbouring solutions*. That is, for a given current solution  $s \in \mathcal{S}$ ,  $\mathcal{N}(s) \subseteq \mathcal{S}$  defines the set of solutions that can be reached from  $s$  by one application of the search operator.

Note that if the domain of  $\mathcal{N}$  is  $\mathcal{S}$ , then we are dealing with a *single-solution-based* metaheuristic; otherwise, the metaheuristic is *population-based*. The search process is typically guided by the evaluation of  $f$  over the neighbouring solutions, often combined with probabilistic rules or adaptive mechanisms to balance intensification (local search) and diversification (exploration of new regions), aiming to identify a solution  $s^* \in \mathcal{S}$  such that  $f(s^*)$  is as small as possible.

#### 2.3.2.2.1 Single-Solution-Based Metaheuristics

Single-solution-based metaheuristics, also called *local search* metaheuristics, foster learning throughout the search process by accumulating information about the solution landscape. They are typically designed to exploit local information efficiently while incorporating diversification strategies to escape local optima and explore new regions of the solution space. A well-known example is tabu search

(TS) which maintains a tabu list of recently explored solutions to discourage cycling and promote broader exploration beyond immediate improvements. TS has been applied by Ripplinger (2005), Spada et al. (2005), Pacheco and Marti (2006), Ngonyani et al. (2015), Semba and Mujuni (2019), and Díaz-Ramírez et al. (2022).

Another local search metaheuristic is simulated annealing (SA), inspired by the annealing process in metallurgy, where controlled cooling gradually leads a system toward a stable state. It allows occasional acceptance of worse solutions based on a probability that decreases over time, helping the algorithm escape local optima. SA has been applied in works such as those by Spada et al. (2005), Samadi-Dana et al. (2017), Semba and Mujuni (2019), Li and Chow (2021), Han et al. (2022), and Li et al. (2023).

Iterated local search (ILS) is a third single-solution-based metaheuristic that builds upon simple local search by repeatedly perturbing and refining a single solution. After applying local search to reach a local optimum, ILS introduces a controlled modification (referred to as *perturbation*) to escape it, then resumes local search. More detail on the implementation of ILS and its application to the SBRP will be provided in Section 4.4.

### 2.3.2.2 Population-Based Metaheuristics

Among the most widely used techniques for the route generation subproblem are evolutionary algorithms (EAs). Inspired by the principles of natural evolution, EAs evolve a population of solutions over successive generations. Through operations such as selection, crossover, and mutation, they combine and modify candidate solutions to progressively enhance their quality. Many studies refer to their proposed methods as genetic algorithms (GAs); however, in several cases, these approaches more generally fall within the broader class of EAs. Such EA-based approaches have been adopted by Thangiah and Nygard (1992), Sghaier et al. (2013), Minocha and Tripathi (2014), Kang et al. (2015), Chalkia et al. (2016), Kotoula et al. (2017), Oluwadare et al. (2018), Ümit and Kiliç (2019), Komijan et al. (2021), Ozmen and Sahin (2021), Hou et al. (2022b), and Pekel Özmen and Küçükdeniz (2024). An extension of EAs is memetic algorithms, which enhance the evolutionary framework by incorporating local search procedures to avoid premature convergence. Sales et al. (2018) proposed a memetic algorithm for the SBRP with varying fixed costs and vehicle capacities, employing a  $\lambda$ -interchange local search procedure (described in Section 4.3.1) to further improve solution quality.

Another population-based metaheuristic is ant colony optimization (ACO), inspired by the foraging behaviour of ants, which use pheromone trails to dis-

cover and reinforce short paths to food sources. In ACO, artificial ants construct solutions by making a sequence of probabilistic decisions. These decisions are guided by a so-called pheromone matrix, which encodes information about previously constructed solutions and may also incorporate problem-specific heuristics. When an ant constructs a high-quality solution, the components or decisions that make up that solution are reinforced in the pheromone matrix. Over time, the matrix guides the search to concentrate on high-quality regions of the solution space. In this sense, ACO might be considered a type of neighbourhood-based search, in that it uses a set of previously constructed solutions to probabilistically build new ones. ACO has been applied to the SBRP in works such as those by Arias-Rojas et al. (2012), Eldrandaly and Abdallah (2012), Euchi and Mraihi (2012), Huo et al. (2014), Yao et al. (2016), Mokhtari and Ghezavati (2018), Ren et al. (2019), and Sharma (2024).

In addition to GAs and ACO, other less common population-based metaheuristics have been explored for the SBRP. These include particle swarm optimization, which simulates the social behaviour of swarms to guide solutions through the search space using velocity and position updates (Fulin and Yueguang, 2012); estimation of distribution algorithms, which build probabilistic models to capture and exploit patterns in high-quality solutions (Pérez-Rodríguez and Hernández-Aguirre, 2016); harmony search, a technique inspired by the improvisation process of musicians seeking pleasing harmonies (Kim and Park, 2013); and scatter search, a method that systematically combines and refines a selected set of diverse solutions to generate improved candidates (Corberán et al., 2002). Although less prevalent in the literature, these approaches offer alternative frameworks that may be tailored to the specific characteristics and constraints of the routing subproblem.

### **2.3.2.2.3 Hybrid Metaheuristics and Metaheuristic Effectiveness**

Different metaheuristic approaches can be combined to create hybrid metaheuristics that leverage the strengths of each method while compensating for their individual weaknesses. This often enables more efficient handling and better scalability of complex optimization tasks. For example, Parvasi et al. (2017) proposed two hybrid approaches for the bus stop selection and route generation subproblems: one combining a GA, exact methods, and TS, and another combining SA, exact methods, and TS. Similarly, Shafahi et al. (2018a,b) employed a hybrid method integrating SA and TS. In other studies, Dang et al. (2019) and Hou et al. (2021) combined ILS with a set partitioning procedure.

Ideally, a metaheuristic should be validated against alternative approaches to build confidence in its effectiveness. This validation can be performed in one

of three ways: by comparing different versions of the same metaheuristic, by comparing it with other metaheuristics, or by benchmarking it against an integer programming model. Additionally, comparisons may involve manually designed routes or simple heuristics, such as those discussed in Section 4.2.1.

For instance, Mokhtari and Ghezavati (2018) benchmarked their ACO algorithm against exact methods and a GA, demonstrating the superiority of their proposed approach. Komijan et al. (2021) showed that their GA produced solutions that were, on average, only 0.62% worse than those from a MILP model, while reducing the number of buses from 44 to 40 and decreasing total travel time by 13% to 28%. Other benchmarking examples include Chen et al. (2015) (SA coupled with local search versus heuristics and exact methods), Pérez-Rodríguez and Hernández-Aguirre (2016) (estimation of distribution algorithm versus GA), Yao et al. (2016) (two ACO variants), Souza Lima et al. (2017) (four ILS variants), Sales et al. (2018) (memetic algorithm versus GA and exact methods), Shafahi et al. (2018a,b) (hybrid versus exact methods), Hou et al. (2022a) (ILP versus ACO and exact methods), and Li et al. (2023) (SA versus GA).

### 2.3.3 School Bell Time Adjustment and Route Scheduling

As noted in Section 1.1, the school bell time adjustment and route scheduling subproblems arise only in the context of multiple schools and are therefore only reviewed briefly here. While most studies treat school start and end times as fixed, some consider them as decision variables. This flexibility allows for optimizing schedules to serve more routes sequentially with the same bus, thereby reducing the number of buses required. For instance, the school bell time adjustment model proposed by Bertsimas et al. (2019), along with the proposed bi-objective routing decomposition algorithm, led to annual savings of \$5 million in Boston by reducing the fleet by 50 buses (31%), while maintaining the same level of service quality.

Desrosiers et al. (1981) employed a CG approach to determine school start and end times. Fügenschuh (2009) focused on scheduling school start times while allowing student transshipments between routes, developing a MIP model solved using a B&C method. Banerjee and Smilowitz (2019) simultaneously modelled the school bell time adjustment and route scheduling subproblems, incorporating equity considerations in changes to school start times. Miranda et al. (2021) proposed three different bell time adjustment strategies, each embedded within different phases of a memetic algorithm.

While school bell time adjustments are relatively uncommon, route scheduling, typically performed after route generation, has been more extensively studied

in multischool settings. Newton and Thomas (1974), for example, developed a multischool model that assumes schools begin at different time periods. Similarly, Bodin and Berman (1979) partitioned schools' time windows using defined time periods. However, their approach may be invalid if school bell times or time periods overlap. Braca et al. (1997) jointly formulated the route generation and scheduling subproblems as a set partitioning problem. More recently, Spada et al. (2005) sequenced schools by start time and attempted to merge routes that were constructed greedily.

Mandujano et al. (2012) suggested a two-stage framework: the first stage determines which schools to close, expand, or maintain and assigns students to each school; the second stage identifies the most efficient bus routes and allocates courses across school shifts, leveraging students' spatial distribution to minimize teacher requirements. Chen et al. (2015) addressed the route scheduling subproblem using two methods: a MIP model that minimizes fleet size and total travel distance, and a metaheuristic combining SA with local search. Shafahi et al. (2017) proposed a school-decomposition heuristic that embeds trip compatibility considerations into the routing phase, demonstrating a reduction of up to 13% in fleet size over traditional routing methods.

In recent years, Babaei and Rajabi-Bahaabadi (2019) have proposed a bi-level heuristic that simultaneously addresses route generation and scheduling, allowing student transfers within a single trip. Unlike prior work focused solely on dispatch times, their approach extends scheduling to include pickup times at all stops. Köksal Ahmed et al. (2021) introduced a reinforcement learning-enabled GA for integrated bus route generation and scheduling. In this method, reinforcement learning dynamically guides the parameter selection of the GA, enabling it to achieve near-optimal solutions more efficiently and with fewer iterations.

For more information on the school bell time adjustment and route scheduling subproblems, the reader is referred to works such as those by Kim et al. (2012), Park et al. (2012), Bögl et al. (2015), Campbell et al. (2015), Yao et al. (2016), Souza Lima et al. (2017), Miranda et al. (2018), Shafahi et al. (2018b), Wang and Haghani (2020), and Díaz-Ramírez et al. (2022).

## 2.4 Summary

This chapter has presented a structured review of the literature relevant to the SBRP, with the intention of establishing a solid foundation for the research that follows. It began with a brief overview of the VRP, providing context for how the SBRP fits within the broader family of routing problems. This background is essential, as many of the assumptions, formulations, and solution approaches

in SBRP research are influenced by developments in the VRP domain.

The focus then shifted to the SBRP, exploring the core characteristics that define the problem in existing studies. These include underlying assumptions related to school policies and the operational context, as well as the objectives that guide solution strategies – most commonly minimizing travel time, cost, or the number of routes, while also acknowledging less frequent goals related to effectiveness and equity in service provision. A discussion of commonly applied constraints was also included, with examples drawn from a wide range of studies to illustrate how different combinations of these constraints are handled in practice.

Finally, the chapter reviewed the main solution methods proposed for tackling the various subproblems that comprise the SBRP. Approaches were categorized based on their methodological frameworks, ranging from exact optimization techniques to metaheuristics. The review highlights the diversity of problem formulations and methodological trends found in the literature, reflecting the complexity and practical relevance of the SBRP.

In the course of this review, we identified notable research gaps that informed the direction of this thesis. Specifically, as will be further evidenced in the subsequent chapters, contributions in the vehicle routing literature that consider both split deliveries and stochastic travel times are rare and, to our knowledge, entirely absent within the scope of the SBRP. Furthermore, a majority of existing studies addressing the S-SBRP are limited to small-scale instances, typically involving a relatively low number of bus stops. Additionally, some of these works rely on overly simplistic or artificially generated data that fail to capture the complexity of real-world scenarios. It is these gaps that this research aimed to address.

# Chapter 3

## The Single-School Bus Routing Problem

### 3.1 Introduction

In this chapter, we define and mathematically formulate our S-SBRP, where a set of routes is constructed independently for each school. We focus on the S-SBRP since mixed loads are typically not permitted in the locations under consideration. Our research addresses the morning problem, where students are picked up from designated stops and transported to school. The problem constitutes a special case of the OVRP, as in the studied locations, routes are contracted out through public tenders, and buses are housed in private company garages rather than at the schools.

Our S-SBRP covers the first three subproblems described in Section 1.1, namely data preparation, bus stop selection, and route generation. The former subproblem is tackled in Section 3.2.1, where we present a network representation of our problem and describe the input parameters and assumptions related to this representation. The bus stop selection and route generation subproblems are intertwined, with their constraints and objectives addressed in Sections 3.2.2 and 3.2.3, respectively. We highlight that, without the bus stop selection subproblem, our problem is equivalent to the open SDVRP with capacity and time constraints. Therefore, any material presented in this research related to the route generation subproblem can also be applied to instances of the open capacitated and time-constrained SDVRP.

A MIP formulation of our S-SBRP is presented in Section 3.2.4. This formulation allows us to get an insight into the computational effort needed to solve different instances of our problem. It can also be used to obtain feasible solutions to problem instances that are either small (for example, a few tens of potential

bus stop locations) or have a simple network structure (for example, clustered student home addresses). This is considered in Section 3.4, where we present some solutions of the MIP model. This section also includes an explanation of the data generation process, a justification of the parameter selections, and a summary of the problem instances considered in this research. Before this section, we establish the computational complexity of the S-SBRP in Section 3.3.

## 3.2 Modelling the S-SBRP

In this section, we define our S-SBRP in terms of its characteristics, constraints and objectives in Sections 3.2.1 to 3.2.3, respectively. We then move on to formulating the problem mathematically in Section 3.2.4.

### 3.2.1 Definitions and Assumptions

At the outset, the S-SBRP can be represented via two vertex sets  $V_1$  and  $V_2$ . The vertex set  $V_1$  consists of one school,  $v_0$ , and  $n$  potential bus stops  $v_1, v_2, \dots, v_n$ . The major consideration when selecting the locations of the potential bus stops is safety. Bus stops should be away from roadways with high traffic and/or speed limits, road intersections, railroad crossings, and areas with other safety threats such as high crime and dangerous wildlife. They should have sufficient space for students to wait safely out of the roadway and for buses to pull in and out. Moreover, visibility is essential from multiple viewpoints; visibility of incoming traffic by the students, visibility of the students by approaching drivers, and visibility of the students and surrounding environment by the bus drivers. Ideally, bus stops should also be connected to well-constructed and well-maintained walking paths. Bus stops may be positioned at central locations such as town squares, or next to landmarks such as churches and community centres. Such locations are easy to access and help to reduce the number of stops along routes and, in turn, the students' riding times. In the case of a door-to-door service, the bus stop locations correspond to the student home addresses.

The vertex set  $V_2$  comprises all student home addresses. As remarked earlier, this set coincides with  $V_1$  in the case of a door-to-door service. Here, it is assumed that this is not the case. Each address  $w \in V_2$  must be at a walking distance of at least  $m_e$  units from the school  $v_0$ . The parameter  $m_e$  denotes the minimum walking distance that students should live from the school to be eligible for school transport service. This parameter may differ depending on the local policies in place, taking into account the students' age or grade level as well as any community variations. In Malta, for example, school transport is provided

free of charge to all kindergarten, primary, and secondary school students residing at least 1 km from the school of their locality of residence (or of their respective college in the case of secondary school students). In the UK, statutory guidance specifies a walking distance of 2 miles for students under the age of 8 attending their nearest suitable school, and 3 miles for those aged 8 to 16. Note that the walking distance between a student's home and the school is typically measured by the shortest safe walking route. Walking conditions in rural environments are prone to be more hazardous than those in urban areas due to the vast expanse of undeveloped land, the absence of paved sidewalks, and the presence of roadways with high speed limits. If a student lives within  $m_e$  walking distance from the school but the walking environment poses safety risks, then the student is commonly still entitled to school transport. The same often applies if the student has special educational needs, mobility problems, or comes from a family with a low income. Moreover, when a student is ordinarily resident at more than one eligible address, a pick-up and drop-off schedule may be agreed upon with the parents/guardians. For the sake of simplicity and since such cases are rare and individualized, they will be excluded from this research. In addition, each address  $w \in V_2$  is assumed to have a corresponding number  $s(w)$  of students residing at that address who require school transport.

Associated with vertex sets  $V_1$  and  $V_2$  are two edge sets  $E_1$  and  $E_2$ , respectively.  $E_1$  is the set of all  $n(n+1)$  arcs  $(u, v)$  where  $u, v \in V_1$  and  $u \neq v$ . This edge set includes an arc between the school and each bus stop (in both directions) and between each pair of distinct bus stops, yielding a complete directed graph  $(V_1, E_1)$ . Each arc  $(u, v) \in E_1$  is weighted by the driving time  $t(u, v)$  and corresponding driving distance  $d(u, v)$  from vertex  $u$  to vertex  $v$ . In the road network, there may exist several paths that lead from vertex  $u$  to vertex  $v$ . Here, the path that satisfies certain compliance criteria and has the shortest driving time is considered. For instance, transport administrators may want to minimize or avoid certain road types or driving manoeuvres such as highways, toll roads, and U-turns. Web mapping platforms such as Bing Maps and Google Maps allow such criteria to be set when establishing driving times and distances. Note that driving times and distances may be asymmetric, i.e., for arc  $(u, v) \in E_1$ , we can have  $t(u, v) \neq t(v, u)$  or  $d(u, v) \neq d(v, u)$ , or both.

Meanwhile,  $E_2$  is the set  $\{\{w, v\} : w \in V_2 \wedge v \in (V_1 \setminus \{v_0\}) \wedge D(w, v) \leq m_w\}$ , where  $D(w, v)$  denotes the shortest walking distance from address  $w$  to bus stop  $v$ . Again, safety measures can be considered when determining the shortest walking distances. The parameter  $m_w$  denotes the maximum distance that students are expected to walk to a bus stop. For instance, a walk of 1.6 km is often deemed reasonable in Wales. Here, the undirected bipartite graph  $(V_2, V_1 \setminus \{v_0\}, E_2)$  is

assumed to have no isolated vertices. To be otherwise implies that either an address has no bus stop within walking distance  $m_w$  (and a new bus stop must hence be added to  $V_1$ ), or a bus stop has no address within walking distance  $m_w$  (and can consequently be deleted from  $V_1$ ). Moreover, a bus stop  $v \in V_1 \setminus \{v_0\}$  for which there exists an address  $w \in V_2$  with the single incident edge  $\{w, v\}$  shall be called a *compulsory stop*. This is because such a stop  $v$  is the only stop within walking distance  $m_w$  to students living at address  $w$  and must, therefore, always be selected to accommodate these students.

### 3.2.2 Constraints

A solution to the S-SBRP is given by a set of routes  $\mathcal{R} = \{R_1, R_2, \dots, R_k\}$ , where  $k$  denotes the total number of routes needed to serve all students. Each route  $R \in \mathcal{R}$  uses one bus that must have adequate seating capacity for all students boarding that route. The selection of the optimal seating capacity for each route  $R$  will be discussed in Section 3.2.3, however, at this point, it will be assumed that the maximum available seating capacity is defined by the constant  $C_{\max}$ .

In a solution, each bus (route) is assumed to start at a bus stop, visit a subset of other bus stops, and terminate at the school  $v_0$ . This type of route design – used, for example, by Bennett and Gazis (1972) and Lewis and Smith-Miles (2018) – is referred to as a *shoestring route*, as the bus starts away from the school and picks up students along the way before ending the route at the school. An alternative approach is the *circular route* design, employed by Newton and Thomas (1969) and Dulac et al. (1980), in which the bus departs from the school, visits a subset of bus stops in a circular pattern, and then returns to the school, where it is parked. The latter design offers the advantage of being easily reversed in the afternoon, thereby balancing students' riding times since those picked up first in the morning are dropped off first in the afternoon. However, the circular design is relatively uncommon, as buses are typically parked at depots near the drivers' homes. Consequently, it is often more practical for drivers to begin routes by picking up students directly rather than commuting to the school to start the route. In some settings (e.g., Malta), the service provider may also bear the cost of travel from the depot to the first pick-up point, further discouraging the circular design.

Certain constraints should be satisfied by a solution  $\mathcal{R}$  for it to be considered a feasible solution. These are mathematically expressed as follows:

$$\bigcup_{i=1}^k R_i = V'_1 \tag{3.1}$$

$$\exists v \in V_1' \mid \{w, v\} \in E_2 \qquad \forall w \in V_2 \qquad (3.2)$$

$$s(R) \leq C_{\max} \qquad \forall R \in \mathcal{R} \qquad (3.3)$$

$$t(R) \leq m_t \qquad \forall R \in \mathcal{R}. \qquad (3.4)$$

In Constraint (3.1), the subset of bus stops visited by at least one route in  $\mathcal{R}$  is denoted by  $V_1' \subseteq V_1 \setminus \{v_0\}$ . Note that although each bus terminates at the school  $v_0$ , the latter is excluded from the vertex union to simplify the notation. The set of visited bus stops  $V_1'$  should *cover* each address in  $V_2$  at least once. This means that each address  $w \in V_2$  must be within walking distance  $m_w$  of at least one bus stop in  $V_1'$ , as shown in Constraints (3.2). Such a covering shall be referred to as a *complete covering* of  $V_2$ , whereas a covering that does not satisfy this property shall be referred to as an *incomplete covering* of  $V_2$ .

Constraints (3.3) state that the total number  $s(R)$  of students boarding the bus on each route  $R \in \mathcal{R}$  must not exceed the maximum bus capacity  $C_{\max}$ . It is assumed here that administrators first construct the bus routes and then put them out to public tender. Therefore, during the route generation phase, the bus fleet is assumed to have differing vehicle capacities, each of which is available in an unlimited quantity.

Finally, Constraints (3.4) state that the total journey time  $t(R)$  of each route  $R \in \mathcal{R}$  must not exceed the maximum journey time  $m_t$ . Similar to  $m_e$  and  $m_w$ , this parameter may differ based on the students' age and school location. In Wales, for example, a maximum 45-minute and 60-minute journey time is recommended for primary and secondary school students, respectively. This follows the Learner Travel Operational Guidance issued by the Welsh government and promotes the positive well-being and safety of students, particularly younger ones. An example of an S-SBRP solution is presented in Figure 3.1.

The total journey time  $t(R)$  of route  $R \in \mathcal{R}$  is assumed to be composed of two components: the total *bus travel time* and the total *bus dwell time*. Each dwell time within a route captures the time spent servicing a designated bus stop; i.e., the time spent decelerating the bus, opening the doors, boarding the students, closing the doors, and accelerating to merge back into traffic. Here, the dwell time at stop  $v$  in route  $R$  is estimated using the linear function  $d(v, R) = d_1 + d_2 s(v, R)$ , where  $s(v, R)$  represents the number of boarding students at stop  $v$  onto route  $R$ ,  $d_2$  represents the boarding time per student, and  $d_1$  is a parameter which accounts for the remaining service time. An analogous function that has an additional term catering for alighting passengers is used in the Transit Capacity and Quality of Service Manual (TCQSM) (Transportation Research Board and National Academies of Sciences, Engineering, and Medicine, 2013) and Highway Capacity Manual (HCM) (Transportation Research Board and National Academies of

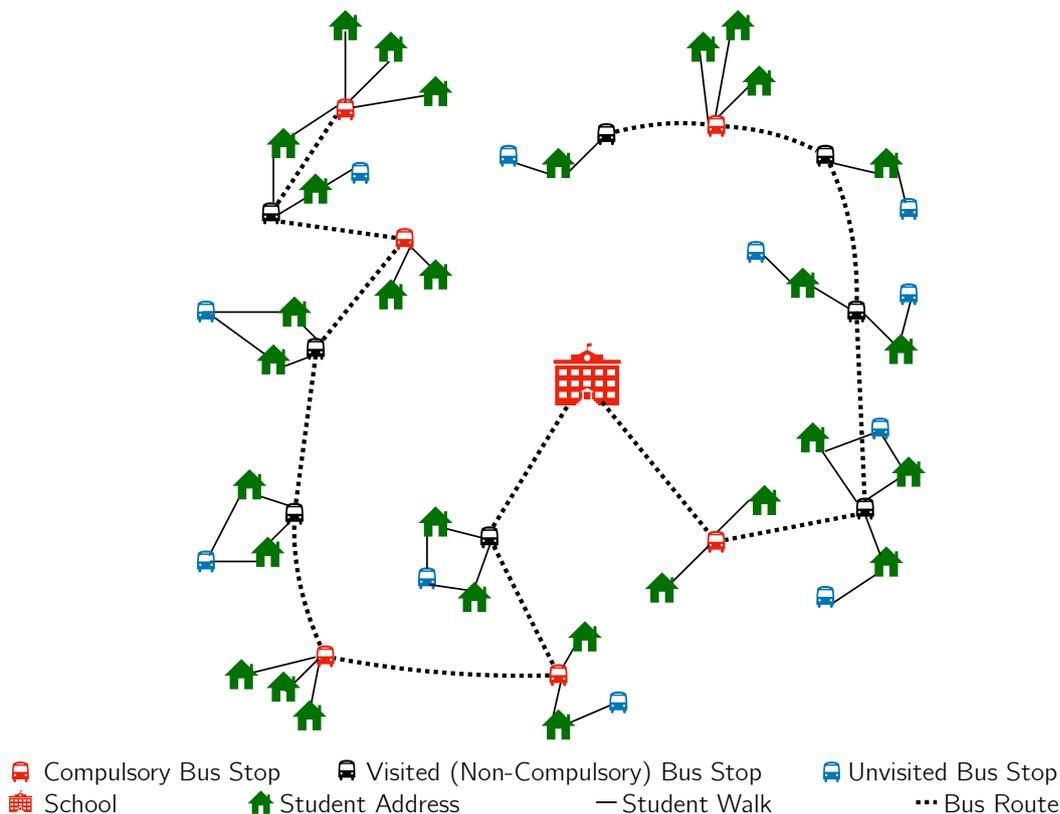


Figure 3.1: An S-SBRP solution with  $k = 2$ ,  $|V_1| = 26$ ,  $|V_1'| = 14$  and  $|V_2| = 30$ .

Sciences, Engineering, and Medicine, 2022).

Several studies analyse and predict bus dwell times based on factors such as the number of boarding/alighting passengers, time of the day, fare collection method, crowding level at the bus stop, number of bus doors and their widths, and bus platform height. A review of research conducted on this topic can be found in the work of Rashidi et al. (2023), who classify bus dwell time estimation models in three categories, namely, regression-based models, probabilistic and time series models, and artificial intelligence-based models. They found that the majority of scholars have used regression-based models and that the number of boarding/alighting passengers has been proven to be the most significant factor.

We now discuss some study findings related to parameters  $d_1$  and  $d_2$ . One of the earliest studies is that by Guenther and Sinha (1983) who propose a simple linear-log regression model with the number of boarding/alighting passengers as the factor. They report, based on data from Lafayette, Indiana, a bus dwell time of 10 to 20 seconds plus 3 to 5 seconds per boarding/alighting passenger. Aash-tiani and Iravani (2002) compare seven different dwell time regression models for the city of Tehran, Iran, that include four factors; these being, the number of boarding passengers, number of alighting passengers, number of bus doors, and load factor (number of passengers in the bus before visiting the bus stop

divided by the vehicle capacity). They fix  $d_1 = 12$  seconds and estimate from the model with the best overall goodness-of-fit index ( $R^2$ ), that  $d_2$  lies between 4.19 (for a standard bus) and 4.55 (for a double-decker bus) seconds. Furthermore, Dunlop (2009) propose a multiple linear regression model with the number of boarding passengers and number of alighting passengers as factors. Using data from Wellington, New Zealand,  $d_1$  and  $d_2$  were estimated to be 7.5 and 5.9 seconds, respectively. Dai et al. (2018) present a probabilistic model to estimate bus dwell time based on the number of buses queuing at the bus stop, number of boarding/alighting passengers, and the time each bus needs to merge back into traffic. The estimates they provide for the city of Hangzhou, China, are  $d_1 = 7.2$  and  $d_2 = 3.14$  seconds.

In this thesis,  $d_1$  and  $d_2$  are taken to be 15 and 5 seconds, respectively. The reason for the former selection is that the passengers are minors in our case and, therefore, the driver should be extra careful when pulling the bus in or out of the bus stop, and opening or closing the doors. The walking time of students from the bus waiting area to the bus is also considered in this parameter. The selection of  $d_2$  follows the study findings discussed above. Note that although Tirachini (2013) argues that young passengers board faster than senior passengers, 5 seconds are allocated to each student. This accounts for the fact that once a student gets onto the bus, they typically have to tap their bus pass or fob against the bus reader.

As a result, given a route  $R = (v_1, v_2, \dots, v_j, v_0)$ , the route journey time  $t(R)$  (in seconds) is calculated as

$$t(R) = \left( \sum_{i=1}^{j-1} t(v_i, v_{i+1}) + t(v_j, v_0) \right) + \left( \sum_{i=1}^j (15 + 5s(v_i, R)) \right), \quad (3.5)$$

where the first parenthesis gives the total bus travel time and the second parenthesis gives the total bus dwell time.

Note that visited bus stops in  $V'_1$  can occur in more than one route in  $\mathcal{R}$ . For example, there may not be enough spare capacity in a bus to serve all students waiting at bus stop  $v \in V'_1$ . In that case, bus stop  $v$  must be visited by multiple buses. Here, a bus stop occurring on two or more routes in a solution is called a *multistop*. Recall from Chapter 2 that this characteristic leads to the SDVRP variant, which was formally studied by Dror and Trudeau (1989, 1990). Compared to the CVRP, they found that the SDVRP yields significant savings in terms of the total distance and number of vehicles used, especially when the customer demands are high relative to the vehicle capacity  $C$  (homogeneous fleet assumed). For example, for problem instances with customer demands between  $0.3C$  and  $0.7C$ ,

the average percentage gains in distance were found to be 9.01, 9.82, and 9.36 for 75, 115, and 150 customers, respectively. The corresponding average numbers of vehicles saved were 8.37, 13.07, and 15.27. A comprehensive review of the SDVRP can be found in the work of Archetti and Speranza (2012). In the context of the S-SBRP, Chen et al. (2016) were the first to consider splitting of demands at bus stops. They formulated a bi-objective problem, minimizing the number of buses and total travelling distance, and proposed a metaheuristic algorithm for this problem. The algorithm incorporates constructive and improvement heuristics, a ruin and recreate mechanism, and an SA rule for accepting worse neighbouring solutions. Their computational experiments confirm the efficiency of the proposed algorithm and showcase the savings brought about by splitting demands.

In our case, it is assumed that each student at a multistop is only allowed to board one specific bus serving that stop since, otherwise, a bus stopping there may be too full to serve subsequent stops in its route. In practice, this rule can be implemented by giving a bus pass to all students. Generally, students living at the same address (particularly siblings) are preferably assigned to the same bus so that they can walk to the bus stop together and board the bus simultaneously. However, there may be cases where students from the same address are assigned to different buses. For example, if there are tens of students waiting at a multistop, these can be split by age to avoid large age gaps between students on the same bus.

### 3.2.3 Objectives

The optimization objective of this research is hierarchical, with the primary goal being to minimize the number  $k$  of routes included in a solution. This is achieved by attempting to construct feasible solutions that use the minimal number  $\underline{k} = \lceil \sum_{w \in V_2} s(w) / C_{\max} \rceil$  of buses needed to serve all students. Under the assumption of having enough buses of capacity  $C_{\max}$  to cater for all students, a solution satisfying Constraints (3.1)-(3.3) and meeting this lower bound of  $k$  is always guaranteed since multistops are allowed. However, any one of the routes could potentially violate Constraint (3.4) relating to the students' maximum riding time. Thus, there may be cases where more than  $\underline{k}$  routes are required. The secondary goal is to minimize route journey times. This ensures that the service is both efficient from the government's perspective and effective from the students' and guardians' perspectives.

Here, a candidate solution  $\mathcal{R}$  is evaluated according to the cost function

$$f(\mathcal{R}) = \sum_{R \in \mathcal{R}} t'(R), \quad (3.6)$$

where

$$t'(R) = \begin{cases} t(R) & \text{if } t(R) \leq m_t, \\ m_t + m_t(1 + t(R) - m_t) & \text{otherwise.} \end{cases} \quad (3.7)$$

This means that if route  $R \in \mathcal{R}$  satisfies Constraint (3.4), then its journey time  $t(R)$  is unaltered in the total route journey time. On the other hand, if the journey time  $t(R)$  exceeds  $m_t$ , then this journey time is scaled up heavily to ensure that Constraint (3.4) is treated as a hard constraint. For example, if  $t(R) = 2701$  seconds and  $m_t = 2700$  seconds (45 minutes), then  $t'(R) = 8100$  seconds. Observe that the addition of the value 1 in the second case of Equation (3.7) guarantees that two routes both with journey time at most  $m_t$  are always preferred over one route with journey time exceeding  $m_t$ . Indeed, without this term, a solution with one route of 2701 seconds would have the same cost as a solution with two routes of 2700 seconds each, for  $m_t = 2700$  seconds.

In addition to minimizing the number of routes and the cost function (3.6), two measures that encourage equity among students are also taken into account. First, students are assigned to their closest stop in a selected subset of bus stops. This avoids situations where students living close to each other have unbalanced walking distances. Note that minimizing students' walking distances is not incorporated directly into Equation (3.6). This extension is worth considering as it improves service effectiveness, at the expense of having longer routes due to more visited bus stops. It is not tackled here because governments are mainly interested in ensuring that the service is available to all eligible students and efficient. What is crucial to them is compliance with the guidelines relating to the students' maximum walking distance and maximum riding time. Second, if multiple solutions have the same number of routes and minimum cost, then one with the smallest time discrepancy between the shortest and longest routes is preferred. This helps to balance students' riding times as much as possible.

Once a solution  $\mathcal{R}$  is achieved after the optimization process, one final task is to assign buses to the routes such that the total number of spare seats is minimized. Such an assignment is desirable because larger buses have higher maintenance and operating costs (for example, they consume more fuel). Furthermore, they usually have worse manoeuvrability, especially on high-traffic and narrow roads (which often occur in places like Malta). In our case, each route is thus assigned a bus with the smallest available seating capacity required to perform that route. Recall from Section 3.2.2 that each vehicle capacity under consideration is assumed to be available in an unlimited quantity. Therefore, this strategy is feasible. If buses of each capacity were to be limited in number, then the problem of assigning buses to routes could be modelled as a *minimum-weight*

*maximum-cardinality bipartite matching problem.* The bipartition would consist of a set of  $k$  vertices representing the routes in the solution and a set of  $b \geq k$  vertices representing all available buses. A route vertex would then be linked by an edge to a bus vertex if and only if the route load was at most the bus capacity. Such an edge would be weighted by the difference between the bus capacity and the route load, corresponding to the number of spare seats. The aim would then be to find a matching of size  $k$  with minimum total weight. An example of an optimal matching in this regard is presented in Figure 3.2. For large  $k$  and  $b$ , such an optimal matching can be found using the Hungarian algorithm, which has worst-case run-time complexity  $\mathcal{O}(b^3)^a$ .

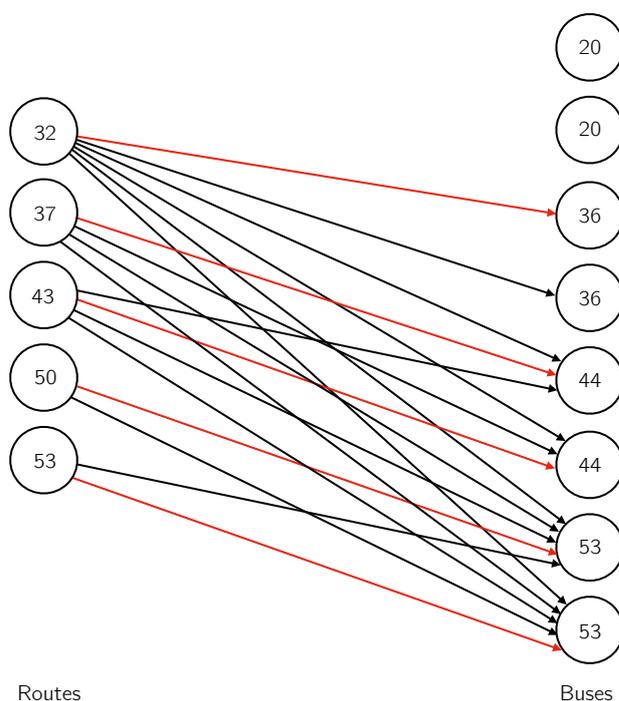


Figure 3.2: A bipartition with  $k = 5$  and  $b = 8$ . The node labels represent the route loads and bus capacities. An example of a minimum-weight maximum-cardinality matching is displayed in red. Note that edge weights have been excluded for clarity.

### 3.2.4 Mathematical Formulation of the S-SBRP

The S-SBRP described in Sections 3.2.1 to 3.2.3 can be formulated as a MIP which we present in this section. For simplicity, routes produced by the following MIP have a circular design, starting and ending at the school. The arc from the school to the first visited bus stop in each route is then excluded. This is made

<sup>a</sup>Function  $f$  has order of function  $g$ , denoted as  $f(x) = \mathcal{O}(g(x))$ , if there exists positive constants  $c$  and  $x_0$  such that  $f(x) \leq cg(x) \forall x \geq x_0$ . In essence,  $f(x)$  does not grow faster than  $g(x)$ .

possible by assuming that the driving time from the school to any potential bus stop is zero, as suggested by Letchford et al. (2007) who proved that any closed VRP can be transformed into an equivalent OVRP.

In addition to this, the MIP requires the number  $k$  of routes to be given as an input parameter. The idea is to start with the lower bound  $\underline{k}$  specified in Section 3.2.3 and increment  $k$  by one whenever no feasible solution is found within a certain time limit.

Let  $K = \{1, 2, \dots, k\}$ . The decision variables in the MIP are as follows:

- $x_{uvi}$  is a binary variable equal to 1 if route  $i \in K$  travels from  $u$  to  $v$ , where  $u, v \in V_1$  and  $u \neq v$ , and 0 otherwise;
- $y_{vi}$  is a binary variable equal to 1 if route  $i \in K$  visits  $v \in V_1$ , and 0 otherwise;
- $z_{wv}$  is a binary variable equal to 1 if students at address  $w \in V_2$  walk to bus stop  $v \in V_1 \setminus \{v_0\}$ , and 0 otherwise;
- $s_{vi}$  is an integer variable indicating the number of students boarding route  $i \in K$  at bus stop  $v \in V_1 \setminus \{v_0\}$ ;
- $l_{vi}$  is an integer variable indicating the total load of route  $i \in K$  just after visiting bus stop  $v \in V_1 \setminus \{v_0\}$  (note that  $l_{0i} = 0 \forall i \in K$  where necessary);
- $t_i$  is a continuous variable giving the total journey time of route  $i \in K$ .

The objective of the S-SBRP is to

$$\text{minimize } \sum_{i \in K} t_i$$

subject to

$$\sum_{u \in V_1 \setminus \{v\}} x_{uvi} = y_{vi} \quad \forall v \in V_1, i \in K \quad (3.8)$$

$$\sum_{u \in V_1 \setminus \{v\}} x_{vui} = y_{vi} \quad \forall v \in V_1, i \in K \quad (3.9)$$

$$\sum_{i \in K} y_{v_0 i} = k \quad (3.10)$$

$$\sum_{i \in K} y_{vi} \geq z_{wv} \quad \forall v \in V_1 \setminus \{v_0\}, w \in V_2 \quad (3.11)$$

$$\sum_{\substack{v \in V_1 \setminus \{v_0\} \\ D(w,v) > m_w}} z_{wv} = 0 \quad \forall w \in V_2 \quad (3.12)$$

$$\sum_{\substack{v \in V_1 \setminus \{v_0\} \\ D(w,v) \leq m_w}} z_{wv} = 1 \quad \forall w \in V_2 \quad (3.13)$$

$$\frac{1}{k} \left( \sum_{i \in K} y_{ui} \right) z_{wv} \leq z_{wu} \quad \forall w \in V_2, u, v \in V_1 \setminus \{v_0\} \mid D(w, u) < D(w, v) \leq m_w \quad (3.14)$$

$$\sum_{w \in V_2} s(w) z_{wv} - \sum_{i \in K} s_{vi} = 0 \quad \forall v \in V_1 \setminus \{v_0\} \quad (3.15)$$

$$y_{vi} \leq s_{vi} \quad \forall v \in V_1 \setminus \{v_0\}, i \in K \quad (3.16)$$

$$C_{\max} y_{vi} \geq s_{vi} \quad \forall v \in V_1 \setminus \{v_0\}, i \in K \quad (3.17)$$

$$l_{ui} + s_{vi} - C_{\max}(1 - x_{uvi}) \leq l_{vi} \quad \forall u \in V_1, v \in V_1 \setminus \{v_0, u\}, i \in K \quad (3.18)$$

$$\sum_{\substack{u, v \in V_1 \\ u \neq v}} t(u, v) x_{uvi} + \sum_{v \in V_1 \setminus \{v_0\}} (15y_{vi} + 5s_{vi}) = t_i \quad \forall i \in K \quad (3.19)$$

$$x_{uvi} \in \{0, 1\} \quad \forall u \in V_1, v \in V_1 \setminus \{u\}, i \in K \quad (3.20)$$

$$y_{vi} \in \{0, 1\} \quad \forall v \in V_1, i \in K \quad (3.21)$$

$$z_{wv} \in \{0, 1\} \quad \forall v \in V_1 \setminus \{v_0\}, w \in V_2 \quad (3.22)$$

$$l_{vi} \in \{0, 1, \dots, C_{\max}\} \quad \forall v \in V_1 \setminus \{v_0\}, i \in K \quad (3.23)$$

$$s_{vi} \in \{0, 1, \dots, C_{\max}\} \quad \forall v \in V_1 \setminus \{v_0\}, i \in K \quad (3.24)$$

$$t_i \in [0, m_t] \quad \forall i \in K. \quad (3.25)$$

In the above, Constraints (3.8) to (3.10) relate to stop and school visits. Specifically, (3.8) and (3.9) ensure that if route  $i \in K$  visits  $v \in V_1$ , then it should enter and leave  $v$  exactly once. Constraints (3.10) guarantee that the number of routes leaving the school  $v_0$  is exactly  $k$ .

Next, Constraints (3.11) to (3.15) relate to student walks and pickups. Constraints (3.11) ensure that no student walks to an unvisited stop. Observe that the sum on the left-hand side gives the number of routes visiting bus stop  $v \in V_1 \setminus \{v_0\}$ . Meanwhile, Constraints (3.12) and (3.13) guarantee that students living at each address  $w \in V_2$  walk to exactly one bus stop within walking distance  $m_w$ . Constraints (3.14) are associated with the assignment of students to their closest bus stop from the visited ones. To illustrate, imagine that the students living at address  $w \in V_2$  have two alternative bus stops they can walk to;  $u$  with  $D(w, u) = 0.5$ , and  $v$  with  $D(w, v) = 0.7$ . By (3.14), if bus stop  $u$  is visited by at least one route, then  $\frac{1}{k} z_{wv} \leq z_{wu}$ . Thus, this inequality excludes solutions with  $z_{wv} = 1$  and  $z_{wu} = 0$  from the solution space. If, on the other hand, bus stop  $u$  is not visited, then no restrictions are imposed on  $z_{wv}$ . Moving on to Constraints (3.15), one can note that these constraints ensure that the total number of students boarding at stop  $v \in V_1 \setminus \{v_0\}$  is equal to the total number of students walking to that stop.

Constraints (3.16) and (3.17) relate to student boardings. Together, they force the number of students boarding route  $i \in K$  at bus stop  $v \in V_1 \setminus \{v_0\}$  to be zero if route  $i$  does not visit stop  $v$ . If route  $i$  visits stop  $v$ , then (3.16) updates the lower bound on the number of boarding students from zero to one.

In addition, Constraints (3.18) relate to route loads and also serve as subtour elimination constraints as proposed by Miller et al. (1960). These constraints, known as the Miller-Tucker-Zemlin (MTZ) constraints, guarantee that no route contains a subtour disconnected from the school  $v_0$ , and that each route load increases by the number of students boarding the bus on that route. If route  $i \in K$  goes from  $u \in V_1$  to stop  $v \in V_1 \setminus \{v_0, u\}$  (i.e.,  $x_{uv} = 1$ ), then the load of route  $i$  immediately after visiting stop  $v$  must be at least the load of route  $i$  just after visiting  $u$  plus the number of students boarding route  $i$  at stop  $v$ . Otherwise, the upper bound  $C_{\max}$  is imposed on  $l_{ui} + s_{vi} - l_{vi}$ , which is always satisfied since  $l_{ui} \leq C_{\max}$  and  $s_{vi} \leq l_{vi}$ . Figure 3.3 demonstrates how a subtour results in a violation of an MTZ constraint.

Constraints (3.19) calculate the total journey time of each route  $i \in K$  in accordance with Equation (3.5). Finally, Constraints (3.20) to (3.25) state the domains of the decision variables.

We now highlight five important points about our proposed MIP. First, our MIP is quadratically constrained due to Constraints (3.14), which include products of binary variables. A product of binary variables  $\eta_{uiwv} := y_{ui} \cdot z_{wv}$  can be

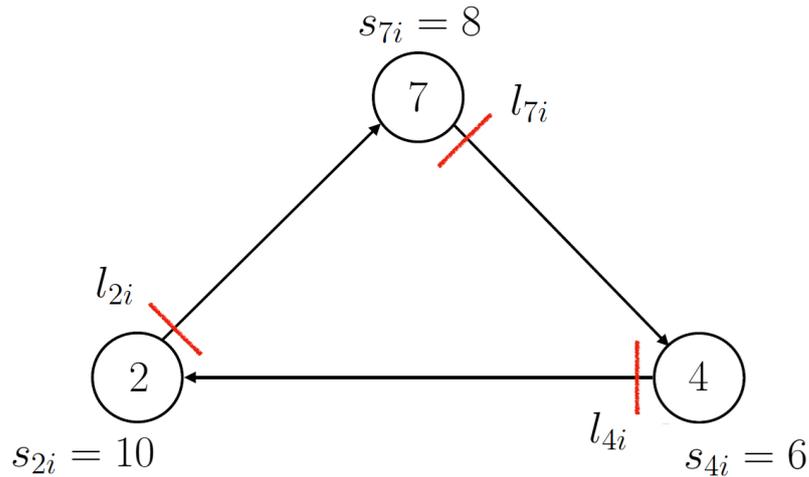


Figure 3.3: A subtour demonstrating a violation of an MTZ constraint. Without loss of generality, setting  $l_{7i} = 8$ ,  $l_{4i} = 14$ , and  $l_{2i} = 24$  results in a violation of the MTZ constraint for arc  $(2, 7)$ .

linearized as follows (Asghari et al., 2022):

$$\begin{aligned}\eta_{uiwv} &\leq y_{ui}, \\ \eta_{uiwv} &\leq z_{wv}, \\ \eta_{uiwv} &\geq y_{ui} + z_{wv} - 1, \\ \eta_{uiwv} &\in \{0, 1\}.\end{aligned}$$

Decision variable  $\eta_{uiwv}$  should be binary-valued since it is a product of binary variables. The first two inequalities guarantee that  $\eta_{uiwv}$  is zero if either  $y_{ui}$  or  $z_{wv}$  is zero. Also, the third inequality ensures that  $\eta_{uiwv}$  takes value one if both  $y_{ui}$  and  $z_{wv}$  are one. In our case, linearizing the quadratic constraints was unnecessary, as the solver used for the computational experiments presented in Section 3.4 could handle quadratic constraints directly.

Second, in conjunction with Constraints (3.18), one might also choose to include reverse inequalities to impose a lower bound on  $l_{ui} + s_{vi} - l_{vi}$ . Along the lines suggested by Kek et al. (2008), these would take the form

$$l_{ui} + s_{vi} + C_{\max}(1 - x_{uvi}) \geq l_{vi} \quad \forall u \in V_1, v \in V_1 \setminus \{v_0, u\}, i \in K. \quad (3.26)$$

If route  $i \in K$  goes from  $u \in V_1$  to stop  $v \in V_1 \setminus \{v_0, u\}$ , then the load of route  $i$  immediately after visiting stop  $v$  must be at most the load of route  $i$  just after visiting  $u$  plus the number of students boarding route  $i$  at stop  $v$ . Otherwise, the lower bound  $-C_{\max}$  is imposed on  $l_{ui} + s_{vi} - l_{vi}$ , which is always satisfied since  $0 \leq l_{ui}, s_{vi}, l_{vi} \leq C_{\max}$ . Note that we did not incorporate Constraints (3.26) in our MIP because our computational experiments demonstrated that their inclusion tended to slow down the optimization process without yielding better solutions within the allotted computational time.

Third, Kek et al. (2008) recommend using the coefficient  $2C_{\max}$  in place of our coefficient  $C_{\max}$  in Constraints (3.18) and (3.26). The latter constraints have the “big  $M$ ” formulation. As noted by several researchers, it is important in such a formulation to use the smallest value of the coefficient  $M$  that is valid within the context of the model. Thus,  $M$  should be as small as possible, yet large enough to preserve all feasible solutions. Rubin (2011) states that mixing very large constraint coefficients  $M$  with other considerably smaller coefficients causes numerical instability in the simplex method due to the ensuing rounding errors. A very large constraint coefficient  $M$  can make the basis matrix seem singular or make a column ineligible for entry into the basis matrix seem eligible. The rounding errors may also lead to loss of precision when computing basic feasible solutions. Furthermore, Rubin (2011) claims that when employing B&B methods,

large values of  $M$  lead to loose bounds on the objective function at the nodes of the search tree, thereby slowing down the solution process. If the bounds are loose (i.e., very low for a minimization problem or very high for a maximization problem), then a few nodes can be pruned by bounds, and hence, the search tree will have more nodes. Note that pruning by bounds occurs when a node bound is not better than the objective function value of the best MIP solution found; thus, the node has no potential of generating a better MIP solution and does not need to be branched on further. For the above reasons, we opted for the coefficient  $C_{\max}$ , which is the smallest possible coefficient in our case.

The fourth point concerns alternative formulations of the SDVRP. Munari and Savelsbergh (2022) stated that formulations that use variables indexed by arcs and vehicles exhibit inherent symmetry, as they allow many equivalent (symmetric) solutions. This redundancy unnecessarily enlarges the solution space and can hinder the efficiency of exact solution methods. In our case, this symmetry manifests as multiple solutions with identical routes but differing only in the assignment of routes to vehicles (since all buses are assumed to be of capacity  $C_{\max}$  in our MIP). Therefore, Munari and Savelsbergh (2022) proposed alternative formulations of the SDVRP that use variables indexed by arcs only (binary variable  $x_{hij}$  equal to 1 if there is a route that traverses arcs  $(h, i)$  and  $(i, j)$  consecutively, and 0 otherwise). Although their formulations have shown superior performance over other formulations, they are based on the fundamental assumption that the travelling times satisfy the triangle inequality.

The travelling times in our problem instances, presented in Section 3.4.1, do not all satisfy the triangle inequality. This arises from the procedure used to generate the data from Bing Maps: for each arc, the shortest travel time is computed by summing the times of individual road segments along the selected path. Roads are divided into smaller segments, often at intersections, lane changes, or changes in speed limit, and each segment has a base travel time calculated from its length and typical speed. These segment times are then adjusted to account for expected traffic, turns, and control devices, including stop signs, traffic lights, pedestrian crossings, and speed bumps, while turns between segments may incur additional time penalties. Since each arc is treated independently, the specific segments chosen may differ slightly due to routing heuristics or internal tie-breaking rules, and small numerical differences can arise from the segment aggregation procedure. Consequently, the sum of real-world travel times along two arcs can occasionally be smaller than the direct travel time between the same endpoints, producing apparent violations of the triangle inequality. For example, for the smallest instance with 59 bus stops, 8.15% of the arc 3-tuples  $((h, i), (i, j), (h, j)), h \neq i, i \neq j, h \neq j$ , violate the triangle inequality by an aver-

age of 5.73 seconds travelling time. For this reason and because our “customer demands” (number of students waiting at each bus stop) are variables and thus not known in advance, it proved challenging to model our S-SBRP using only arc-indexed variables.

If the travelling times satisfy the triangle inequality, established structural properties of SDVRP solutions can be exploited in the MIP constraints, as has been done by Munari and Savelsbergh (2022) to reduce the solution space and strengthen the mathematical formulation. It has been proved that for any SDVRP feasible instance with integer demands/capacities and travelling costs/times that satisfy the triangle inequality, there exists an optimal solution with the following properties:

1. Each route is elementary, i.e., it visits each customer at most once (Feillet et al., 2006);
2. Two routes share at most one customer (Dror and Trudeau, 1989);
3. Each arc between customer nodes is traversed at most once (Feillet et al., 2006);
4. For each pair of reverse arcs between two customers, at most one of them is traversed (Desaulniers, 2010);
5. All delivery quantities are integers (Archetti et al., 2011).

In our case, for example, Properties 2 and 3 may not hold and thus the binary nature of variable  $x_{hij}$  in the formulations by Munari and Savelsbergh (2022) is inapplicable. Modelling  $x_{hij}$  as an integer variable rather than a binary variable would make it impractical to identify the distribution of the students waiting at a multistop among the different vehicles visiting that stop.

One final point relates to symmetry-breaking constraints that can be used to tackle the symmetry issue mentioned previously. These constraints eliminate some symmetric solutions while preserving at least one solution from each batch of symmetric solutions. In our case, we attempted to incorporate symmetry-breaking constraints in our MIP by forcing the total journey times to be sorted in non-increasing order as follows:

$$t_i \geq t_{i+1} \quad \forall i \in K \setminus \{k\}. \quad (3.27)$$

However, our computational experiments showed that including these constraints generally leads to a worse solution in the same amount of computational time and, in some cases, no identification of a feasible solution. Thus, Constraints (3.27) were not used in our final experiments.

We end this section by discussing the size of our proposed MIP. The formulation given in (3.8)-(3.25) contains  $\mathcal{O}(n^2k)$  variables and  $\mathcal{O}(n^2|V_2|)$  constraints.

Due to its large size and the exponential increase in the running time of exact methods with an increase in the model size, the MIP can only be solved optimally for very small problem instances (with a few tens of bus stops) within a reasonable time frame. That is why in Section 3.4 we set a time limit in our computational experiments, after which the best feasible solution found through a B&C algorithm (if any) is returned. The optimality gap (in percentage terms) between the objective value of the best feasible solution found and the best lower bound on the objective value across all unpruned nodes in the search tree is then presented.

### 3.3 Computational Complexity of the S-SBRP

It is interesting to establish the computational complexity of the S-SBRP with multistops and route duration constraints. Recall from Section 3.2.2 that this problem is equivalent to the time-constrained SDVRP. Dror and Trudeau (1990) proved that the SDVRP is an NP-hard problem. This can be proved by showing that the decision version of the SDVRP<sup>b</sup> is NP-complete, i.e., it is in class NP, and a known NP-complete problem is polynomially reducible to it.

First, the decision version of the SDVRP is in NP since, given a candidate solution, one can efficiently verify that it satisfies all constraints (for example, capacity and subtour elimination constraints). Second, the TSP can be viewed as a special case of the SDVRP by considering the following two simplifications: (i) set the vehicle capacity to be at least equal to the total demand of all customers, ensuring that all customers can be visited in a single route, and (ii) relax or set sufficiently large the maximum route duration, so that the time constraint does not restrict the tour. Under these conditions, minimizing the total travel time in the SDVRP reduces exactly to the classical TSP objective of finding the shortest possible tour visiting all customers. Now, the TSP has been proved to be NP-hard. This follows by a chain of polynomial-time reductions (Garey and Johnson, 1979), where  $P_1 \leq_p P_2$  denotes that problem  $P_1$  is polynomially reducible to  $P_2$ :

- Boolean Satisfiability (SAT) Problem<sup>c</sup> is NP-complete (Cook-Levin Theorem);
- SAT Problem  $\leq_p$  3-Satisfiability (3-SAT) Problem<sup>d</sup>;
- 3-SAT Problem  $\leq_p$  Vertex Cover Problem<sup>e</sup>;

---

<sup>b</sup>The decision version of an optimization problem is the yes/no question “Does there exist a feasible solution with objective value no worse than a prescribed value  $\Gamma$ ?”

<sup>c</sup>Given a collection of clauses (disjunctions of literals), is there a way to assign truth values to the variables to make all the clauses true?

<sup>d</sup>Special case of SAT in which each clause has exactly three literals.

<sup>e</sup>Given a graph and a positive integer  $m$ , is there a subset of vertices of size at most  $m$  such

- Vertex Cover Problem  $\leq_p$  Hamiltonian Circuit Problem<sup>f</sup>;
- Hamiltonian Circuit Problem  $\leq_p$  TSP.

Seeing that one of its special cases is NP-hard, it follows that the SDVRP is also NP-hard.

Another level of complexity is introduced if the bus stop selection subproblem is tackled. As noted by Park and Kim (2010), if bus stops have capacity restrictions, this subproblem can be transformed to the Generalized Assignment Problem (GAP)<sup>g</sup>. The GAP decision version has been proved to be NP-hard:

- 3-SAT Problem  $\leq_p$  3-Dimensional Matching (3DM) Problem<sup>h</sup> (Garey and Johnson, 1979);
- 3DM Problem  $\leq_p$  3-Partition Problem<sup>i</sup> (Garey and Johnson, 1975);
- 3-Partition Problem  $\leq_p$  GAP (Yuan et al., 2014).

## 3.4 Computational Results

In this section, we provide computational results for the MIP formulated in Section 3.2.4, which was coded in Python version 3.10.9 and solved using Gurobi Optimizer version 10.0.1. All experiments reported in this dissertation were conducted on a personal computer with a 2.8 GHz Quad-Core 11th Generation Intel Core i7-1165G7 CPU and 40 GB of DDR4 SDRAM. The MIP source code and output files of our experiments can be downloaded from (Sciortino, 2024c).

### 3.4.1 Problem Instances

A set of twenty real-world problem instances is considered here, summarized in Tables 3.1 and 3.2. The problem instances pertaining to the UK and Australia originate from the work of Lewis and Smith-Miles (2018) and can be found at (Lewis, 2017). The remaining problem instances were generated for this research and can be downloaded from (Sciortino, 2024d). The software used to create these problem instances is the same one used by Lewis and Smith-Miles (2018).

Each problem instance contains the following data:

- the cardinality of the sets  $V_1$ ,  $V_2$ , and  $E_2$ ;
- the minimum eligibility distance  $m_e$  (km);

---

that every edge in the graph is connected to a vertex in this subset?

<sup>f</sup>Given a graph, is there a Hamiltonian circuit (cycle that visits each vertex exactly once) in this graph?

<sup>g</sup>GAP minimizes the cost of assigning a set of items to a set of agents such that each item is assigned to precisely one agent, and subject to a capacity constraint for each agent.

<sup>h</sup>Given sets  $W, X, Y$  and a set of triples  $M \subseteq W \times X \times Y$ , is there a subset  $M' \subseteq M$  such that each element of  $W, X, Y$  appears exactly once?

<sup>i</sup>Given a set of  $3m$  integers, can it be partitioned into  $m$  disjoint three-element subsets such that each subset has exactly the same sum?

Table 3.1: Summary of the twenty real-world problem instances, listed in increasing order of  $|V_1|$ .  $S$  represents the total number of students, calculated as  $\sum_{w \in V_2} s(w)$ . Distances  $m_e$  and  $m_w$  are given in km.

Location	Country/State	$ V_1 $	$ V_2 $	$S$	$ E_2 $	$ E_1 $	$m_e$	$m_w$
Mġarr	Malta	60	110	190	657	3,540	1.0	1.0
Mellieħa	Malta	86	98	171	617	7,310	1.0	1.0
Porthcawl	Wales	153	42	66	2,246	23,256	3.2	1.6
Qrendi	Malta	158	150	255	2,528	24,806	1.0	1.0
Suffolk	England	174	123	209	1,276	30,102	4.8	1.6
Senglea	Malta	186	158	266	3,980	34,410	1.0	1.0
Victoria	Gozo	316	99	171	1,333	99,540	1.0	1.0
Pembroke	Malta	322	200	335	4,601	103,362	1.0	1.0
Canberra	ACT	331	296	499	3,925	109,230	4.8	1.0
Ħandaq	Malta	393	170	285	3,471	154,056	1.0	1.0
Valetta	Malta	445	159	268	4,518	197,580	1.0	1.0
Birkirkara	Malta	469	181	306	4,626	219,492	1.0	1.0
Ħamrun	Malta	518	192	321	5,028	267,806	1.0	1.0
Cardiff	Wales	552	90	156	5,753	304,512	4.8	1.6
Milton Keynes	England	579	149	274	9,612	334,662	4.8	1.6
Bridgend	Wales	633	221	381	10,019	400,056	4.82	1.6
Edinburgh-2	Scotland	917	190	320	48,666	839,972	1.6	1.6
Edinburgh-1	Scotland	959	409	680	34,594	918,722	1.6	1.6
Adelaide	S. Australia	1188	342	565	34,192	1,410,156	1.6	1.6
Brisbane	Queensland	1817	438	757	34,708	3,299,672	3.2	1.6

Table 3.2: Descriptive statistics on the travelling times (rounded to the nearest second), namely the mean  $\pm$  standard deviation, lower quartile (LQ), median and upper quartile (UQ). These are calculated across all arcs in  $E_1$ .

Location	Mean $\pm$ Std.	LQ	Median	UQ
Mġarr	336 $\pm$ 190	179	330	463
Mellieħa	465 $\pm$ 253	262	461	644
Porthcawl	428 $\pm$ 273	234	353	532
Qrendi	422 $\pm$ 223	255	398	558
Suffolk	894 $\pm$ 439	580	894	1203
Senglea	386 $\pm$ 189	248	369	509
Victoria	538 $\pm$ 244	362	529	711
Pembroke	445 $\pm$ 186	311	446	576
Canberra	732 $\pm$ 362	421	770	1022
Ħandaq	506 $\pm$ 210	358	500	647
Valetta	483 $\pm$ 187	357	489	611
Birkirkara	499 $\pm$ 203	356	496	638
Ħamrun	501 $\pm$ 188	374	509	633
Cardiff	618 $\pm$ 396	319	510	837
Milton Keynes	379 $\pm$ 155	262	377	496
Bridgend	883 $\pm$ 445	566	863	1133
Edinburgh-2	491 $\pm$ 230	317	478	648
Edinburgh-1	702 $\pm$ 329	446	693	947
Adelaide	703 $\pm$ 308	471	692	918
Brisbane	775 $\pm$ 325	540	774	999

- the maximum walking distance  $m_w$  (km);
- the name and GPS coordinates of the school;
- the name and GPS coordinates of each potential bus stop;
- the family name, number of students requiring school transport, and GPS coordinates of each student address;
- the shortest driving time (s) and distance (km) along each arc in  $E_1$ ;
- the shortest walking time (s) and distance (km) of each feasible walk in  $E_2$ .

We highlight that vertices in  $V_1$  are indexed from 0 to  $n = |V_1| - 1$ , with 0 representing the school and the other indices corresponding to potential bus stops. Likewise, student addresses are indexed from 0 to  $|V_2| - 1$ .

Each problem instance was generated as follows. First, the location of a school was identified and several student addresses were selected within the local area of the school, but at least  $m_e$  km from the school. These addresses were chosen randomly since data protection regulations prohibit the publishing of personal information about students. The number of students requiring school transport at each address was then generated randomly according to the following discrete probability distribution: 1, 2, 3, and 4 with probabilities 0.45, 0.4, 0.14, and 0.01, respectively. As stated by Lewis and Smith-Miles (2018), this distribution approximates the relevant statistics in the locations considered. Next, potential bus stops were identified through public records such that each stop has at least one address within walking distance  $m_w$  and each address has at least one stop within walking distance  $m_w$ . Finally, the shortest driving time and distance along each arc in  $E_1$  and the shortest walking time and distance of each feasible walk in  $E_2$  were determined using the Bing Maps Routes application programming interface. Two example problem instances are visualized in Figure 3.4<sup>j</sup>.

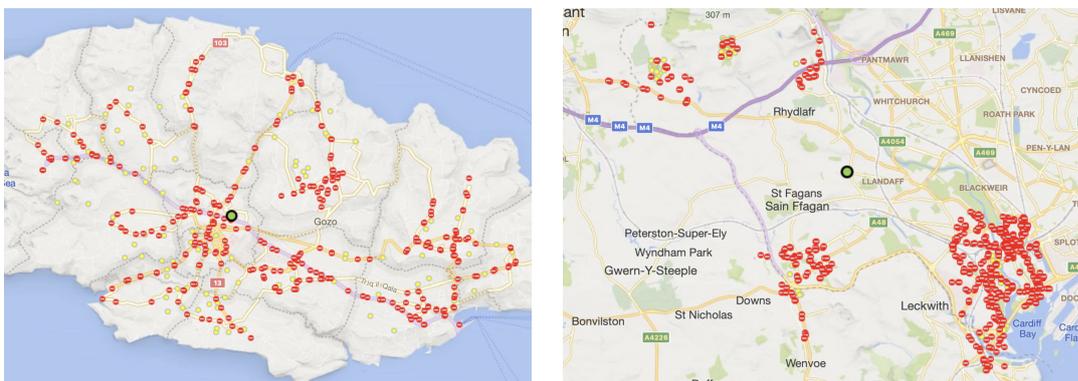


Figure 3.4: Victoria (left) and Cardiff (right) problem instances. The lime dot represents the school, the red dots represent the potential bus stops, and the yellow dots represent the student addresses.

<sup>j</sup>A visualization tool, provided by my supervisor Prof. Lewis, was used to generate the visual outputs presented in this dissertation.

Here, we assume that the maximum journey time  $m_t = 2700$  seconds (45 minutes) for all the locations, and that the vehicle capacities differ across the locations. The selection of the capacities was based on publicly available fleet information from school transport operators in the locations under study. For locations in Malta, the available capacities (excluding the driver) are taken to be 8, 14, 16, 18, 20, 36, 44, and 53. Note that capacities above 53 are excluded since very long and double-decker buses are not currently used for school transport in Malta due to unsuitable road infrastructure. For locations in Australia, the assumed capacities are 11, 13, 18, 21, 24, 28, 33, 35, 37, 39, 43, 45, 47, 49, 51, 53, 55, 57, 59, 61, 65, 70, 78, and 80, while those for locations in the UK are 8, 12, 16, 23, 25, 27, 29, 33, 37, 39, 43, 45, 49, 51, 53, 55, 57, 61, 63, 65, 70, 74, 78, and 80.

### 3.4.2 Analysis of Results

Tables 3.3 and 3.4 contain some model statistics and the results of our computational experiments, respectively. Each experiment was allowed to run for up to six hours. As mentioned in Section 3.2.4, we initially attempted to find feasible solutions with the number  $k$  of routes fixed to  $\underline{k} = \lceil \sum_{w \in V_2} s(w) / C_{\max} \rceil$  for each instance. Here,  $C_{\max}$  denotes the assumed maximum vehicle capacity for the given instance (Malta: 53, Australia and UK: 80). If no feasible solution was found within six hours,  $k$  was incremented to  $\underline{k} + 1$  and the experiment was restarted. If difficulty in achieving feasible solutions was yet again observed,  $k$  was incremented to  $\underline{k} + 2$  and the experiment was repeated for a final time. The results presented in Tables 3.3 and 3.4 therefore correspond to the model with the highest attempted value for  $k$ . Note that a feasible solution with  $k \in \{\underline{k}, \underline{k} + 1, \underline{k} + 2\}$  was only found for eight out of twenty instances. Out of memory errors were also encountered for the last four instances.

Table 3.3 presents the number of decision variables, constraints, and explored nodes in the B&C process. The optimality gap (%) at the end of each two-hour interval is also displayed in the fourth column. A hyphen in this column indicates that no feasible solution was found during the respective interval. For the Senglea instance, for example, the first feasible solution was found in the second two-hour interval and the best feasible solution has an optimality gap of 46.7%. The objective value of the best feasible solution (presented in Table 3.4) is 77.30 minutes, while the best lower bound on the objective value is 41.23 minutes (this result can be found in the console output file at (Sciortino, 2024c)). The latter represents an objective value that may be attained or closely attained if the search tree is further explored. However, there is, of course, no guarantee

that an unexplored solution with such an objective value exists.

The model statistics in Table 3.3 show that the computational effort depends on the number of potential bus stops, the number of buses, the number of student addresses, and the structure of the considered instance. By structure, we refer to the positioning of potential bus stops and student addresses. For example, feasible solutions were found for the Cardiff instance but not for instances with comparable student addresses and fewer potential bus stops, such as the Victoria instance. This may be because the student addresses in the Cardiff instance are arranged in clusters, as shown in Figure 3.4, which facilitates routing since students in localized areas can walk to the same bus stop. Moreover, the number of explored nodes is relatively small for some instances since the corresponding model size is very large, and thus, nodes take longer to process.

The best feasible solutions are displayed in Table 3.4 and visualized in Appendix A. For ease of reference, one visualization – the best solution for the M̄garr instance – is also provided in Figure 3.5. From Table 3.4, we see that feasible solutions using the lower bound of  $k$  routes were found for five instances. The Cardiff and Milton Keynes instances required one additional route beyond  $k$ , while the Suffolk instance required two additional routes. The latter is likely a result of the relatively high travel times in the data, as observed in Table 3.2. The ‘Load’ value gives the number of boarding students on each route, while the ‘Cap.’ value shows the capacity of the smallest bus that can fit all boarding students. In addition, the ‘JT’ column gives each route’s journey time (minutes), and the ‘TJT’ column shows the total journey time (minutes) across all routes.

Note that most of the journey times are less than 20 minutes and that the percentage of capacity utilized is generally high (ranging from 63.89% to 100%). Regarding the Suffolk and Senglea instances, the journey times are not as balanced. In the former case where the imbalance is more evident, 9 students have a riding time of approximately 10 minutes, and 62 students have a riding time of approximately 36 minutes. The Suffolk solution also shows that a higher num-

Table 3.3: MIP model statistics.

Location	#Variables	#Constraints	Gap (% every 2 hours)	#Nodes
M̄garr	21,370	21,650	(16.5, 16.1, 15.9)	798,840
Mellieħa	38,602	38,884	(14.5, 13.6, 13.2)	749,290
Porthcawl	30,099	30,336	(35.0, 33.0, 31.6)	86,998
Qrendi	149,950	150,408	(50.0, 42.8, 35.0)	20,866
Suffolk	174,399	174,819	(20.0, 20.0, 18.1)	100,162
Senglea	239,038	239,540	( – , 48.4, 46.7)	12,938
Cardiff	967,014	967,746	(34.6, 34.2, 32.1)	1,165
Milton Keynes	1,768,117	1,768,994	( – , 27.8, 21.3)	3

ber of stops (typically to cater for more boarding students) generally increases a route's journey time.

Table 3.4: MIP solutions.  $k$  denotes the number of routes in the solution, Cap. is short for Capacity, JT stands for Journey Time (minutes), and TJT stands for Total Journey Time (minutes). The routes are presented in terms of the indices of the vertices in  $V_1$ , and the subscripts represent the number of boarding students.

Location	$\underline{k}$ $k$	Routes	Load/Cap.	JT	TJT
Mġarr	4 4	21 <sub>32</sub> → 19 <sub>11</sub> → 0	43/44	9.83	54.02
		42 <sub>10</sub> → 39 <sub>21</sub> → 35 <sub>12</sub> → 15 <sub>1</sub> → 0	44/44	14.33	
		47 <sub>18</sub> → 52 <sub>20</sub> → 57 <sub>15</sub> → 0	53/53	15.40	
		48 <sub>1</sub> → 33 <sub>9</sub> → 11 <sub>10</sub> → 12 <sub>29</sub> → 15 <sub>1</sub> → 0	50/53	14.45	
Mellieħa	4 4	30 <sub>8</sub> → 20 <sub>16</sub> → 27 <sub>8</sub> → 0	32/36	14.70	56.30
		66 <sub>5</sub> → 83 <sub>9</sub> → 80 <sub>14</sub> → 5 <sub>21</sub> → 0	49/53	15.80	
		43 <sub>14</sub> → 49 <sub>3</sub> → 46 <sub>14</sub> → 59 <sub>19</sub> → 0	50/53	15.82	
		33 <sub>7</sub> → 7 <sub>5</sub> → 9 <sub>28</sub> → 0	40/44	9.98	
Porthcawl	1 1	16 <sub>5</sub> → 27 <sub>19</sub> → 151 <sub>17</sub> → 38 <sub>25</sub> → 0	66/70	26.87	26.87
		129 <sub>12</sub> → 10 <sub>15</sub> → 126 <sub>16</sub> → 68 <sub>10</sub> → 0	53/53	19.77	
Qrendi	5 5	155 <sub>4</sub> → 152 <sub>9</sub> → 133 <sub>21</sub> → (cont.)			80.13
		50 <sub>13</sub> → 146 <sub>2</sub> → 0	49/53	17.30	
		112 <sub>27</sub> → 135 <sub>24</sub> → 0	51/53	15.70	
		29 <sub>37</sub> → 76 <sub>16</sub> → 0	53/53	13.75	
		68 <sub>8</sub> → 63 <sub>20</sub> → 137 <sub>12</sub> → 121 <sub>9</sub> → 0	49/53	13.62	
Suffolk	3 5	129 <sub>5</sub> → 73 <sub>5</sub> → 122 <sub>12</sub> → 0	22/23	15.18	114.25
		43 <sub>4</sub> → 116 <sub>18</sub> → 172 <sub>2</sub> → 5 <sub>22</sub> → 98 <sub>4</sub> → 0	80/80	31.08	
		70 <sub>29</sub> → 130 <sub>6</sub> → 173 <sub>1</sub> → 102 <sub>2</sub> → (cont.)			
		67 <sub>4</sub> → 9 <sub>8</sub> → 65 <sub>12</sub> → 0	62/63	36.32	
		81 <sub>9</sub> → 0	9/12	10.00	
		158 <sub>9</sub> → 111 <sub>17</sub> → 140 <sub>8</sub> → 43 <sub>2</sub> → 0	36/37	21.67	
Senglea	6 6	14 <sub>28</sub> → 140 <sub>16</sub> → 19 <sub>9</sub> → 0	53/53	15.35	77.30
		19 <sub>48</sub> → 0	48/53	10.22	
		123 <sub>21</sub> → 181 <sub>6</sub> → 150 <sub>10</sub> → 171 <sub>16</sub> → 0	53/53	19.17	
		42 <sub>48</sub> → 53 <sub>1</sub> → 0	49/53	11.45	
		59 <sub>16</sub> → 53 <sub>7</sub> → 0	23/36	6.25	
		182 <sub>9</sub> → 121 <sub>18</sub> → 67 <sub>13</sub> → 0	40/44	14.87	
Cardiff	2 3	347 <sub>32</sub> → 502 <sub>18</sub> → 116 <sub>12</sub> → 0	62/63	24.93	57.25
		384 <sub>41</sub> → 474 <sub>6</sub> → 0	47/49	14.18	
		534 <sub>45</sub> → 92 <sub>2</sub> → 0	47/49	18.13	
Milton Keynes	4 5	418 <sub>52</sub> → 0	52/53	12.18	66.80
		564 <sub>40</sub> → 373 <sub>32</sub> → 0	72/74	13.97	
		191 <sub>42</sub> → 0	42/43	13.58	
		341 <sub>13</sub> → 110 <sub>37</sub> → 0	50/51	13.92	
		395 <sub>41</sub> → 365 <sub>17</sub> → 0	58/61	13.15	

As a case study, we now delve into the MIP solution of the Mġarr instance, as visualized in Figure 3.5. Some subroutes in the figure are not visible because

they overlap with other subroutes. Additionally, certain roads in the blue route are traversed in both directions. The solution contains 13 visited stops, with two buses serving Stop 15. For the subset  $V'_1$  of bus stops, there are 158 feasible walks in  $E_2$ , yielding an average of 1.4 feasible walks per address. A total of 67, 38, and 5 addresses have one, two, and three feasible walks, respectively. For instance, the two students at Address 58 can only walk to Stop 15, whereas the student at Address 63 can walk to Stops 12, 15, and 35 with walking distances of 0.211, 0.967, and 0.909 km, respectively. Due to Constraint (3.14), the latter student is assigned to Stop 12. The same holds for other students with more than one feasible walk. Observe from the routes in Table 3.4 that the two students at Address 58 board different buses. As stated at the end of Section 3.2.2, this is not an ideal situation and is a limitation of our MIP. By allowing bus stops to be visited by multiple buses in the mathematical formulation, we risk having students living at the same address being assigned to different routes. For the Mġarr solution, this issue can be easily resolved by removing Stop 15 from the second route and having both students at Address 58 board the fourth bus. However, this results in an increase in the total journey time by five seconds since the triangle inequality is violated for the arc 3-tuple  $((35,15),(15,0),(35,0))$  (driving times in seconds are  $d(35, 15) = 119, d(15, 0) = 6, d(35, 0) = 145$ ).

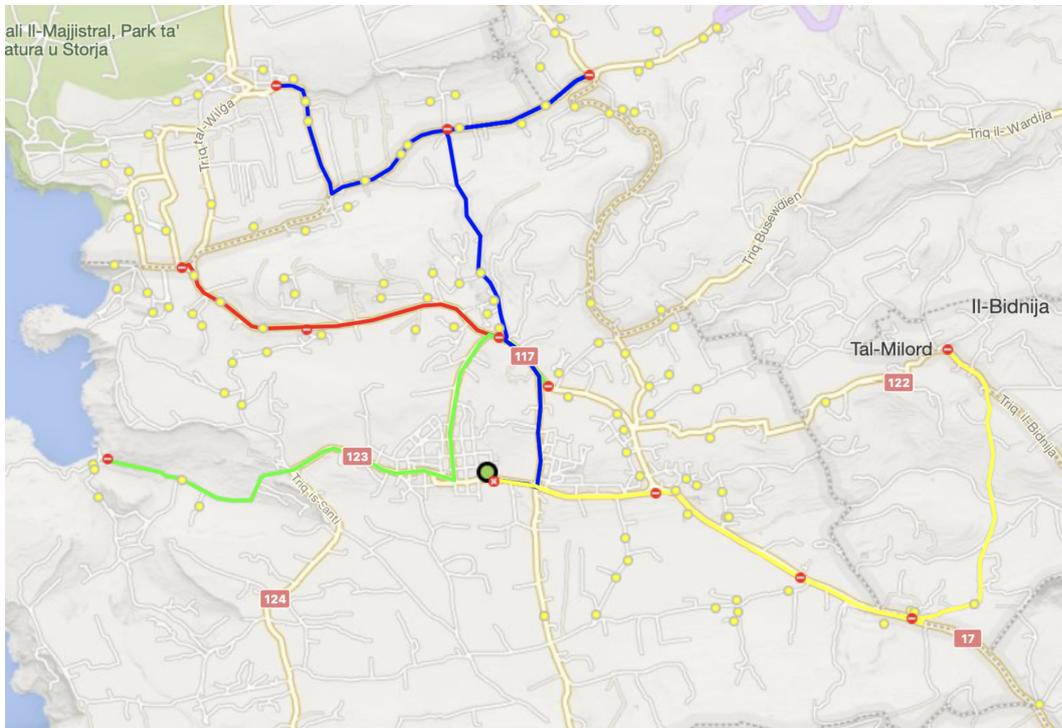


Figure 3.5: Mġarr MIP solution with 4 routes (red, lime, blue, yellow), 190 students, average walk 6.06 minutes, and average journey time 13.50 minutes. The lime dot represents the school, the red dots represent the visited bus stops, and the yellow dots represent the student addresses.

## 3.5 Summary

In this chapter, we discussed the definitions, assumptions, constraints, and objectives of our S-SBRP. This problem includes many realistic features such as student eligibility for school transport, maximum walking distances, maximum riding times, different types of buses, bus stop selection, and bus dwell times. The S-SBRP is inherently multi-objective due to considerations of efficiency, effectiveness, and equity. However, it was highlighted that our top priority is to minimize the number  $k$  of routes because the major costs are the bus acquisition costs and drivers' salaries, which are both heavily dependent on  $k$ . In this chapter, a MIP was proposed for the S-SBRP that integrates the bus stop selection and bus route generation subproblems. This MIP takes  $k$  as an input parameter and attempts to construct feasible solutions having exactly  $k$  routes. Here,  $k$  is initially set to the lower bound  $\underline{k}$ .

A set of real-world problem instances from Malta, the UK, and Australia, with sizes upwards of 1800 potential bus stops and 750 students, was presented. This set will be used throughout the remainder of this thesis. Computational experiments using the MIP were also performed to demonstrate the high computational effort necessary to solve real-world S-SBRP instances. From these experiments, it was noted that finding a feasible, let alone an optimal, solution using  $\underline{k}$  routes (or a number slightly higher) is challenging, especially when the number of potential bus stops exceeds approximately 200. Moreover, the S-SBRP was shown to be NP-hard, suggesting that finding an optimal solution for a larger problem instance can take an infeasibly large amount of time.

For this reason, large real-life problem instances of the S-SBRP are commonly addressed using heuristic approaches that provide reasonably good solutions in a more practicable time frame. What is vital in practice is to successfully provide the transport service to all students concerned, in which case finding an approximate solution quickly is preferred over spending a much longer period attempting to find an optimal solution. This motivates our focus on heuristic approaches for the S-SBRP in the next chapters.



# Chapter 4

## A Heuristic Algorithm for the Single-School Bus Routing Problem

### 4.1 Introduction

The NP-hard nature of the S-SBRP, coupled with the typically large problem sizes encountered in real-world scenarios, suggests that heuristics are a suitable solution strategy. In this chapter, we develop a specialized heuristic algorithm capable of producing good solutions to our S-SBRP within reasonable computational time. Vehicle routing heuristics are categorized into three main classes, namely *constructive heuristics*, *improvement heuristics*, and *metaheuristics* (Liu et al., 2023). The latter have been described in Section 2.3.2.2. Therefore, the focus in this section will be on the former two classes.

Constructive heuristics build solutions by iteratively inserting a new stop into a currently incomplete route or a new route. Construction of routes can be performed either sequentially, one route at a time, or in parallel, with multiple routes being constructed simultaneously. In the former strategy, a new route is started only if a stop that is to be visited cannot be feasibly inserted into an existing route. Meanwhile, in parallel construction, multiple routes are devised concurrently by prespecifying the number of vehicles used. Drawbacks of sequential construction are that it is difficult to predict the number of routes, and the last route is often markedly less loaded than the other routes (Labadie et al., 2016). Moreover, the last route tends to have poor quality due to the nodes being scattered over a large geographic area (Larsen, 1999). In parallel construction, the prescribed number of routes is not guaranteed at the end; instead, extra routes may be needed to satisfy all constraints.

Improvement heuristics aim to enhance a given solution by exploring its surrounding neighbourhood. A neighbourhood is typically defined by a specific move operator, which determines the set of candidate solutions reachable from the current one. These neighbouring solutions are evaluated using the objective function, and if a better solution is found, it replaces the current solution. This process generally continues until no further improvement is possible. While improvement heuristics are effective at quickly locating a local optimum, they can easily get trapped in such an optimum (Laporte and Semet, 2002).

Constructive and improvement heuristics are delved into more deeply in Sections 4.2 and 4.3, respectively, as both types are applied in our heuristic algorithm. Section 4.4 describes the ILS metaheuristic that serves as the foundation of our algorithm. This metaheuristic embeds improvement heuristics within an iterative process to generate a sequence of local optima. In Section 4.5, we then present the overall framework of our algorithm. Finally, in Section 4.6, we provide and compare the results achieved by different variants of our proposed algorithm on the twenty real-world problem instances summarized in Tables 3.1 and 3.2.

## 4.2 Constructive Heuristics

We begin this section with an overview of some constructive heuristics commonly applied in VRP literature, particularly SBRP publications.

### 4.2.1 Related Works

In the seminal paper by Newton and Thomas (1969), a two-step procedure was proposed for constructing S-SBRP solutions. First, a giant tour through each stop is determined through the nearest neighbour (NN) approach. It is then improved by three link changes and two link changes (if the time matrix is symmetric). This tour is then partitioned into a series of routes that adhere to the order previously determined, and the capacity and riding time constraints. This procedure follows the route-first cluster-second principle (Beasley, 1983) and was also employed by Bodin and Berman (1979) and Dulac et al. (1980), who form the tour using the three-opt procedure (Lin, 1965) and the Clarke-Wright savings heuristic (Clarke and Wright, 1964), respectively. Route-first cluster-second approaches are not as common, because CVRP research has shown that they tend to produce poor solutions (Cordeau et al., 2007).

One of the most common CVRP constructive heuristics is the Clarke-Wright savings (CWS) heuristic. This heuristic starts from a solution in which each stop features in a separate route. Then, for each pair of stops (one being the first

stop in a route and the other the last stop in another route), the decrease in cost (saving) resulting from merging the two routes that contain the stops is calculated. The merge with the highest saving is performed, given that the capacity constraint is respected. This is repeated iteratively until no further merging is possible. In the context of the SBRP, the CWS heuristic was applied in the works of Bennett and Gazis (1972), Desrosiers et al. (1981), Russell et al. (1986), Riera-Ledesma and Salazar-González (2012, 2013), Campbell et al. (2015), Souza Lima et al. (2016), and Dang et al. (2019), amongst others. For the CVRP, Avdoshin and Beresneva (2019) compared six constructive heuristics, namely sequential insertion (Laporte and Semet, 2002), improved parallel insertion (Laporte et al., 1985), NN, CWS, Yellow (1970)'s variant of CWS, and sweep (described on the next page). They found that CWS is the best except when nodes are located in concentric rays, in which case NN is better.

An improved version of the CWS heuristic is the greedy randomized adaptive search procedure (GRASP) (Feo and Resende, 1989, 1995). GRASP builds a restricted candidate list (RCL) of feasible merges in a greedy fashion and, in each iteration, one merge is randomly selected from this list. The size of this list is a parameter that controls the balance between greediness and randomness. The random component prevents the heuristic from getting stuck in a local optimum. Some SBRP researchers who have used GRASP are Corberán et al. (2002), Schittekat et al. (2013), Siqueira et al. (2016), Faraj et al. (2014), and Silva et al. (2015). In the latter two works, an RCL of feasible stop insertions was built using the NN algorithm. Analogously, Sales et al. (2018) chose the first stop randomly and then used an RCL of the three closest stops to the most recently added stop.

Dulac et al. (1980) and Spada et al. (2005) applied Rosenkrantz et al. (1974)'s sequential insertion heuristic to the SBRP. Each route is initiated at the farthest unrouted stop from the school. Then, at each step, the cost to insert each unrouted stop at the best position in the current route is calculated. The stop with the lowest such cost is inserted at its respective position, provided that the constraints are satisfied. If no stop can be feasibly inserted into the current route, a new route is created. A similar strategy was employed by Braca et al. (1997) and Bertsimas et al. (2019) who used a randomized variant of the location-based CVRP heuristic of Bramel and Simchi-Levi (1995). The difference here is that the first stop of each route is selected at random. Siqueira et al. (2016) showed that GRASP is better than Braca et al. (1997)'s approach with a reduction of 36.29% in total cost, 32.11% in number of buses, and 54.17% in run time. Chapleau et al. (1985) compared Rosenkrantz et al. (1974)'s insertion and the CWS heuristics with their developed cluster-first route-second heuristic. They first generated a minimum number of student clusters, with each cluster including the furthest

unassigned student node from the school and the closest unassigned nodes to it that fit in the route capacity. At the end of this procedure, some nodes may not belong to any cluster, in which case the cluster with the smallest number of students is considered and the closest such node is added iteratively. Two types of node exchanges are attempted for improving a capacity measure and a set of stops is then selected for each cluster. A route through these stops is finally determined using a two-opt procedure (Croes, 1958) and stop exchanges are possibly applied to reduce the total distance. Their results showed that their heuristic is superior in decreasing the number of routes without increasing the average walking distance and the average route length per student.

Another constructive heuristic, which follows the cluster-first route-second principle, is the sweep algorithm first proposed by Wren and Holliday (1972) and further developed by Gillett and Miller (1974). Clusters of nodes are created to form the routes by using the polar-coordinate angle of each node relative to the depot. Routes are generated sequentially by starting with the node with the smallest angle, adding nodes in increasing order of their angle, and closing up a route (cluster) when the capacity or distance/time constraint is violated. Then an improvement procedure is applied whereby one considers replacing one node in a route with one or more nodes in the succeeding route until no further improvement is found. Gillett and Miller (1974) considered route generation from different starting nodes and in clockwise and counter-clockwise directions. According to Dulac et al. (1980), in some cases, pairs of stops such as two stops on opposite sides of a river, cannot belong to the same cluster. Moreover, Labadie et al. (2016) argued that the last cluster is often too small and that “the algorithm does not work well when the depot is excentered, for instance in a port at the extremity of a peninsula”. Park et al. (2012) used a variant of the sweep algorithm for the mixed load SBRP. Furthermore, Dulac et al. (1980) showed that Newton and Thomas (1969)’s two-step procedure dominates Gillett and Miller (1974)’s sweep algorithm. Thangiah and Nygard (1992) compared a system using the sweep algorithm and CWS with another system that clusters the student locations using genetic sectoring and constructs the routes using the cheapest insertion heuristic. The latter system was shown to attain consistently better results.

Na et al. (2011) developed a hybrid algorithm, called the sweep nearest algorithm (SNA), that combines the sweep and NN algorithms. SNA assigns nodes to clusters via the NN heuristic rather than in increasing/decreasing order of their polar angles. Xue et al. (2023) applied SNA to cluster stops, accounting for the capacity and distance constraints. First, the stops are sorted in increasing order of their polar angles. A new cluster is started from the unassigned stop with the smallest angle, and NN is iteratively used to select the next unassigned stop until

all stops are assigned.

Bowerman et al. (1995) developed another cluster-first route-second algorithm which uses the space-filling curve with optimal partitioning heuristic (Bowerman et al., 1994). A space-filling curve is used to transform the routing problem from the plane to a closed curve and a giant tour is constructed by visiting all nodes in the order dictated by the curve. This tour is then optimally partitioned using dynamic programming techniques with respect to a weighted objective function. The resulting routes are used to define the student clusters. For each of these clusters, a subset of bus stops is selected such that every student has a bus stop within maximum walking distance. Students are then assigned to their nearest selected bus stop and a route is generated using a two-opt procedure. This procedure is carried out several times for each cluster and the route with the least total weighted distance (walking distance and route length) is selected.

Corberán et al. (2002) also clustered nodes by creating sectors of a fixed size around sequentially selected nodes. First, the nodes are ordered in decreasing order of their travelling time to the school. At each step, the next node is randomly selected so that its travelling time to the school is at least a prespecified portion of that of the furthest unassigned node. If the selected node can be feasibly added to an existing sector, then the node is inserted in the position that causes the least increase in route time; otherwise, a new sector is created. Corberán et al. (2002) showed that their heuristic generates less dispersed solutions concerning the maximum route time when compared to GRASP; however, the latter tends to produce better best solutions. Recently, Ansari et al. (2021) used a capacitated clustering problem methodology (Mulvey and Beck, 1984) in which random students are selected as centroids and clusters are formed around them, minimizing the total distance between students and the centroids and ensuring that clusters adhere to the bus capacities. Bus routes are then devised via mathematical modelling to minimize the total travel distance.

Pacheco and Marti (2006) compared four constructive procedures, two of which are Corberán et al. (2002)'s approach and GRASP. The third procedure addresses a GAP by a modified version of the heuristic of Martello and Toth (1981) to assign locations to clusters. At each step, the location with the greatest difference between the lowest and second-lowest feasible assignment costs – considering capacity and time constraints – is selected for clustering. This location is assigned to the cluster corresponding to the lowest cost and inserted in the position that minimizes the route length. The fourth procedure first orders the locations. Five ordering strategies are compared, with the best one being ordering the locations in decreasing order of their distance to the school. When considering the next location, an RCL with feasible best insertions resulting in

route lengths that are at most a threshold value is constructed. A random insertion is then selected and performed. Pacheco and Marti (2006) showed that the third procedure outperforms the other three and that the fourth procedure also produces relatively good solutions.

Ripplinger (2005) sorts the homes in decreasing order of their distance from the school. The furthest home is put at the start of the first route and all homes in its shadow (i.e., homes which lie at an angle relative to the school within a range above and below that of the reference home relative to the school) are removed. This is repeated until a predefined number of routes are assigned a starting location, ignoring capacity and time constraints. Once all routes have a starting location, the cost of inserting all remaining homes within the routes is calculated and the insertion with the minimal cost is selected until all homes are assigned to a route. Parallel cheapest insertion was also used by Li and Chow (2021). Another work that constructs routes from pick-up points that are the furthest away from the school is that of Li and Fu (2002). In this work, routes are constructed sequentially and the only route possibly not full is the last one. For each route generation, the furthest pick-up point from the school is selected as the starting location and students are assigned in turn along the shortest, second shortest, third shortest, etc. route to the school until the route is full or no more students are left to be assigned (in the case of the last route). More recently, Calvete et al. (2022) adapted ideas from those of Fischetti et al. (1997) for the generalized TSP. First, each route is initialized by sequentially selecting the pick-up point furthest from the depot and the pick-up points that have already been considered. Then, each remaining node (a pick-up point or an individual's location) is selected randomly. If the selected node is a pick-up point, it is inserted in the route that yields the minimum insertion cost; otherwise, the selected individual is allocated to the closest routed pick-up point that has spare capacity.

Ochoa-Zezzatti et al. (2020) grouped students to bus stops such that each student does not walk more than 500 metres to reach the stop. If a student lies between the intersection of two bus stops' radii of walking distance, they are assigned to the closer group. The NN algorithm is then implemented, starting from different initial points of each route to find the best solution. Huo et al. (2014) use a similar clustering idea and then generate routes by ACO to minimize the total travel distance.

Shafahi et al. (2018a) proposed a matching-based constructive heuristic, whereby stops are first clustered using the  $k$ -means algorithm with  $k$  being the number of routes. From each cluster, the furthest stop from the school is added to a separate route. This is done because such stops are influential to the journey

time constraints. The unrouted stops are then feasibly inserted into the existing or new routes by iteratively solving minimum-cost matching problems by the Hungarian algorithm. Kotoula et al. (2017) also used the  $k$ -means algorithm to cluster stops. A stop order is then achieved by initially selecting the furthest stop from the school and then iteratively selecting the closest unrouted stop to the last added stop. Following this, a GA for the TSP is applied that constructs a population of initial solutions by the NN heuristic, CWS heuristic, and completely random solutions. Ren et al. (2019) employed an iterative clustering method based on  $k$ -means for the stops' generation and students' allocation. After this, an ACO routing algorithm with two local search operators is applied to generate bus routes. Another work that uses a graph-based method is that of Bögl et al. (2015) who adopted Prim's algorithm to construct a minimum spanning tree for the set of visited bus stops. The starting node of their tree represents the school, while the edges point toward the previously selected bus stop. To adhere to the time constraints, they limit the maximum length of a path in the minimum spanning tree.

In their CG approach, Caceres et al. (2019) employed both a savings-based and a probabilistic insertion heuristic. Lewis and Smith-Miles (2018) used a constructive heuristic inspired by the first-fit descending heuristic for bin packing problems. Bus stops are ordered in decreasing order of the number of boarding students. At each step, a bus stop is inserted in the least loaded route having enough spare capacity for the boarding students. If no such route exists, the least loaded route is filled up with some of the boarding students, and a multistop is created with the remaining boarding students. Ümit and Kiliç (2019) employed random construction, considering the capacity constraints. In each step, one unrouted stop is randomly selected and added to the route. If there is no unrouted stop that fits in the spare capacity, then a new route is created. Kinable et al. (2014) and Ozmen and Sahin (2021) also constructed randomly generated initial bus routes.

## 4.2.2 Implemented Constructive Heuristics

Recall that our S-SBRP formulation tackles both the bus stop selection and route generation subproblems. Therefore, in our case, it is fundamental to select a subset of visited bus stops  $V_1'$  before constructing the routes. The common formulations of the bus stop selection subproblem have been discussed in Section 2.3.1. In our case, we seek to minimize the number of visited bus stops since this approach appears advantageous in reducing the length (and thus the financial cost) of routes.

Here, the selection of  $V'_1$  is performed using the following greedy heuristic, based on the Kolesar-Walker heuristic (Kolesar and Walker, 1972) for the uni-cost set covering problem. First, all compulsory stops are included in  $V'_1$ . The non-compulsory stops are then arranged in non-increasing order according to the number of currently uncovered addresses they serve. The stop with the largest such value is then added to  $V'_1$ , breaking ties randomly. This ordering and selection procedure is repeated until a complete covering of  $V_2$  is obtained, similar to what is done by Dulac et al. (1980), Chapleau et al. (1985), Bowerman et al. (1995), and Lewis and Smith-Miles (2018). Each address in  $V_2$  is then assigned to the closest bus stop in  $V'_1$ . The assignment of addresses to stops determines the number  $s_{V'_1}(v)$  of boarding students at each stop  $v \in V'_1$ . It may be the case that some stops have no boarding students, in which case they are removed from  $V'_1$ .

Having selected  $V'_1$ , each bus stop in this subset is then assigned to one of the  $k$  routes such that each bus is not overloaded. During the construction phase, the maximum journey time constraints (3.4) are not enforced as these are tackled later in the improvement phase. In Sections 4.2.2.1 to 4.2.2.3, we discuss three different constructive heuristics for performing the stop-route assignments. These will be compared later in the computational results presented in Section 4.6. To facilitate explanation, the selected bus stops and the corresponding pairwise travel times presented in Table 4.1 will be used as a running example throughout Sections 4.2.2.1 to 4.2.2.3.

Table 4.1: Example of a subset of visited bus stops and their pairwise travel times (in minutes). The first column gives the number of boarding students.

#Boarding Students	Stop	0	64	253	318	353	412
2	64	9.45	0	10.62	14.50	25.05	16.60
45	253	10.88	11.22	0	3.43	23.48	13.42
47	318	11.17	13.55	3.02	0	22.73	13.70
50	353	17.30	25.38	23.47	22.97	0	10.18
12	412	8.87	16.97	12.98	13.25	10.95	0

#### 4.2.2.1 Parallel Nearest Neighbour Heuristic (PNNH)

The first considered strategy is a parallel backward implementation of the NN heuristic. To start,  $k$  empty routes are defined and the remaining capacity  $c_i$  of each route  $R_i, i \in \{1, 2, \dots, k\}$ , is set to  $C_{\max}$ . The  $k$  closest bus stops to the school are inserted at the front of the routes, one in each route. Closeness to school is measured by the dwell time at the stop plus the shortest driving time from the stop to the school. To calculate the dwell time at stop  $v \in V'_1$

in route  $R_i$ , the minimum of  $c_i$  and  $s_{V'_1}(v)$  is considered as there may be more than  $c_i$  students boarding at stop  $v$ . In this case, a multistop is created since the remaining  $s_{V'_1}(v) - c_i$  students boarding at stop  $v$  must be assigned to a different route  $R_j, j \neq i$ . The remaining capacities  $c_i$  are then updated accordingly. This iterative procedure of determining the closest stop to the most recently added stop in route  $R_i$ , inserting it at the front of the route, and updating the remaining capacity  $c_i$ , is repeated until all stops in  $V'_1$  are assigned to a route. Note that routes are processed in increasing order of  $i$ , with a route being skipped once it becomes fully loaded.

As an example, consider the data in Table 4.1 and assume  $C_{\max} = 80$ . The total number of boarding students is 156; therefore, two routes are required to satisfy the capacity constraints. Table 4.2 presents the PNNH heuristic applied step by step. At each step, the stop closest (in terms of both dwell time and travel time) to the current starting stop of each route is added to the front of that route. The first route is processed first, followed by the second.

Table 4.2: Example of PNNH Initial Routes. Cap. is short for Capacity, DT stands for Dwell Time (minutes), TT stands for Travel Time (minutes), and JT stands for Journey Time (minutes). The routes are presented in terms of the indices of the vertices in  $V_1$ , and the subscripts represent the number of boarding students.

Step	Routes	Load/Cap.	DT	TT	JT
1	$64_2 \rightarrow 0$	2/80	0.42	9.45	9.87
	$412_{12} \rightarrow 0$	12/80	1.25	8.87	10.12
2	$253_{45} \rightarrow 64_2 \rightarrow 0$	47/80	4.42	20.67	25.09
	$353_{50} \rightarrow 412_{12} \rightarrow 0$	62/80	5.67	19.05	24.72
3	$318_{33} \rightarrow 253_{45} \rightarrow 64_2 \rightarrow 0$	80/80	7.42	23.69	31.11
	$318_{14} \rightarrow 353_{50} \rightarrow 412_{12} \rightarrow 0$	76/80	7.08	41.78	48.86

#### 4.2.2.2 Sequential Random Heuristic (SRH)

The second strategy constructs routes sequentially such that only the last  $(k - \underline{k} + 1)$  routes may have spare capacity. To start,  $k$  empty routes are defined and the remaining capacity  $c_i$  of each route  $R_i, i \in \{1, 2, \dots, k\}$ , is set to  $C_{\max}$ . A bus stop  $v \in V'_1$  is selected randomly and inserted at the end of  $R_1$ . If the number  $s_{V'_1}(v)$  of boarding students is at most  $c_1$ , then  $c_1$  is updated accordingly and another stop is selected randomly. Otherwise,  $c_1$  is exhausted and a multistop is created with the remaining  $s_{V'_1}(v) - c_1$  students and added to  $V'_1$ . Once a route is full, the procedure is repeated for the next route. This process continues until all stops are assigned to a route. In cases where  $k > \underline{k}$ , this process will

only construct  $k$  non-empty routes.

As an example, consider the data in Table 4.1 and assume  $C_{\max} = 80$ . Suppose the stops are selected in the following random sequential order: 412, 318, 353, 253, 353 (multistop), and 64. This selection yields the routes shown stepwise in Table 4.3. When Stop 353 is first selected, the load of Route 1 is 59, leaving a remaining capacity of 21. Consequently, a multistop is created to accommodate the remaining 29 students.

Table 4.3: Example of SRH Initial Routes. Cap. is short for Capacity, DT stands for Dwell Time (minutes), TT stands for Travel Time (minutes), and JT stands for Journey Time (minutes). The routes are presented in terms of the indices of the vertices in  $V_1$ , and the subscripts represent the number of boarding students.

Step	Routes	Load/Cap.	DT	TT	JT
1	412 <sub>12</sub>	12/80	1.25	-	1.25
	-	0/80	-	-	-
2	412 <sub>12</sub> → 318 <sub>47</sub>	59/80	5.42	13.25	18.67
	-	0/80	-	-	-
3	412 <sub>12</sub> → 318 <sub>47</sub> → 353 <sub>21</sub> → 0	80/80	7.42	53.28	60.70
	-	0/80	-	-	-
4	412 <sub>12</sub> → 318 <sub>47</sub> → 353 <sub>21</sub> → 0	80/80	7.42	53.28	60.70
	253 <sub>45</sub>	45/80	4.00	-	4.00
5	412 <sub>12</sub> → 318 <sub>47</sub> → 353 <sub>21</sub> → 0	80/80	7.42	53.28	60.70
	253 <sub>45</sub> → 353 <sub>29</sub>	74/80	6.67	23.48	30.15
6	412 <sub>12</sub> → 318 <sub>47</sub> → 353 <sub>21</sub> → 0	80/80	7.42	53.28	60.70
	253 <sub>45</sub> → 353 <sub>29</sub> → 64 <sub>2</sub> → 0	76/80	7.08	58.31	65.39

### 4.2.2.3 Regret Insertion Heuristic (RIH)

The third constructive heuristic is inspired by the work of Pacheco and Marti (2006), which we extend to cater for the inclusion of multistops. Like the other strategies,  $k$  empty routes are defined and the remaining capacity  $c_i$  of each route  $R_i, i \in \{1, 2, \dots, k\}$ , is set to  $C_{\max}$ . For each bus stop  $v \in V'_1$  and route  $R_i$  with positive  $c_i, i \in \{1, 2, \dots, k\}$ , we calculate the increase in  $R_i$ 's bus travel time if stop  $v$  had to be inserted in the position that causes the least travel time increase. These increases are then navigated in non-decreasing order and the bus dwell times corresponding to them are tackled as follows.

Starting from the lowest increase for stop  $v \in V'_1$ , the remaining capacity of the corresponding route  $R_{i^*}$  is checked to see whether it fits all  $s_{V'_1}(v)$  boarding students. If so,  $R_{i^*}$ 's increase plus the bus dwell time for  $s_{V'_1}(v)$  boarding students is appended to a new list  $\Omega_v$ . Otherwise,  $R_{i^*}$ 's increase plus the bus dwell time for

$c_{i^*}$  boarding students is appended to  $\Omega_v$ . To cater for the remaining  $s_{V'_1}(v) - c_{i^*}$  students, we renavigate the increases in non-decreasing order (skipping the one for  $R_{i^*}$ ) and continue until all  $s_{V'_1}(v)$  students are assigned to a route. For each multistop created, we update  $R_{i^*}$ 's entry in  $\Omega_v$  by incorporating both the bus dwell time and the increase in bus travel time.

Once the procedure is done for each increase, the list  $\Omega_v$  is sorted in non-decreasing order. For each stop  $v \in V'_1$ , the difference between the first and second entries in  $\Omega_v$  is calculated (the first entry is taken when  $|\Omega_v| = 1$  at the end of the solution construction). Such a difference is commonly referred to as a *regret value* (refer to Liu et al. (2023)) and measures the cost difference between the best and second-best insertion plans. One stop is randomly chosen for insertion from the stops having the maximum regret value since such stops should ideally be inserted first to prevent the unavailability of their better insertion plans. This stop is inserted in a route corresponding to the lowest increase in  $\Omega_v$ , breaking ties randomly, at the position that causes the least travel time increase. The whole procedure is repeated until all stops are assigned to a route.

As an example, consider the data in Table 4.1 and assume  $C_{\max} = 80$ . In Step 1, the lists  $\Omega_v$ ,  $v \in V'_1$ , are as follows:

$\Omega_{64} = [9.87 \text{ (Route 1), } 9.87 \text{ (Route 2)}]$	Regret = 0;
$\Omega_{253} = [14.88 \text{ (Route 1), } 14.88 \text{ (Route 2)}]$	Regret = 0;
$\Omega_{318} = [15.34 \text{ (Route 1), } 15.34 \text{ (Route 2)}]$	Regret = 0;
$\Omega_{353} = [21.72 \text{ (Route 1), } 21.72 \text{ (Route 2)}]$	Regret = 0;
$\Omega_{412} = [10.12 \text{ (Route 1), } 10.12 \text{ (Route 2)}]$	Regret = 0.

Since all regrets are equal, suppose Stop 412 is selected at random to initialize Route 2, with the route chosen randomly due to the tie. In Step 2, the lists  $\Omega_v$ ,  $v \in V'_1$ , are updated as follows:

$\Omega_{64} = [9.87 \text{ (Route 1), } 17.02 \text{ (Route 2)}]$	Regret = 7.15;
$\Omega_{253} = [14.88 \text{ (Route 1), } 17.42 \text{ (Route 2)}]$	Regret = 2.54;
$\Omega_{318} = [15.34 \text{ (Route 1), } 17.87 \text{ (Route 2)}]$	Regret = 2.53;
$\Omega_{353} = [14.60 \text{ (Route 2), } 21.72 \text{ (Route 1)}]$	Regret = 7.12.

Having the maximum regret value, Stop 64 is next selected to initialize Route 1. In Step 3, the lists  $\Omega_v$ ,  $v \in V'_1$ , are updated as follows:

$\Omega_{253} = [15.22 \text{ (Route 1), } 17.42 \text{ (Route 2)}]$	Regret = 2.20;
--	----------------

$$\begin{aligned}\Omega_{318} &= [17.72 \text{ (Route 1)}, 17.87 \text{ (Route 2)}] && \text{Regret} = 0.15; \\ \Omega_{353} &= [14.60 \text{ (Route 2)}, 29.80 \text{ (Route 1)}] && \text{Regret} = 15.20.\end{aligned}$$

Having the maximum regret value, Stop 353 is next selected to be placed at its best position in Route 2 (at the front of the route). In Step 4, the lists  $\Omega_v$ ,  $v \in V'_1$ , are updated as follows:

$$\begin{aligned}\Omega_{253} &= [15.22 \text{ (Route 1)}, 30.46 \text{ (Route 2)}] && \text{Regret} = 15.24; \\ \Omega_{318} &= [17.72 \text{ (Route 1)}, 33.52 \text{ (Route 2)}] && \text{Regret} = 15.80.\end{aligned}$$

In this step, both entries for Route 2 correspond to creating a multistop and inserting it into the best position in Route 1, since the remaining capacity in Route 2 is only 18. The next selected stop is Stop 318 to be placed at its best position in Route 1 (at the front of the route). In the final step, the list  $\Omega_{253}$  is updated as follows:

$$\Omega_{253} = [19.93 \text{ (Route 1)}, 19.93 \text{ (Route 2)}] \quad \text{Regret} = 0.$$

Given the tie, Route 2 is randomly selected to be filled to full capacity, and a multistop is created in Route 1 to accommodate the remaining 27 students. Table 4.4 summarizes the step-by-step application of the RIH heuristic.

Table 4.4: Example of RIH Initial Routes. Cap. is short for Capacity, DT stands for Dwell Time (minutes), TT stands for Travel Time (minutes), and JT stands for Journey Time (minutes). The routes are presented in terms of the indices of the vertices in  $V_1$ , and the subscripts represent the number of boarding students.

Step	Routes	Load/Cap.	DT	TT	JT
1	-	0/80	-	-	-
	$412_{12} \rightarrow 0$	12/80	1.25	8.87	10.12
2	$64_2 \rightarrow 0$	2/80	0.42	9.45	9.87
	$412_{12} \rightarrow 0$	12/80	1.25	8.87	10.12
3	$64_2 \rightarrow 0$	2/80	0.42	9.45	9.87
	$353_{50} \rightarrow 412_{12} \rightarrow 0$	62/80	5.67	19.05	24.72
4	$318_{47} \rightarrow 64_2 \rightarrow 0$	49/80	4.58	23.00	27.58
	$353_{50} \rightarrow 412_{12} \rightarrow 0$	62/80	5.67	19.05	24.72
5	$318_{47} \rightarrow 253_{27} \rightarrow 64_2 \rightarrow 0$	76/80	7.08	23.69	30.77
	$353_{50} \rightarrow 412_{12} \rightarrow 253_{18} \rightarrow 0$	80/80	7.42	34.04	41.46

## 4.3 Improvement Heuristics

Recall from Section 4.1 that improvement heuristics rely on neighbourhood operators to generate and evaluate local modifications to a given solution. In the context of VRPs, neighbourhood operators are classified into two categories: *intra-route operators* and *inter-route operators*. Intra-route operators make moves that affect a single route, while inter-route operators make moves that affect two or more routes simultaneously. In the literature, operators from both categories are commonly combined within the same solution method. Here, we use a combination of three intra- and three inter-route operators, which are discussed in Sections 4.3.2 and 4.3.3, respectively. Before this discussion, we review some VRP and SBRP works that employ neighbourhood operators in Section 4.3.1.

### 4.3.1 Related Works

The most trivial neighbourhood operators involve removing a node to be reinserted at a different position in the same or a different route (*node relocation* or *one-point move*) or swapping two nodes in the same or different routes (*node exchange* or *two-point move*) (Osman, 1993). Examples of SBRP works that feature node relocations and/or exchanges are those of Dulac et al. (1980), Chapleau et al. (1985), Corberán et al. (2002), Spada et al. (2005), Schittekat et al. (2013), Faraj et al. (2014), Kinable et al. (2014), Chen et al. (2015), Silva et al. (2015), Miranda et al. (2018, 2021), and Ren et al. (2019). Both operators generate a neighbourhood of size  $\mathcal{O}(N^2)$ , where  $N$  is the number of nodes.

The most common neighbourhood operators in routing are edge and chain exchanges. The most widely used intra-route edge exchange is *r-opt* which removes  $r$  edges from a route and reconnects the obtained subroutes with  $r$  other edges. The neighbourhood size of an *r-opt* operator is  $\mathcal{O}(N^r)$ . Usually,  $r$  is set to two or three to maintain low computational cost. The two-opt method was first proposed by Croes (1958) and its generalization was put forward by Lin (1965). In the context of the SBRP, two-opt was used by Newton and Thomas (1969), Desrosiers et al. (1981), Chapleau et al. (1985), Bowerman et al. (1995), Park et al. (2012), Siqueira et al. (2016), Souza Lima et al. (2016), Yao et al. (2016), Ozmen and Sahin (2021), Calvete et al. (2022), and Hou et al. (2022b), amongst others. Three-opt has not been commonly applied for the SBRP, with applications by Bennett and Gazis (1972), Orloff (1976), Bodin and Berman (1979), and Russell et al. (1986).

A popular intra-route chain exchange is *Or-opt* (Or, 1976). This moves chains of three consecutive nodes (in the same or reversed order) to a different position until no further improvement is possible, then repeats with chains of two consec-

utive nodes, and finally with single nodes. Or-opt generates a neighbourhood of size  $\mathcal{O}(N^2)$ . Notice that Or-opt exchanges generalize intra-route node relocations and are a subset of three-opt exchanges since three edges are exchanged with three others. Pacheco and Marti (2006) and Bögl et al. (2015) employed Or-opt in their SBRP frameworks. A generalization of Or-opt that tackles chains of more than three nodes exists (Babin et al., 2007). Or-opt or its generalization can also be applied between two routes by moving chains of consecutive nodes from one route to another. Here, the inter-route variant of the generalized Or-opt is termed *Or-exchange*. Both the intra- and inter-route variants of the generalized Or-opt have been used by Lewis and Smith-Miles (2018) and Dang et al. (2019). The inter-route variant also features in the work of Pacheco et al. (2013).

A widely used inter-route chain exchange is *cross-exchange* (Taillard et al., 1997), which swaps two chains of consecutive nodes between two routes, generalizing inter-route node exchange. Cross-exchange was used in the works of Pacheco and Marti (2006), Pacheco et al. (2013), Bögl et al. (2015), and Chen et al. (2015). A *two-opt\** or *crossover* move (Potvin and Rousseau, 1995) is a special case of cross-exchange, which swaps the ending subroutes of two routes. This means that arcs  $(i, j)$  and  $(i', j')$  are removed from two routes, and the nodes are reconnected as  $(i, j')$  and  $(i', j)$ . The neighbourhood size of this move is  $\mathcal{O}(N^2)$ . This was applied for instance by Dulac et al. (1980), Souza Lima et al. (2016), and Souza Lima et al. (2017). Another inter-route neighbourhood structure is  $\lambda$ -*interchange* used by Sales et al. (2018). This was originally developed by Osman (1993) for the VRP. It entails swapping two subsets of size in the range  $\{0, 1, \dots, \lambda\}$ ,  $\lambda \in \mathbb{N}$ , between two routes, with the subsets possibly having different sizes. The elements of the subsets are randomly selected, yielding a generalization of cross-exchange or Or-exchange (if one subset is empty).

We now discuss some less common operators applied in SBRP research. Pacheco et al. (2013) exchanged chains of nodes (consecutive or non-consecutive) in a single route. Souza Lima et al. (2017) merged two random routes by connecting the last node of one route to the first node of another route. They claim that such merging helps avoid artificial balancing, which occurs when routes are generated unnecessarily to balance the routes' journey times. The latter is one of their three objectives, besides minimizing the total routing and fixed costs and the total students' weighted travelling time. Souza Lima et al. (2017) also considered splitting a random route into two new routes by removing the arc closest to the middle of the route. Such a split move does not improve the total cost since an additional bus is needed, but it possibly improves the other objectives.

Campbell et al. (2015) employed a trip reduction heuristic to eliminate routes having more than half spare capacity by relocating stops to other routes. More-

over, they attempted to merge and reorder two routes with low occupancy into one route to reduce the number of buses required. A similar idea was used by Calvete et al. (2022). Shafahi et al. (2018b) also tried to reduce the number of buses by moving some stops from the shorter routes (with respect to the number of stops and the travel time) to the longer routes.

Conversely, Li and Fu (2002) tried to minimize the total travel time and balance the loads and travel times between buses by shifting students from the longer- to the shorter-distance routes. Such a shift reduces the number of stops in the longer route and, hence, the total dwell time. A similar idea was employed by Lewis and Smith-Miles (2018) who propose a *create multistop* operator to handle long routes. This operator, also used in this research, shifts a portion of the number of boarding students from a route exceeding the maximum journey time to another route, creating a multistop.

Bowerman et al. (1995) proposed a bus stop insertion heuristic to decrease the total weighted distance (walking and route distances) by adding an unvisited bus stop to a route and reassigning students to stops. An improvement is possible since the decrease in the walking distance can surpass the increase in the route distance. Likewise, Kinable et al. (2014) updated the bus stop selection during the route generation phase. They attempted to insert unvisited stops or remove stops from routes. Removal of stops and replacement of visited stops by unvisited ones were also considered by Schittekat et al. (2013) and Calvete et al. (2020).

We end this section by highlighting that several moves can be combined into one compound move, creating an *embedded neighbourhood* (Ergun et al., 2006). Examples include the ejection chain procedure developed by Glover (1996) and applied to the CVRP by Rego (2001), and the generalized insertion procedure (GENI) (Gendreau et al., 1992). GENI combines node insertions under the principle that a new node in a route is not necessarily inserted between two currently consecutive nodes. Ejection chain procedures involve linked steps in which changes generated by one move lead to the requirement of other changes in the next step. For example, inter-route node relocation might improve the solution but cause a violation of the capacity constraint. However, this may be satisfied by performing another node relocation from the over-utilized route to another route. To our knowledge, such embedded neighbourhoods have not been applied to the SBRP.

### 4.3.2 Implemented Intra-Route Operators

We now describe the neighbourhood operators used in this research, beginning with the intra-route operators. These operators are applied to a single route

$R$  in a candidate solution  $\mathcal{R}$  with the hope of shortening its journey time  $t(R)$  while maintaining the satisfaction of Constraints (3.1)-(3.3). Without loss of generality, assume that  $R = (v_1, v_2, \dots, v_{l_1}, v_0)$ . Here, we consider three intra-route operators, namely the exchange, two-opt, and generalized Or-opt operators.

#### 4.3.2.1 Exchange

The exchange (or swap) operator chooses two stops  $v_i, v_j$  in  $R$ , where  $1 \leq i < j \leq l_1$ , and swaps their position:

$$\text{Original Route : } (v_1, \dots, v_{i-1}, \underbrace{v_i}, v_{i+1}, \dots, v_{j-1}, \underbrace{v_j}, v_{j+1}, \dots, v_{l_1}, v_0)$$

$$\text{Modified Route : } (v_1, \dots, v_{i-1}, \underbrace{v_j}, v_{i+1}, \dots, v_{j-1}, \underbrace{v_i}, v_{j+1}, \dots, v_{l_1}, v_0).$$

The neighbourhood size corresponding to the exchange operator is  $\mathcal{O}(|V_1'|^2)$ . This follows from the fact that a route contains at most  $|V_1'|$  stops and each exchange move is defined by a pair of distinct indices  $(i, j)$ , yielding a quadratic number of possible exchanges.

#### 4.3.2.2 Two-Opt

The two-opt operator chooses two stops  $v_i, v_j$  in  $R$ , where  $1 \leq i < i + 3 \leq j \leq l_1$ , and inverts the subroute  $v_i, \dots, v_j$ :

$$\text{Original Route : } (v_1, \dots, v_{i-1}, \underbrace{v_i, v_{i+1}, \dots, v_{j-1}, v_j}, v_{j+1}, \dots, v_{l_1}, v_0)$$

$$\text{Modified Route : } (v_1, \dots, v_{i-1}, \underbrace{v_j, v_{j-1}, \dots, v_{i+1}, v_i}, v_{j+1}, \dots, v_{l_1}, v_0).$$

The cases  $j = i + 1$  and  $j = i + 2$  are excluded since these would yield swap moves already considered in the exchange operator. The neighbourhood size corresponding to the two-opt operator is  $\mathcal{O}(|V_1'|^2)$ , since each two-opt move is defined by a pair of distinct indices  $(i, j)$  and, therefore, the number of admissible pairs grows quadratically with  $|V_1'|$ .

#### 4.3.2.3 Generalized Or-Opt

The generalized Or-opt operator first chooses three stops  $v_h, v_i, v_j$  in  $R$ , where  $1 \leq i \leq j \leq l_1$  and  $(1 \leq h < i$  or  $j + 1 < h \leq l_1 + 1)$ . Then it removes the subroute  $v_i, \dots, v_j$  and transfers it before stop  $v_h$ , where  $h = l_1 + 1$  corresponds to the school  $v_0$ . The subroute may also be inverted if this yields a higher improvement. The

following is an example with  $h > j + 1$  and no inversion:

$$\begin{aligned} \text{Original Route : } & (v_1, \dots, v_{i-1}, \underbrace{v_i, \dots, v_j}_{}, v_{j+1}, \dots, v_{h-1}, \underbrace{v_h}_{}, v_{h+1}, \dots, v_{l_1}, v_0) \\ \text{Modified Route : } & (v_1, \dots, v_{i-1}, v_{j+1}, \dots, v_{h-1}, \underbrace{v_i, \dots, v_j}_{}, \underbrace{v_h}_{}, v_{h+1}, \dots, v_{l_1}, v_0). \end{aligned}$$

The neighbourhood size corresponding to the generalized Or-opt operator is  $\mathcal{O}(|V_1'|^3)$ . Each generalized Or-opt move is defined by a triple of indices  $(h, i, j)$  corresponding to the insertion position and the endpoints of the relocated subroute. The number of admissible index triples therefore grows cubically with  $|V_1'|$ .

Observe that for a given route, the intra-route operators do not alter the number of visited stops, the number of boarding students, and the dwell time at each stop.

### 4.3.3 Implemented Inter-Route Operators

Here, inter-route operators are applied to a pair of routes  $R, R' \in \mathcal{R}$  to try to decrease the sum of their costs while ensuring that Constraints (3.1)-(3.3) are still satisfied. Without loss of generality, assume that  $R = (v_1, v_2, \dots, v_{l_1}, v_0)$  and  $R' = (u_1, u_2, \dots, u_{l_2}, v_0)$ . Here, we apply three inter-route operators, these being the Or-exchange, cross-exchange, and create multistop operators.

#### 4.3.3.1 Or-Exchange

The Or-exchange operator takes stops  $v_i, v_j$  in  $R$ , where  $1 \leq i \leq j \leq l_1$ , and stop  $u_h$  in  $R'$ , where  $1 \leq h \leq l_2 + 1$ . Then it removes the subroute  $v_i, \dots, v_j$  from  $R$  and transfers it before stop  $u_h$  in  $R'$ , where  $h = l_2 + 1$  corresponds to the school  $v_0$ . The subroute may also be inverted if this yields a higher improvement. The following shows an example with inversion:

$$\begin{aligned} \text{Original Routes : } & (v_1, \dots, v_{i-1}, \underbrace{v_i, v_{i+1}, \dots, v_{j-1}, v_j}_{}, v_{j+1}, \dots, v_{l_1}, v_0) \\ & (u_1, \dots, u_{h-1}, \underbrace{u_h}_{}, u_{h+1}, \dots, u_{l_2}, v_0) \\ \text{Modified Routes : } & (v_1, \dots, v_{i-1}, v_{j+1}, \dots, v_{l_1}, v_0) \\ & (u_1, \dots, u_{h-1}, \underbrace{v_j, v_{j-1}, \dots, v_{i+1}, v_i}_{}, \underbrace{u_h}_{}, u_{h+1}, \dots, u_{l_2}, v_0). \end{aligned}$$

The neighbourhood size corresponding to the Or-exchange operator is  $\mathcal{O}(|V_1'|^3)$ , following the same reasoning as for the generalized Or-opt operator, since each move is defined by a subroute in one route and an insertion position in another.

Note that an Or-exchange move can cause a violation of the capacity con-

straint (3.3). If  $k > \underline{k}$  and  $R'$  includes only the school, it can also cause a shift of all the stops from  $R$  to  $R'$ , leading to an equivalent solution. In both cases, such a move is ignored. Moreover, an Or-exchange move can result in duplicate stops in the same route. Although a stop may be repeated in different routes, the repetition of the same stop in one route is illogical. Therefore, duplicate stops in the same route are removed as follows. Without loss of generality, assume that subroute  $v_i, \dots, v_j$  is being transferred from route  $R$  to route  $R'$  and that one stop  $v_g, i \leq g \leq j$ , is already present in  $R'$ . Then stop  $v_g$  is removed from the subroute and the students boarding this occurrence of  $v_g$  are all transferred to the occurrence of  $v_g$  in  $R'$ . The same arguments regarding capacity violations, equivalent solutions, and the removal of duplicate stops within a single route also apply to the second inter-route operator, the cross-exchange operator.

#### 4.3.3.2 Cross-Exchange

The cross-exchange operator first chooses stops  $v_{i_1}, v_{j_1}$  in  $R$ , where  $1 \leq i_1 \leq j_1 \leq l_1$ , and stops  $u_{i_2}, u_{j_2}$  in  $R'$ , where  $1 \leq i_2 \leq j_2 \leq l_2$ . Subsequently, it swaps subroutes  $v_{i_1}, \dots, v_{j_1}$  and  $u_{i_2}, \dots, u_{j_2}$ , possibly inverting either subroute if this yields a higher improvement. The following is an example with inversion in the subroute being transferred to  $R$  and no inversion in that being transferred to  $R'$ :

$$\begin{aligned}
 \text{Original Routes : } & (v_1, \dots, v_{i_1-1}, \underbrace{v_{i_1}, v_{i_1+1}, \dots, v_{j_1-1}, v_{j_1}}_{\text{subroute}}, v_{j_1+1}, \dots, v_{l_1}, v_0) \\
 & (u_1, \dots, u_{i_2-1}, \underbrace{u_{i_2}, u_{i_2+1}, \dots, u_{j_2-1}, u_{j_2}}_{\text{subroute}}, u_{j_2+1}, \dots, u_{l_2}, v_0) \\
 \\ 
 \text{Modified Routes : } & (v_1, \dots, v_{i_1-1}, \underbrace{u_{j_2}, u_{j_2-1}, \dots, u_{i_2+1}, u_{i_2}}_{\text{subroute}}, v_{j_1+1}, \dots, v_{l_1}, v_0) \\
 & (u_1, \dots, u_{i_2-1}, \underbrace{v_{i_1}, v_{i_1+1}, \dots, v_{j_1-1}, v_{j_1}}_{\text{subroute}}, u_{j_2+1}, \dots, u_{l_2}, v_0).
 \end{aligned}$$

The neighbourhood size corresponding to the cross-exchange operator is  $\mathcal{O}(|V_1|^4)$ . This follows from the fact that each cross-exchange move is defined by the end-point pairs of two subroutes,  $(i_1, j_1)$  in  $R$  and  $(i_2, j_2)$  in  $R'$ , and the number of admissible index quadruples grows quartically with  $|V_1|$ .

#### 4.3.3.3 Create Multistop

The create multistop operator is employed when two routes  $R$  and  $R'$  satisfy  $t(R) > m_t$  and  $s(R') < C_{\max}$ , i.e.,  $R$  has journey time exceeding the maximum journey time and  $R'$  has at least one empty seat. This operator starts by choosing stop  $v_i$  in  $R$ , where  $1 \leq i \leq l_1$ , for which the number of boarding students  $s(v_i, R) \geq 2$ . If  $v_i$  is not already in  $R'$ , then a copy of  $v_i$  is inserted in  $R'$  before

the stop  $u_h$ , where  $1 \leq h \leq l_2$ , (or the school  $v_0$ ) that causes the smallest increase in  $t(R')$ . Next, the operator transfers  $z := \min\{s(v_i, R) - 1, C - s(R')\}$  students from the occurrence of stop  $v_i$  in  $R$  to the occurrence of stop  $v_i$  in  $R'$ . Here, the value  $z$  gives the maximum number of students who can be transferred (hence, decreasing  $t(R)$  as much as possible) such that both occurrences of  $v_i$  have at least one boarding student and both  $R$  and  $R'$  satisfy the capacity constraint (3.3). The following is an example with a copy of  $v_i$  being inserted in  $R'$ :

$$\begin{aligned}
 \text{Original Routes : } & (v_1, \dots, v_{i-1}, \underbrace{v_i}_{s(v_i, R) \text{ boarding students}}, v_{i+1}, \dots, v_{l_1}, v_0) \\
 & (u_1, \dots, u_{h-1}, \underbrace{u_h}_{s(v_i, R) \text{ boarding students}}, u_{h+1}, \dots, u_{l_2}, v_0) \\
 \text{Modified Routes : } & (v_1, \dots, v_{i-1}, \underbrace{v_i}_{(s(v_i, R) - z) \text{ boarding students}}, v_{i+1}, \dots, v_{l_1}, v_0) \\
 & (u_1, \dots, u_{h-1}, \underbrace{v_i}_{z \text{ boarding students}}, \underbrace{u_h}_{z \text{ boarding students}}, u_{h+1}, \dots, u_{l_2}, v_0).
 \end{aligned}$$

The neighbourhood size corresponding to the create multistop operator is  $\mathcal{O}(|V_1'|^2)$ . Each move is defined by a pair of indices  $(h, i)$ , representing the insertion position in route  $R'$  and the stop in route  $R$ , so the number of admissible pairs grows quadratically with  $|V_1'|$ .

### 4.3.4 Application of Neighbourhood Operators

The six neighbourhood operators described in Sections 4.3.2 and 4.3.3 are used in combination as follows. All possible applications of these six operators that do not cause violations of the capacity constraints (3.3) are evaluated. The move that yields the greatest improvement in cost is then performed, following the direction of *steepest descent* (SD). This strategy is also employed in other SBRP works such as those of Pacheco and Marti (2006) and Lewis and Smith-Miles (2018).

In our case, if multiple moves within the union of the neighbourhoods give the best cost improvement, the one that results in the smallest discrepancy between the longest and shortest routes is performed. As remarked in Section 3.2.3, such breakage of ties aims to balance journey times across routes. The process of evaluating the neighbourhoods and selecting the best move is repeated until a local optimum is reached – i.e., a solution for which no improving move exists in any of the neighbourhoods.

An alternative and more common direction that can be followed is *variable neighbourhood descent* (VND) (Mladenović and Hansen, 1997). In VND, the different neighbourhoods are considered sequentially, in the order defined

by a prespecified permutation of the neighbourhoods (often sorted in increasing neighbourhood size). If a neighbourhood fails to find an improving move, the next neighbourhood is applied. Otherwise, the first-improving or best-improving move in the current neighbourhood is executed. Examples of SBRP works that employ VND are those by Euchi and Mraïhi (2012), Kim et al. (2012), Schittekat et al. (2013), Bögl et al. (2015), Campbell et al. (2015), and Chen et al. (2015).

A random ordering of the neighbourhoods instead of a deterministic ordering gives rise to random VND (RVND). The latter was used by Souza Lima et al. (2016), Souza Lima et al. (2017), and Miranda et al. (2018, 2021), amongst others. As will be seen in the computational results in Section 4.6, RVND considerably reduces the computational effort for performing one move and thus speeds up the search process. However, it does not always reach local optima that are as good as those of SD.

## 4.4 Perturbation Mechanisms

An algorithm that takes an initial solution (generated by a constructive heuristic such as those discussed in Section 4.2, or manually using prior knowledge and experience) and improves it via the application of neighbourhood operators (such as those in Section 4.3), is usually called a *local search algorithm*. The issue with such an algorithm is that it explores only the solutions in the combined neighbourhoods of the current solution and lacks a diversification strategy. Therefore, it typically converges to a local optimum that is generally not a global optimum (Crama et al., 1995), as the search is confined to a limited portion of the solution space.

The ILS metaheuristic provides a remedy for this issue. The idea of this metaheuristic is that “one iteratively builds a sequence of solutions generated by the embedded heuristic, leading to far better solutions than if one were to use repeated random trials of that heuristic” (Lourenço et al., 2003). In essence, ILS includes a perturbation mechanism that diversifies the output of one run of a local search algorithm before the latter is restarted, extending the local search algorithm to an iterative one. This idea has a long history, with first approaches by Baxter (1981) to a depot-location problem and Baum (1986) to a TSP. An overview of the early historical developments of ILS can be found in Johnson and McGeoch (1997).

The question we ask is “How can one achieve a sequence of solutions?”. The simplest strategy would be to repeat the search from a different initial solution, generating a set of independent local optima. However, such random restarts increase computational time and have a lower probability of finding good solutions

as the instance size increases (Lourenço et al., 2003). Therefore, for better overall performance, the standard approach is to apply perturbations to a current local optimum  $s^*$  to achieve another solution  $s'$  and then improve  $s'$  to reach another local optimum  $s^{*'}.$  The perturbations should neither be too small nor too large. In the former case, one will likely fall back to  $s^*$ , leading to limited diversification, while the latter behaves like a random restart. As suggested by Lourenço et al. (2003), perturbations should be randomized or adaptive (for example, depending on any of the previous  $s^*$ ) and not deterministic to avoid short solution cycles.

The acceptance or rejection of the new local optimum  $s^{*'}$  is another important factor in ILS. One may incorporate an acceptance criterion against which  $s^{*'}$  is checked. If  $s^{*'}$  satisfies this criterion, it becomes the current local optimum  $s^*$ ; otherwise, one returns to  $s^*$ . The acceptance criterion strongly influences the search, affecting the balance between intensification and diversification. For instance, accepting  $s^{*'}$  only if it is better than  $s^*$  achieves strong intensification, while always accepting  $s^{*'}$  achieves strong diversification. It is possible to choose an acceptance criterion between these two extreme examples. For instance, an SA type acceptance criterion, like the ones used in Spada et al. (2005), Chen et al. (2015), Parvasi et al. (2017), Hu et al. (2018), Shafahi et al. (2018a,b), and Hou et al. (2022a), always accepts  $s^{*'}$  if it is better than  $s^*$ . Otherwise,  $s^{*'}$  is typically accepted with probability  $\exp\{(f(s^*) - f(s^{*'}))/\tau\}$ , where  $f$  is the cost function to be minimized and  $\tau$  is a parameter called temperature. Initially, the temperature and, in turn, the probability of accepting a worse solution is high, focusing on diversification. Along the process, the temperature is gradually lowered based on a cooling schedule scheme to shift the focus to intensification.

Given that our S-SBRP embeds both the bus stop selection and route generation subproblems, the perturbation mechanisms considered in this research alter not only the routes in a current solution  $\mathcal{R}$ , but also the set of visited bus stops  $V_1'$ . Notice that  $V_1'$  remains fixed when applying the improvement heuristics presented in Section 4.3. For a fixed  $V_1'$ , it may be the case that no solution satisfying the maximum journey time constraints (3.4) is achieved after applying the improvement heuristics. It may also be the case that different subsets of bus stops will lead to better solutions. Hence, our algorithm contains a destroy-and-repair operator that generates a new subset of bus stops  $V_1''$  from  $V_1'$ , assigns students to these bus stops, and then repairs and optimizes the current routes accordingly.

We consider four perturbation mechanisms, which differ in the way they generate  $V_1''$ . These are:

- (PM 1) Generate  $V_1''$  from scratch (random restart);
- (PM 2) Generate  $V_1''$  from the subset  $V_1'$  used in the previous iteration (focuses on diversification);

- (PM 3) Generate  $V_1''$  from the most recent subset  $V_1'$  that yielded a feasible solution with the lowest cost found so far (focuses on intensification);
- (PM 4) Generate  $V_1''$  via a trade-off between PM 2 and 3, whereby  $V_1''$  has a 50% chance of being generated according to PM 2 and a 50% chance of being generated according to PM 3.

Note that in PM 3, the subset of stops generated in the previous iteration is used if no subset has yet yielded a feasible solution.

In PM 1, the generation of  $V_1''$  and the construction of the routes follow the same strategies as those discussed in Section 4.2. For the other mechanisms, the non-compulsory stops in  $V_1'$  are identified and a random selection of these is removed. Assuming a total number  $\eta$  of non-compulsory stops, the number of removals is selected according to a Binomial distribution with parameters  $\eta$  and  $3/\eta$ . This is done so that, on average, a moderate amount of three stops are removed. Upon removal, if we have an incomplete covering of  $V_2$ , additional stops must be added to  $V_1'$ . If all addresses not covered by the stops in  $V_1'$  are covered by stops which were not originally in  $V_1'$ , then, at each stage, a stop from the latter set of stops which serves the largest number of uncovered addresses is added, breaking ties randomly. If, on the other hand, some address is also uncovered by the stops that were not originally in  $V_1'$ , then at least one of the removed stops must be added back. The same selection strategy is applied in this case and the whole procedure is repeated until a new complete covering  $V_1''$  of  $V_2$  is achieved. Each address is then reassigned to the closest stop in  $V_1''$  and stops with no addresses assigned to them are removed from  $V_1''$ .

After determining a new subset of bus stops, repairs must be made to  $\mathcal{R}$  so that only bus stops in  $V_1''$  feature in the solution. To do this, all occurrences of stops in  $V_1' \setminus V_1''$  are first removed from  $\mathcal{R}$ . For stops  $v \in V_1'' \cap V_1'$  for which  $s_{V_1''}(v) < s_{V_1'}(v)$ ,  $s_{V_1'}(v) - s_{V_1''}(v)$  students are removed from occurrences of  $v$  in  $\mathcal{R}$ , processing the occurrences in non-increasing order of boarding count at  $v$ . If this results in an occurrence of  $v$  with no boarding students, then this occurrence is removed from  $\mathcal{R}$ . For stops  $v \in V_1'' \cap V_1'$  for which  $s_{V_1''}(v) > s_{V_1'}(v)$ , an attempt is made to add students to occurrences of  $v$  in  $\mathcal{R}$ . If not all  $s_{V_1''}(v) - s_{V_1'}(v)$  students can be added, then a new occurrence of  $v$  must be added to  $\mathcal{R}$ . Stops  $v \in V_1'' \setminus V_1'$  must also be added to the solution. In our case, a new stop is inserted in a route having the lowest load, at the position which causes the least increase in the route journey time. If this insertion does not cater for all students boarding that stop, the procedure is repeated.

Having repaired solution  $\mathcal{R}$  (or generated new routes in the case of PM 1), the neighbourhood operators are reapplied. This destroy-repair-and-improve process is repeated for a predefined number of iterations, discussed further in Section 4.5.

## 4.5 The Algorithmic Framework

In this section, we provide an overview of our heuristic algorithm, whose pseudocode is presented in Algorithm 1. To begin, a subset of bus stops  $V_1'$  that covers  $V_2$  is selected (Step 2) and a constructive heuristic is employed to construct an initial solution  $\mathcal{R}$  using a fixed number  $k$  of routes (Step 3). As mentioned in Chapter 3,  $k$  is initially taken to be the lower bound  $\underline{k}$  (Step 1). Recall also from Section 4.2 that the initial assignment of stops to routes allows the violation of the maximum journey time constraints (3.4).

A local search routine involving the six neighbourhood operators described in Section 4.3 is then invoked on  $\mathcal{R}$  to attempt to shorten the routes that use the current subset of bus stops  $V_1'$  (Step 6). Recall that this routine may follow the direction of SD, VND, or RVND. Once the local search routine is complete, the current subset  $V_1'$  of visited bus stops is altered and the current solution  $\mathcal{R}$  is repaired, as discussed in Section 4.4 (Step 7). The local search routine is reapplied and the process is repeated for  $T$  iterations (Steps 4, 5, 8, and 9). If no solution satisfying both Constraints (3.3) and (3.4) is found in  $T$  iterations, the number  $k$  of routes is incremented by one, and the algorithm is restarted (Steps 10 to 13).

---

### Algorithm 1 S-SBRP Heuristic Algorithm

---

- 1: Set the number of routes  $k = \underline{k} = \lceil \sum_{w \in V_2} s(w) / C_{\max} \rceil$ .
  - 2: Select an initial subset of bus stops  $V_1' \subseteq V_1$  that covers  $V_2$ .
  - 3: Construct an initial solution  $\mathcal{R}$  with  $k$  routes via a constructive heuristic. This solution must satisfy Constraints (3.3) but is permitted to violate Constraints (3.4).
  - 4: *counter*  $\leftarrow$  0.
  - 5: **while** *counter*  $<$   $T$  **do**  $\triangleright T$  is a predefined number of iterations
  - 6:     Apply improvement heuristics to  $\mathcal{R}$  to improve the total journey time of its routes.
  - 7:     Alter the subset of visited bus stops  $V_1'$  and repair  $\mathcal{R}$  accordingly.
  - 8:     *counter*  $\leftarrow$  *counter* + 1.
  - 9: **end while**
  - 10: **if** no feasible solution satisfying both Constraints (3.3) and (3.4) is found **then**
  - 11:      $k \leftarrow k + 1$ .
  - 12:     **goto** Step 2.
  - 13: **end if**
-

## 4.6 Computational Results

The computational experiments were performed on the twenty problem instances summarized in Tables 3.1 and 3.2. The C++ source code can be downloaded from (Sciortino, 2024b). Twelve variants of the algorithm were run 25 times each, with each variant composed of one constructive heuristic (PNNH, SRH, or RIH) and one perturbation mechanism (PM 1, 2, 3, or 4).

The number of iterations per run,  $T$ , was determined individually for each instance to ensure fair and comparable results. Its value was based on the PNNH heuristic, which was the first constructive heuristic we employed. In preliminary time-based runs of our algorithm – each limited to five minutes – we observed that perturbation mechanisms PM 3 and PM 4 consistently produced statistically significantly better solutions than PM 1 and PM 2 for most instances. To select the number  $T$  of iterations for better comparability throughout the study, we ran a more sophisticated version of our algorithm, described in Chapter 5, fifty times under a five-minute time limit per run: 25 runs with PM 3 and 25 with PM 4. For each run, the total number of iterations completed was recorded. The average number of iterations across these fifty runs was then computed and rounded to the nearest ten, defining the value of  $T$  for that instance. The  $T$  values are presented in the second column of Table 4.5, which summarizes the results. This approach was chosen for two main reasons. First, by basing  $T$  on the actual behaviour of the algorithm under controlled five-minute runs, we ensured that the iteration counts reflected realistic computational effort for each instance. Second, a five-minute per-run limit was selected so that each instance would require approximately 2-3 hours of computation across all runs, balancing adequate exploration of the solution space with practical runtime considerations.

From Column 4 of Table 4.5, we see that solutions using the lower bound of  $k$  routes, and having journey times not exceeding 45 minutes in all cases, were found for nineteen of the twenty instances in all runs. The only instance that required one additional route beyond the lower bound in all runs was the Bridgend instance. This is probably because the pairwise travel times have relatively high percentiles, as seen in Table 3.2, leading to a great chance of violating the maximum riding time constraints (3.4). In particular, this problem instance involves several geographically dispersed clusters of student addresses and others located near two parallel and distant roads, one of which is a dead-end running into the South-Wales mountains.

From the last three columns of Table 4.5 and the boxplots displaying the TJT values for the twelve algorithm variants (found in Appendix B.1), it is clear that PM 3 and PM 4 perform much better than PM 1 and PM 2, and that SRH and

Table 4.5: Results are averaged across 25 runs. For each constructive heuristic, the best total journey time (TJT) average (in minutes) across all perturbation mechanisms is presented, with the mechanism number(s) in superscript. Bold results represent the overall best averages.

Location	$T$	$k$	$k$	Best TJT Average		
				PNNH	SRH	RIH
Mgarr	85,450	4	4	54.10 <sup>1,2,3,4</sup>	<b>54.03</b> <sup>1</sup>	54.09 <sup>1</sup>
Mellieħa	43,960	4	4	<b>56.30</b> <sup>3,4</sup>	<b>56.30</b> <sup>3,4</sup>	<b>56.30</b> <sup>3,4</sup>
Porthcawl	284,260	1	1	<b>27.88</b> <sup>4</sup>	<b>27.88</b> <sup>4</sup>	<b>27.88</b> <sup>4</sup>
Qrendi	72,300	5	5	<b>75.73</b> <sup>4</sup>	<b>75.73</b> <sup>4</sup>	<b>75.73</b> <sup>4</sup>
Suffolk	9,380	3	3	<b>116.14</b> <sup>3</sup>	116.70 <sup>4</sup>	116.92 <sup>3</sup>
Senglea	70,750	6	6	72.55 <sup>3</sup>	72.44 <sup>4</sup>	<b>72.43</b> <sup>4</sup>
Victoria	5,980	4	4	93.04 <sup>3</sup>	<b>92.93</b> <sup>3</sup>	93.03 <sup>3</sup>
Pembroke	32,610	7	7	<b>103.53</b> <sup>3</sup>	103.71 <sup>3,4</sup>	103.55 <sup>3</sup>
Canberra	6,610	7	7	179.63 <sup>3</sup>	179.55 <sup>3</sup>	<b>179.46</b> <sup>3</sup>
ħandaq	26,020	6	6	96.91 <sup>3</sup>	96.84 <sup>3</sup>	<b>96.72</b> <sup>3</sup>
Valletta	32,090	6	6	104.69 <sup>3</sup>	<b>104.65</b> <sup>3</sup>	<b>104.65</b> <sup>4</sup>
Birkirkara	43,580	6	6	97.85 <sup>4</sup>	97.94 <sup>3</sup>	<b>97.56</b> <sup>4</sup>
ħamrun	21,620	7	7	<b>99.67</b> <sup>4</sup>	99.69 <sup>4</sup>	99.92 <sup>4</sup>
Cardiff	82,320	2	2	<b>67.12</b> <sup>1,2,3,4</sup>	<b>67.12</b> <sup>1,2,3,4</sup>	<b>67.12</b> <sup>1,2,3,4</sup>
Milton Keynes	63,740	4	4	59.86 <sup>3</sup>	<b>59.78</b> <sup>4</sup>	59.88 <sup>3</sup>
Bridgend	21,260	5	6	170.43 <sup>4</sup>	170.72 <sup>4</sup>	<b>170.34</b> <sup>4</sup>
Edinburgh-2	19,430	4	4	<b>59.13</b> <sup>1,2,4</sup>	<b>59.13</b> <sup>1,2,4</sup>	<b>59.13</b> <sup>1,2,4</sup>
Edinburgh-1	31,390	9	9	144.49 <sup>4</sup>	144.26 <sup>3</sup>	<b>144.22</b> <sup>3</sup>
Adelaide	25,070	8	8	128.57 <sup>4</sup>	<b>128.06</b> <sup>3</sup>	128.52 <sup>3</sup>
Brisbane	15,120	10	10	213.25 <sup>3</sup>	<b>212.43</b> <sup>3</sup>	213.01 <sup>3</sup>

RIH are, overall, better than PNNH. Statistically significant differences between the variants were confirmed for each instance by an independent-samples Kruskal-Wallis test ( $p < 0.001$ , refer to Appendix B.2). Based on these observations, the choice of the best-performing variant was narrowed down to SRH-3 (SRH and PM 3), SRH-4 (SRH and PM 4), RIH-3 (RIH and PM 3), and RIH-4 (RIH and PM 4). For both SRH and RIH, a Mann-Whitney test was then performed to determine the better perturbation mechanism for each instance. The results of these tests can be found in Columns 2 to 9 of Table 4.6. Statistically significant differences at the 0.05 level were only found for the Porthcawl and Edinburgh-2 instances, and also the Senglea instance in the case of SRH. The mean ranks corresponding to the significant  $p$ -values highlight that the SRH-4 results are significantly lower (better) than those for SRH-3, the same holding for RIH. Therefore, the final selection was between SRH-4 and RIH-4. Mann-Whitney test results (Columns 10 to 13 of Table 4.6) did not reveal statistically significant differences between SRH-4 and RIH-4. However, the SRH-4 mean rank was found to be better or equally good for sixteen instances. For this reason, SRH-4 was selected as the best-performing variant to continue with.

Table 4.6: Mann-Whitney test results. The values under the SRH-3, SRH-4, RIH-3, and RIH-4 columns represent the mean ranks,  $U$  represents the test statistic, while bold  $p$ -values are significant at the 0.05 level.

Location	SRH-3 vs SRH-4			RIH-3 vs RIH-4			SRH-4 vs RIH-4					
	SRH-3	SRH-4	$U$	RIH-3	RIH-4	$U$	SRH-4	RIH-4	$U$	$p$ -value		
Mgarr	25.50	25.50	312.5	1.000	25.50	25.50	312.5	1.000	25.50	25.50	312.5	1.000
Mellieha	25.50	25.50	312.5	1.000	25.50	25.50	312.5	1.000	25.50	25.50	312.5	1.000
Porthcawl	36.68	14.32	33.0	< <b>0.001</b>	36.68	14.32	33.0	< <b>0.001</b>	25.50	25.50	312.5	1.000
Qrendi	26.00	25.00	300.0	0.317	26.00	25.00	300.0	0.317	25.50	25.50	312.5	1.000
Suffolk	24.48	26.52	287.0	0.570	22.96	28.04	249.0	0.197	22.58	28.42	239.5	0.142
Senglea	30.42	20.58	189.5	<b>0.013</b>	29.38	21.62	215.5	0.052	25.20	25.80	305.0	0.881
Victoria	22.64	28.36	241.0	0.138	22.32	28.68	233.0	0.103	23.34	27.66	258.5	0.271
Pembroke	25.60	25.40	310.0	0.961	23.84	27.16	271.0	0.419	25.46	25.54	311.5	0.984
Canberra	21.76	29.24	219.0	0.069	21.64	29.36	216.0	0.061	26.18	24.82	295.5	0.741
Handaq	24.50	26.50	287.5	0.627	21.92	29.08	223.0	0.082	25.28	25.72	307.0	0.915
Valletta	25.96	25.04	301.0	0.820	27.26	23.74	268.5	0.381	26.42	24.58	289.5	0.650
Birkirkara	24.50	26.50	287.5	0.627	25.94	25.06	301.5	0.831	27.88	23.12	253.0	0.248
Hamrun	25.86	25.14	303.5	0.859	25.80	25.20	305.0	0.884	23.96	27.04	274.0	0.450
Cardiff	25.50	25.50	312.5	1.000	25.50	25.50	312.5	1.000	25.50	25.50	312.5	1.000
Milton Keynes	27.94	23.06	251.5	0.178	24.66	26.34	291.5	0.647	23.50	27.50	262.5	0.253
Bridgend	27.04	23.96	274.0	0.442	24.82	26.18	295.5	0.729	27.50	23.50	262.5	0.314
Edinburgh-2	37.50	13.50	12.5	< <b>0.001</b>	37.00	14.00	25.0	< <b>0.001</b>	25.50	25.50	312.5	1.000
Edinburgh-1	24.94	26.06	298.5	0.786	23.50	27.50	262.5	0.332	24.56	26.44	289.0	0.648
Adelaide	22.90	28.10	247.5	0.205	25.16	25.84	304.0	0.869	24.98	26.02	299.5	0.800
Brisbane	23.90	27.10	272.5	0.438	23.84	27.16	271.0	0.421	24.20	26.80	280.0	0.528

The results for SRH-4 are given in Table 4.7, while those for the other algorithm variants can be found in Appendix B.3. Detailed results for each run of each variant are available at (Sciortino, 2024a). A point worth noting is that PM 1 requires a considerably longer computational time because this perturbation mechanism does not use information from previous iterations when altering the current subset of bus stops. Hence, local search takes longer in each iteration as it operates on a completely new set of routes. Another observation is that PM 4 is overall slightly faster than PM 3. The reason for this is that 50% of the time, PM 4 uses the subset of bus stops from the previous iteration. In this case, local search takes a shorter time than when the subset corresponding to the best feasible solution is used because fewer changes are made.

From Table 4.7, it is evident that our algorithm has a high average feasibility rate, meaning that it finds feasible solutions in most of the iterations performed. Constraints (3.1)-(3.4) were satisfied in all runs for sixteen instances. For the Suffolk instance, a notably low average feasibility rate (18.59%) was observed.

Table 4.7: PM 4 results for SRH. Results are averaged across 25 runs and presented as mean  $\pm$  standard deviation. Time gives the computational time (seconds), FR stands for Feasibility Rate (percentage of iterations yielding a solution with all route journey times being at most 45 minutes), and TJT stands for Total Journey Time (minutes). Bold TJT values represent the best averages across all SRH perturbation mechanisms.

Location	SRH-4			
	Time (s)	FR (%)	TJT (mins.)	
Mgarr	166.50 $\pm$ 1.49	100.00 $\pm$ 0.00	54.10 $\pm$ 0.00	
Mellicha	150.85 $\pm$ 28.33	100.00 $\pm$ 0.00	<b>56.30</b> $\pm$ 0.00	
Porthcawl	637.15 $\pm$ 15.19	99.98 $\pm$ 0.03	<b>27.88</b> $\pm$ 0.38	
Qrendi	183.83 $\pm$ 2.52	100.00 $\pm$ 0.00	<b>75.73</b> $\pm$ 0.00	
Suffolk	69.55 $\pm$ 6.73	18.79 $\pm$ 4.34	<b>116.70</b> $\pm$ 0.99	
Senglea	232.94 $\pm$ 7.96	100.00 $\pm$ 0.00	<b>72.44</b> $\pm$ 0.39	
Victoria	167.60 $\pm$ 7.75	100.00 $\pm$ 0.00	93.02 $\pm$ 0.44	
Pembroke	218.53 $\pm$ 25.92	100.00 $\pm$ 0.00	<b>103.71</b> $\pm$ 0.98	
Canberra	103.04 $\pm$ 10.07	100.00 $\pm$ 0.00	180.01 $\pm$ 1.11	
Handaq	232.80 $\pm$ 24.18	100.00 $\pm$ 0.00	96.90 $\pm$ 0.72	
Valletta	321.41 $\pm$ 11.24	100.00 $\pm$ 0.00	104.77 $\pm$ 0.45	
Birkirkara	345.90 $\pm$ 19.12	100.00 $\pm$ 0.00	98.09 $\pm$ 1.50	
Hamrun	260.62 $\pm$ 76.85	100.00 $\pm$ 0.00	<b>99.69</b> $\pm$ 1.18	
Cardiff	733.17 $\pm$ 119.16	66.79 $\pm$ 0.32	<b>67.12</b> $\pm$ 0.00	
Milton Keynes	537.76 $\pm$ 33.20	100.00 $\pm$ 0.00	<b>59.78</b> $\pm$ 0.56	
Bridgend	218.93 $\pm$ 10.28	99.37 $\pm$ 0.12	<b>170.72</b> $\pm$ 0.79	
Edinburgh-2	540.81 $\pm$ 102.37	100.00 $\pm$ 0.00	<b>59.13</b> $\pm$ 0.00	
Edinburgh-1	297.83 $\pm$ 51.08	100.00 $\pm$ 0.00	144.38 $\pm$ 1.77	
Adelaide	326.73 $\pm$ 96.99	100.00 $\pm$ 0.00	128.47 $\pm$ 1.45	
Brisbane	323.09 $\pm$ 53.18	100.00 $\pm$ 0.00	213.07 $\pm$ 4.04	

Lewis and Smith-Miles (2018) also encountered feasibility issues for this rural-based instance.

Sensitivity analysis was performed to assess the impact of each neighbourhood operator on the algorithm’s overall performance. This was conducted by removing one of the six operators at a time and comparing the algorithm’s performance without one operator to that obtained with all operators. Table 4.8 presents the relative percentage increase in the TJT average for each instance-operator pair. For each such pair, we also show the  $p$ -value resulting from a paired-samples non-parametric test that we performed to test for a statistically significant increase in the TJT median upon removal of the operator (one-tailed test). A non-parametric test was used since it yields more robust results than its parametric equivalent for a small sample size like ours (size 25). The test performed was a Wilcoxon signed-rank test when the distribution of the differences between the two result sets was approximately symmetrical (skewness between -1 and 1). Otherwise, a sign test was performed. The latter does not rely on distributional assumptions but is less powerful than the Wilcoxon signed-rank test because it only uses the signs of the differences rather than their magnitudes.

The results in Table 4.8 highlight some important observations. From the intra-route operators, the generalized Or-opt operator was identified as having the most impact, with a statistically significant increase in the TJT median being confirmed for the Hamrun instance. All other paired-samples tests for the intra-route operators did not yield significant findings at the 0.05 level. Although, in some cases, the removal of the exchange or generalized Or-opt operator decreases the TJT average slightly, the worsening of other TJT averages does not justify their removal. The two-opt operator was not found to impact the algorithm’s performance, but it was retained in the final version since its influence may manifest in longer runs. From the inter-route operators, the cross-exchange operator was found to have a great impact on the algorithm’s performance, especially for the Suffolk instance. No feasible solution with  $\underline{k} = 3$  routes was found for the Suffolk instance when this operator was removed. Hence, we do not present any results for this instance-operator pair as the number of routes is not comparable. A statistically significant increase in the TJT median upon the cross-exchange removal was found for ten instances, with relative increases of up to 20.96%. A significant result was also found for the Brisbane instance and the Or-exchange operator, while the remaining tests did not yield significant findings.

Further sensitivity analyses were performed to compare the algorithm’s performance using SD to that using RVND and to assess the algorithm’s performance in longer runs. The results are displayed in Table 4.9. Recall that our algorithm

Table 4.8: Sensitivity analysis assessing the impact of each neighbourhood operator on the algorithm’s performance. Inc. represents the relative increase in the TJT average across 25 runs, while bold  $p$ -values are significant at the 0.05 level.

Location	Exchange			Two-Opt			Generalized Or-Opt			Or-Exchange			Cross-Exchange			Create Multistop		
	Inc. (%)	$p$ -value	Inc. (%)	Inc. (%)	$p$ -value	Inc. (%)	Inc. (%)	$p$ -value	Inc. (%)	Inc. (%)	$p$ -value	Inc. (%)	Inc. (%)	$p$ -value	Inc. (%)	Inc. (%)	$p$ -value	
Mgarr	0.00	0.500	0.00	0.00	0.500	0.00	0.00	0.500	0.00	0.500	0.00	0.00	0.500	0.00	0.500	0.00	0.500	
Mellieha	0.00	0.500	0.00	0.00	0.500	0.00	0.00	0.500	0.00	0.500	0.00	0.00	0.500	0.00	0.500	0.00	0.500	
Porthcawl	0.00	0.500	0.00	0.00	0.500	0.00	0.00	0.500	0.00	0.500	0.00	0.00	0.500	0.00	0.500	0.00	0.500	
Qrendi	0.00	0.500	0.00	0.00	0.500	0.00	0.00	0.500	0.00	0.500	0.00	0.00	0.500	0.00	<b>0.001</b>	0.00	0.500	
Suffolk	-0.05	0.364	0.00	0.00	0.500	-0.13	0.402	0.402	0.26	0.095	N/A	N/A	N/A	0.03	0.500	0.03	0.500	
Senglea	0.00	0.500	0.00	0.00	0.500	0.03	0.500	0.500	0.21	0.227	0.50	0.084	0.00	0.500	0.00	0.500		
Victoria	0.10	0.313	0.00	0.00	0.500	0.27	0.105	0.105	0.13	0.148	0.22	<b>0.025</b>	-0.01	0.500	-0.01	0.500		
Pembroke	0.02	0.500	0.00	0.00	0.500	0.01	0.476	0.476	-0.07	0.460	3.09	<b>0.001</b>	0.00	0.500	0.00	0.500		
Canberra	-0.05	0.073	0.00	0.00	0.500	0.12	0.343	0.343	0.00	0.477	4.24	<b>0.001</b>	-0.06	0.228	-0.06	0.228		
Handaq	-0.02	0.436	0.00	0.00	0.500	0.03	0.427	0.427	0.14	0.247	1.22	<b>0.001</b>	0.00	0.500	0.00	0.500		
Valletta	0.01	0.500	0.00	0.00	0.500	-0.07	0.369	0.369	0.08	0.247	0.51	<b>0.002</b>	0.00	0.500	0.00	0.500		
Birkirkara	-0.25	0.134	0.00	0.00	0.500	0.68	0.092	0.092	-0.25	0.339	20.96	<b>0.001</b>	0.00	0.500	0.00	0.500		
Hamrun	-0.10	0.500	0.00	0.00	0.500	0.63	<b>0.032</b>	<b>0.032</b>	-0.10	0.420	0.86	<b>0.026</b>	0.00	0.500	0.00	0.500		
Cardiff	0.00	0.500	0.00	0.00	0.500	0.00	0.500	0.500	0.00	0.500	0.00	0.500	0.00	0.500	0.00	0.500		
Milton Keynes	0.00	0.500	0.00	0.00	0.500	0.00	0.500	0.500	0.06	0.375	-0.06	0.464	0.00	0.500	0.00	0.500		
Bridgend	0.00	0.500	0.00	0.00	0.315	-0.11	0.500	0.500	-0.12	0.153	0.08	0.217	-0.04	0.500	-0.04	0.500		
Edinburgh-2	0.00	0.500	0.00	0.00	0.500	0.00	0.500	0.500	0.00	0.500	0.00	0.500	0.00	0.500	0.00	0.500		
Edinburgh-1	0.00	0.500	0.00	0.00	0.500	-0.26	0.232	0.232	0.08	0.353	3.64	<b>0.001</b>	0.11	0.500	0.11	0.500		
Adelaide	0.00	0.500	0.00	0.00	0.500	-0.02	0.319	0.319	0.13	0.291	-0.10	0.500	0.00	0.500	0.00	0.500		
Brisbane	-0.02	0.459	0.00	0.00	0.500	0.57	0.144	0.144	1.03	<b>0.031</b>	16.72	<b>0.001</b>	0.43	0.115	0.43	0.115		

follows the SD direction, meaning that the best-improving move from all neighbourhoods is executed in each iteration. The algorithm was altered to follow an RVND direction. In each iteration, the best-improving move from a randomly selected neighbourhood out of the six considered is executed. From Column 2 of Table 4.9, we observe that the change from SD to RVND yields an improvement in the TJT average for eight instances, a deterioration for seven instances, and no change for five instances. Mann-Whitney test results (Columns 3 to 6 of Table 4.9) did not reveal statistically significant differences between the SD and RVND solutions. The mean ranks, however, point towards RVND being a marginally better alternative. The algorithm also converges faster with RVND, with relative decreases in the average computational time between 8.27% and 71.64%. For these reasons, RVND will be employed henceforth.

In further experiments, the number of iterations was doubled to analyze the impact of the termination criterion on the algorithm's performance. For fifteen instances, Column 7 of Table 4.9 shows a decrease in the TJT average between 0.06% and 0.71%. A one-tailed paired-samples non-parametric test (Wilcoxon

Table 4.9: Sensitivity analysis replacing SD with RVND, and doubling the number of iterations. Inc. (Dec.) represents the relative increase (decrease) in the TJT average across 25 runs. The values under the SD and RVND columns represent the mean ranks,  $U$  represents the Mann-Whitney test statistic, while bold  $p$ -values are significant at the 0.05 level.

Location	RVND					Double Iterations	
	Inc. (%)	SD	RVND	$U$	$p$ -value	Dec. (%)	$p$ -value
Mgarr	0.00	25.50	25.50	312.5	1.000	0.00	0.500
Mellieħa	0.00	25.50	25.50	312.5	1.000	0.00	0.500
Porthcawl	0.33	24.50	26.50	287.5	0.302	0.49	0.125
Qrendi	0.00	25.50	25.50	312.5	1.000	0.00	0.500
Suffolk	0.03	25.68	25.32	308.0	0.924	0.11	<b>0.008</b>
Senglea	-0.10	26.68	24.32	283.0	0.561	0.50	< <b>0.001</b>
Victoria	0.08	24.44	26.56	286.0	0.578	0.14	<b>0.004</b>
Pembroke	0.24	23.54	27.46	263.5	0.341	0.45	<b>0.004</b>
Canberra	-0.32	29.30	21.70	217.5	0.065	0.18	< <b>0.001</b>
ħandaq	-0.03	26.56	24.44	286.0	0.607	0.27	< <b>0.001</b>
Valetta	-0.09	25.56	25.44	311.0	0.977	0.16	<b>0.001</b>
Birkirkara	0.04	26.56	24.44	286.0	0.607	0.66	<b>0.001</b>
ħamrun	-0.06	26.20	24.80	295.0	0.732	0.33	<b>0.016</b>
Cardiff	0.00	25.50	25.50	312.5	1.000	0.00	0.500
Milton Keynes	0.22	23.74	27.26	268.5	0.317	0.06	0.125
Bridgend	-0.04	25.32	25.68	308.0	0.928	0.06	0.250
Edinburgh-2	0.00	25.50	25.50	312.5	1.000	0.00	0.500
Edinburgh-1	-0.08	25.74	25.26	306.5	0.907	0.39	< <b>0.001</b>
Adelaide	-0.09	26.62	24.38	284.5	0.586	0.14	< <b>0.001</b>
Brisbane	0.88	22.50	28.50	237.5	0.146	0.71	< <b>0.001</b>

signed-rank test or sign test, as applied before) revealed a statistically significant decrease in the TJT median for twelve instances. The average computational time increased between 108.98% and 165.06%. Given that the best solution improved for six instances upon doubling the number of iterations, it is arguable that the improvements gained counterbalance the increase in computational time. Thus, moving forward, the number of iterations for each instance will be kept as double the number presented in Table 4.5. Further details on the RVND and the longer runs are presented in Table 4.10 and are available at (Sciortino, 2024b).

The best SRH-4 solutions are displayed in Table 4.11 and visualized in Appendix B.4. These solutions have the lowest total journey time and, in the case of ties, average walking time. Some of these solutions were achieved in multiple runs (Mgarr: 3, Mellieña: 12, Porthcawl: 4, Qrendi: 5, Suffolk: 2, Senglea: 3, Pembroke: 3, Cardiff: 2, Edinburgh-2: 25 runs). The Mgarr, Qrendi, and Cardiff instances also have additional solutions with the same total journey time and average walking time, but with different subsets of visited bus stops and/or route configurations. This is a merit of our algorithm, as it can provide multiple solutions with equivalent quality, from which the most appropriate one can be identified based on local factors such as bus depot location and stop accessibility.

From Table 4.11, it can be inferred that approximately 80% of the routes have a journey time that is less than half the maximum journey time  $m_t = 45$  minutes and that the percentage of capacity utilized ranges between 69.44% and 100%. We now compare the routes of the Mgarr, Mellieña, Porthcawl, Qrendi, and Senglea instances with the corresponding ones in Table 3.4 (MIP solutions). We consider only these instances as the number of routes is not comparable for the other instances. The SRH-4 routes for the Mgarr instance are almost the same as the MIP routes, with Stop 31 used instead of Stop 12, and Stop 15 featuring in one route instead of two. This is the only instance for which the average journey time is slightly worse for SRH-4 (+0.03 minutes), while the average walking time is slightly better (−0.01 minutes). The Mellieña SRH-4 solution has the same average journey time and slightly better average walking time (−0.01 minutes) than the corresponding MIP solution. The only difference is that the former uses Stop 45, while the latter uses Stop 49. The Porthcawl solutions are equivalent. Moreover, for the Qrendi instance, the solutions are considerably different, with the SRH-4 solution having better average journey time (−0.88 minutes) at the expense of a higher average walking time (+0.29 minutes). Similar observations can be made for the Senglea instance, with a decrease of 0.99 minutes in average journey time and an increase of 0.53 minutes in average walking time.

Table 4.10: Results for SRH-4 with RVND, and SRH-4 with RVND and double iterations. Results are averaged across 25 runs and presented as mean  $\pm$  standard deviation. Time gives the computational time (seconds), FR stands for Feasibility Rate (percentage of iterations yielding a solution with all route journey times being at most 45 minutes), and TJT stands for Total Journey Time (minutes).

Location	SRH-4: RVND				SRH-4: RVND + Double Iterations				
	Time (s)	FR (%)	TJT (mins.)	Time (s)	FR (%)	TJT (mins.)	Time (s)	FR (%)	TJT (mins.)
Mgarr	60.37 $\pm$ 0.83	100.00 $\pm$ 0.00	54.10 $\pm$ 0.00	137.19 $\pm$ 4.80	100.00 $\pm$ 0.00	54.10 $\pm$ 0.00	4.80	100.00 $\pm$ 0.00	54.10 $\pm$ 0.00
Mellieha	52.77 $\pm$ 0.56	100.00 $\pm$ 0.00	56.30 $\pm$ 0.00	128.43 $\pm$ 6.87	100.00 $\pm$ 0.00	56.30 $\pm$ 0.00	6.87	100.00 $\pm$ 0.00	56.30 $\pm$ 0.00
Porthcawl	253.55 $\pm$ 11.39	99.99 $\pm$ 0.01	27.97 $\pm$ 0.23	586.01 $\pm$ 12.47	99.98 $\pm$ 0.03	27.83 $\pm$ 0.43	12.47	99.98 $\pm$ 0.03	27.83 $\pm$ 0.43
Qrendi	69.96 $\pm$ 0.98	100.00 $\pm$ 0.00	75.73 $\pm$ 0.00	180.07 $\pm$ 2.01	100.00 $\pm$ 0.00	75.73 $\pm$ 0.00	2.01	100.00 $\pm$ 0.00	75.73 $\pm$ 0.00
Suffolk	49.92 $\pm$ 6.59	16.01 $\pm$ 4.16	116.739 $\pm$ 1.09	117.08 $\pm$ 9.06	17.50 $\pm$ 3.29	116.61 $\pm$ 0.90	9.06	17.50 $\pm$ 3.29	116.61 $\pm$ 0.90
Senglea	90.21 $\pm$ 3.02	100.00 $\pm$ 0.00	72.37 $\pm$ 0.45	225.90 $\pm$ 6.52	100.00 $\pm$ 0.00	72.01 $\pm$ 0.54	6.52	100.00 $\pm$ 0.00	72.01 $\pm$ 0.54
Victoria	47.54 $\pm$ 1.32	100.00 $\pm$ 0.00	93.10 $\pm$ 0.47	126.00 $\pm$ 3.29	100.00 $\pm$ 0.00	92.96 $\pm$ 0.41	3.29	100.00 $\pm$ 0.00	92.96 $\pm$ 0.41
Pembroke	101.81 $\pm$ 5.48	100.00 $\pm$ 0.00	103.96 $\pm$ 0.96	232.69 $\pm$ 13.97	100.00 $\pm$ 0.00	103.49 $\pm$ 0.68	13.97	100.00 $\pm$ 0.00	103.49 $\pm$ 0.68
Canberra	43.76 $\pm$ 6.64	100.00 $\pm$ 0.00	179.44 $\pm$ 0.68	91.45 $\pm$ 9.98	100.00 $\pm$ 0.00	179.11 $\pm$ 0.53	9.98	100.00 $\pm$ 0.00	179.11 $\pm$ 0.53
Handaq	90.45 $\pm$ 4.59	100.00 $\pm$ 0.00	96.87 $\pm$ 0.56	215.78 $\pm$ 8.99	100.00 $\pm$ 0.00	96.61 $\pm$ 0.57	8.99	100.00 $\pm$ 0.00	96.61 $\pm$ 0.57
Valletta	133.84 $\pm$ 7.93	100.00 $\pm$ 0.00	104.67 $\pm$ 0.74	335.78 $\pm$ 18.52	100.00 $\pm$ 0.00	104.50 $\pm$ 0.78	18.52	100.00 $\pm$ 0.00	104.50 $\pm$ 0.78
Birkirkara	147.66 $\pm$ 5.25	100.00 $\pm$ 0.00	98.13 $\pm$ 2.42	354.98 $\pm$ 14.25	100.00 $\pm$ 0.00	97.49 $\pm$ 2.64	14.25	100.00 $\pm$ 0.00	97.49 $\pm$ 2.64
Hamrun	117.14 $\pm$ 3.94	100.00 $\pm$ 0.00	99.63 $\pm$ 1.38	277.27 $\pm$ 8.16	100.00 $\pm$ 0.00	99.31 $\pm$ 1.40	8.16	100.00 $\pm$ 0.00	99.31 $\pm$ 1.40
Cardiff	362.30 $\pm$ 19.94	59.70 $\pm$ 0.38	67.12 $\pm$ 0.00	844.80 $\pm$ 41.60	59.73 $\pm$ 0.33	67.12 $\pm$ 0.00	41.60	59.73 $\pm$ 0.33	67.12 $\pm$ 0.00
Milton Keynes	250.62 $\pm$ 6.82	100.00 $\pm$ 0.00	59.91 $\pm$ 0.59	602.61 $\pm$ 15.78	100.00 $\pm$ 0.00	59.88 $\pm$ 0.57	15.78	100.00 $\pm$ 0.00	59.88 $\pm$ 0.57
Bridgend	200.84 $\pm$ 1.44	99.37 $\pm$ 0.10	170.65 $\pm$ 0.65	475.82 $\pm$ 4.93	99.38 $\pm$ 0.10	170.55 $\pm$ 0.53	4.93	99.38 $\pm$ 0.10	170.55 $\pm$ 0.53
Edinburgh-2	306.67 $\pm$ 2.99	100.00 $\pm$ 0.00	59.13 $\pm$ 0.00	741.87 $\pm$ 16.36	100.00 $\pm$ 0.00	59.13 $\pm$ 0.00	16.36	100.00 $\pm$ 0.00	59.13 $\pm$ 0.00
Edinburgh-1	249.59 $\pm$ 8.37	100.00 $\pm$ 0.00	144.27 $\pm$ 1.63	593.07 $\pm$ 34.05	100.00 $\pm$ 0.00	143.71 $\pm$ 1.56	34.05	100.00 $\pm$ 0.00	143.71 $\pm$ 1.56
Adelaide	256.59 $\pm$ 8.08	100.00 $\pm$ 0.00	128.36 $\pm$ 1.43	541.72 $\pm$ 26.25	100.00 $\pm$ 0.00	128.18 $\pm$ 1.45	26.25	100.00 $\pm$ 0.00	128.18 $\pm$ 1.45
Brisbane	241.02 $\pm$ 5.53	100.00 $\pm$ 0.00	214.95 $\pm$ 4.17	518.79 $\pm$ 26.33	100.00 $\pm$ 0.00	213.41 $\pm$ 4.31	26.33	100.00 $\pm$ 0.00	213.41 $\pm$ 4.31

Table 4.11: Best SRH-4 (with RVND and double iterations) solutions. Cap. is short for Capacity, JT stands for Journey Time (minutes), and TJT stands for Total Journey Time (minutes). The routes are presented in terms of the indices of the vertices in  $V_1$ , and the subscripts represent the number of boarding students.

Location	Routes	Load/Cap.	JT	TJT
Mgarr	21 <sub>32</sub> → 19 <sub>11</sub> → 0	43/44	9.83	54.10
	48 <sub>1</sub> → 33 <sub>9</sub> → 11 <sub>10</sub> → 31 <sub>29</sub> → 15 <sub>2</sub> → 0	51/53	14.53	
	42 <sub>10</sub> → 39 <sub>21</sub> → 35 <sub>12</sub> → 0	43/44	14.33	
	47 <sub>18</sub> → 52 <sub>20</sub> → 57 <sub>15</sub> → 0	53/53	15.40	
Mellieħa	43 <sub>14</sub> → 45 <sub>3</sub> → 46 <sub>14</sub> → 59 <sub>19</sub> → 0	50/53	15.82	56.30
	33 <sub>7</sub> → 7 <sub>5</sub> → 9 <sub>28</sub> → 0	40/44	9.98	
	66 <sub>5</sub> → 83 <sub>9</sub> → 80 <sub>14</sub> → 5 <sub>21</sub> → 0	49/53	15.80	
	30 <sub>8</sub> → 20 <sub>16</sub> → 27 <sub>8</sub> → 0	32/36	14.70	
Porthcawl	16 <sub>5</sub> → 27 <sub>19</sub> → 151 <sub>17</sub> → 38 <sub>25</sub> → 0	66/70	26.87	26.87
Qrendi	29 <sub>43</sub> → 76 <sub>8</sub> → 0	51/53	13.58	75.73
	89 <sub>27</sub> → 76 <sub>26</sub> → 0	53/53	13.40	
	50 <sub>15</sub> → 99 <sub>19</sub> → 146 <sub>1</sub> → 54 <sub>10</sub> → 0	45/53	11.95	
	67 <sub>3</sub> → 68 <sub>14</sub> → 69 <sub>24</sub> → 137 <sub>12</sub> → 0	53/53	16.72	
	124 <sub>15</sub> → 126 <sub>25</sub> → 155 <sub>4</sub> → 152 <sub>9</sub> → 0	53/53	20.08	
Suffolk	4 <sub>34</sub> → 116 <sub>18</sub> → 172 <sub>2</sub> → 5 <sub>22</sub> → 98 <sub>4</sub> → 0	80/80	31.08	114.73
	121 <sub>8</sub> → 65 <sub>12</sub> → 126 <sub>9</sub> → 155 <sub>9</sub> → (cont.)			
	111 <sub>17</sub> → 140 <sub>8</sub> → 43 <sub>2</sub> → 0	65/65	44.13	
	171 <sub>29</sub> → 130 <sub>6</sub> → 173 <sub>1</sub> → 106 <sub>4</sub> → (cont.)			
	68 <sub>2</sub> → 50 <sub>5</sub> → 73 <sub>5</sub> → 122 <sub>12</sub> → 0	64/65	39.52	
Senglea	44 <sub>49</sub> → 0	49/53	10.63	71.32
	123 <sub>21</sub> → 81 <sub>6</sub> → 151 <sub>16</sub> → 0	43/44	18.10	
	6 <sub>52</sub> → 0	52/53	13.62	
	182 <sub>8</sub> → 100 <sub>20</sub> → 120 <sub>20</sub> → 53 <sub>4</sub> → 0	52/53	15.48	
	55 <sub>17</sub> → 0	17/18	3.48	
	49 <sub>43</sub> → 53 <sub>10</sub> → 0	53/53	10.00	
Victoria	204 <sub>9</sub> → 182 <sub>11</sub> → 111 <sub>5</sub> → 0	25/36	15.17	92.60
	199 <sub>2</sub> → 156 <sub>12</sub> → 271 <sub>5</sub> → 134 <sub>9</sub> → (cont.)			
	137 <sub>9</sub> → 12 <sub>6</sub> → 23 <sub>1</sub> → 0	44/44	22.95	
	79 <sub>7</sub> → 41 <sub>7</sub> → 51 <sub>8</sub> → 29 <sub>12</sub> → (cont.)			
	59 <sub>2</sub> → 73 <sub>5</sub> → 232 <sub>8</sub> → 0	49/53	27.93	
	143 <sub>13</sub> → 141 <sub>8</sub> → 152 <sub>11</sub> → 121 <sub>6</sub> → (cont.)			
	175 <sub>8</sub> → 172 <sub>6</sub> → 168 <sub>1</sub> → 0	53/53	26.55	
Pembroke	22 <sub>47</sub> → 302 <sub>5</sub> → 0	50/53	16.20	102.58
	88 <sub>21</sub> → 97 <sub>11</sub> → 84 <sub>21</sub> → 0	52/53	10.42	
	287 <sub>17</sub> → 15 <sub>22</sub> → 0	53/53	19.25	
	59 <sub>31</sub> → 255 <sub>13</sub> → 206 <sub>6</sub> → 207 <sub>3</sub> → 0	46/53	13.13	
	62 <sub>21</sub> → 243 <sub>16</sub> → 12 <sub>13</sub> → 0	53/53	19.92	
	67 <sub>26</sub> → 68 <sub>20</sub> → 0	42/44	13.10	
	44 <sub>33</sub> → 17 <sub>9</sub> → 0	39/44	10.57	

CHAPTER 4. A HEURISTIC ALGORITHM FOR THE SINGLE-SCHOOL BUS ROUTING PROBLEM

	122 <sub>63</sub> → 142 <sub>17</sub> → 0	80/80	19.53	
	327 <sub>6</sub> → 324 <sub>8</sub> → 328 <sub>27</sub> → 329 <sub>17</sub> → 330 <sub>1</sub> → 0	59/59	38.93	
	214 <sub>50</sub> → 190 <sub>25</sub> → 0	75/78	18.12	
Canberra	155 <sub>7</sub> → 164 <sub>13</sub> → 256 <sub>12</sub> → 206 <sub>28</sub> → 0	60/61	23.93	178.13
	268 <sub>73</sub> → 0	73/78	17.92	
	35 <sub>9</sub> → 27 <sub>18</sub> → 24 <sub>7</sub> → 11 <sub>7</sub> → 142 <sub>31</sub> → 0	72/78	27.07	
	85 <sub>4</sub> → 144 <sub>18</sub> → 246 <sub>9</sub> → 227 <sub>27</sub> → 42 <sub>22</sub> → 0	80/80	32.63	
	87 <sub>34</sub> → 206 <sub>13</sub> → 375 <sub>3</sub> → 0	50/53	17.08	
	66 <sub>11</sub> → 58 <sub>12</sub> → 40 <sub>10</sub> → 41 <sub>7</sub> → 0	40/44	13.73	
	2 <sub>11</sub> → 7 <sub>9</sub> → 338 <sub>14</sub> → 332 <sub>19</sub> → 0	53/53	21.03	
Handaq	223 <sub>13</sub> → 290 <sub>6</sub> → 357 <sub>6</sub> → 113 <sub>18</sub> → 77 <sub>10</sub> → 0	53/53	18.63	95.62
	275 <sub>16</sub> → 35 <sub>25</sub> → 0	41/44	13.95	
	159 <sub>34</sub> → 53 <sub>14</sub> → 0	48/53	11.18	
	268 <sub>9</sub> → 383 <sub>19</sub> → 195 <sub>23</sub> → 0	51/53	18.85	
	141 <sub>19</sub> → 227 <sub>7</sub> → 334 <sub>6</sub> → 86 <sub>20</sub> → 0	52/53	21.63	
	283 <sub>14</sub> → 0	14/14	5.88	
Valetta	43 <sub>14</sub> → 31 <sub>18</sub> → 401 <sub>20</sub> → 0	52/53	16.17	102.38
	427 <sub>19</sub> → 114 <sub>27</sub> → 0	46/53	19.58	
	46 <sub>13</sub> → 207 <sub>13</sub> → 58 <sub>13</sub> → 61 <sub>14</sub> → 0	53/53	20.27	
	152 <sub>31</sub> → 156 <sub>14</sub> → 427 <sub>5</sub> → 0	50/53	12.58	
	397 <sub>14</sub> → 293 <sub>23</sub> → 281 <sub>13</sub> → 0	50/53	17.02	
	393 <sub>16</sub> → 350 <sub>24</sub> → 360 <sub>11</sub> → 0	51/53	14.27	
	371 <sub>17</sub> → 370 <sub>28</sub> → 447 <sub>7</sub> → 0	52/53	11.85	
Birkirkara	104 <sub>11</sub> → 456 <sub>10</sub> → 384 <sub>21</sub> → 410 <sub>11</sub> → 0	53/53	19.48	93.92
	466 <sub>3</sub> → 251 <sub>5</sub> → 316 <sub>12</sub> → 2 <sub>15</sub> → 445 <sub>15</sub> → 0	50/53	18.72	
	497 <sub>20</sub> → 164 <sub>13</sub> → 167 <sub>12</sub> → 0	45/53	16.50	
	243 <sub>17</sub> → 342 <sub>3</sub> → 482 <sub>22</sub> → 0	42/44	14.05	
	436 <sub>23</sub> → 437 <sub>8</sub> → 44 <sub>8</sub> → 0	39/44	12.20	
Hamrun	261 <sub>16</sub> → 386 <sub>13</sub> → 365 <sub>17</sub> → 0	46/53	14.90	97.40
	3 <sub>27</sub> → 382 <sub>23</sub> → 0	50/53	11.17	
	123 <sub>23</sub> → 463 <sub>19</sub> → 11 <sub>10</sub> → 0	52/53	12.23	
	509 <sub>6</sub> → 95 <sub>26</sub> → 74 <sub>2</sub> → 75 <sub>13</sub> → 0	47/53	16.35	
	456 <sub>33</sub> → 508 <sub>45</sub> → 109 <sub>2</sub> → 0	80/80	26.73	
Cardiff	353 <sub>50</sub> → 117 <sub>12</sub> → 456 <sub>14</sub> → 0	76/78	40.38	67.12
	165 <sub>43</sub> → 219 <sub>7</sub> → 0	50/51	11.38	
	293 <sub>43</sub> → 373 <sub>37</sub> → 0	80/80	15.20	
	182 <sub>38</sub> → 143 <sub>29</sub> → 0	67/70	14.95	
Milton Keynes	199 <sub>25</sub> → 354 <sub>52</sub> → 0	77/78	17.50	59.03
	507 <sub>18</sub> → 263 <sub>4</sub> → 74 <sub>8</sub> → 130 <sub>19</sub> → 0	49/49	28.87	
	57 <sub>80</sub> → 0	80/80	23.73	
	567 <sub>72</sub> → 92 <sub>7</sub> → 0	79/80	32.50	
Bridgend	57 <sub>80</sub> → 0	80/80	23.73	169.95
	385 <sub>29</sub> → 178 <sub>23</sub> → 246 <sub>2</sub> → 432 <sub>21</sub> → 0	75/78	36.60	
	363 <sub>8</sub> → 393 <sub>4</sub> → 268 <sub>6</sub> → 0	18/23	24.52	

Edinburgh-2	$142_{80} \rightarrow 0$	80/80	13.40	59.13
	$142_{80} \rightarrow 0$	80/80	13.40	
	$509_{42} \rightarrow 575_{38} \rightarrow 0$	80/80	21.98	
	$575_{80} \rightarrow 0$	80/80	10.35	
Edinburgh-1	$354_{60} \rightarrow 747_{20} \rightarrow 0$	80/80	18.85	140.75
	$129_{12} \rightarrow 6_{56} \rightarrow 0$	68/70	12.72	
	$92_{70} \rightarrow 0$	70/70	14.42	
	$104_{50} \rightarrow 354_3 \rightarrow 747_{26} \rightarrow 0$	79/80	21.83	
	$270_{69} \rightarrow 0$	69/70	11.22	
	$92_{80} \rightarrow 0$	80/80	15.25	
	$337_{55} \rightarrow 274_{20} \rightarrow 0$	75/78	17.22	
	$99_{74} \rightarrow 635_5 \rightarrow 0$	79/80	14.55	
	$143_{25} \rightarrow 293_{38} \rightarrow 635_{17} \rightarrow 0$	80/80	14.70	
Adelaide	$1005_{72} \rightarrow 0$	72/78	10.03	126.40
	$291_{13} \rightarrow 725_{37} \rightarrow 292_{21} \rightarrow 0$	71/78	18.33	
	$146_{62} \rightarrow 1072_{17} \rightarrow 0$	79/80	16.57	
	$1094_{29} \rightarrow 894_{19} \rightarrow 508_{30} \rightarrow 0$	78/78	18.17	
	$452_{80} \rightarrow 0$	80/80	16.32	
	$604_{27} \rightarrow 786_{43} \rightarrow 106_{10} \rightarrow 0$	80/80	22.45	
	$1072_{70} \rightarrow 0$	70/70	8.05	
	$452_{14} \rightarrow 1023_{21} \rightarrow 0$	35/35	16.48	
Brisbane	$726_{77} \rightarrow 0$	77/78	19.28	207.00
	$1098_{52} \rightarrow 1507_{27} \rightarrow 0$	79/80	22.17	
	$1024_{45} \rightarrow 1044_{35} \rightarrow 0$	80/80	18.28	
	$1529_{28} \rightarrow 252_{38} \rightarrow 0$	66/70	15.67	
	$434_{80} \rightarrow 0$	80/80	19.57	
	$647_{30} \rightarrow 644_8 \rightarrow 1243_{23} \rightarrow 398_3 \rightarrow 0$	64/65	23.67	
	$1406_{24} \rightarrow 1407_{35} \rightarrow 603_{20} \rightarrow 0$	79/80	25.62	
	$905_{38} \rightarrow 1744_{42} \rightarrow 0$	80/80	18.73	
$812_{53} \rightarrow 367_4 \rightarrow 288_{21} \rightarrow 0$	78/78	19.32		
$1373_{74} \rightarrow 0$	74/78	24.70		

## 4.7 Summary

In this chapter, we developed a local search approach to the S-SBRP which will form the basis of the next chapters. At the outset, we presented a greedy heuristic, based on the Kolesar-Walker heuristic for the unicast set covering problem, to select the bus stops to be visited. We then assigned students to their closest designated bus stop to promote short walking distances and proposed three route constructive heuristics: PNNH, which builds routes backward and in parallel using the NN heuristic; SRH, which constructs routes sequentially and randomly; and RIH, which assigns stops to routes in non-increasing order of their regret values and is an extension of Pacheco and Marti (2006)'s heuristic catering for

multistops.

We also considered six neighbourhood operators, namely three intra-route operators applying to a single route (exchange, two-opt, generalized Or-opt) and three inter-route operators applying to a pair of routes (Or-exchange, cross-exchange, create multistop). We saw how these operators can be used to improve initial routes generated by the aforementioned constructive heuristics. We embedded the constructive and improvement heuristics in an ILS algorithm that also contains a destroy-and-repair operator as a perturbation mechanism. We proposed four different ways of applying this operator that vary in the way they alter the set  $V'_1$  of visited bus stops: PM 1 uses random restart; PM 2 alters the set  $V'_1$  used in the previous iteration; PM 3 alters the set  $V'_1$  that yielded the most recent best feasible solution; and PM 4 uses either PM 2 or PM 3, both with probability 0.5.

Results obtained from thorough computational experiments using a variety of problem instances were then provided. Initially, we compared twelve algorithm variants and concluded, from the average costs across 25 runs, that the combination of SRH and PM 4 is the best-performing variant. We then conducted complementary analyses to assess the impact of (i) each neighbourhood operator, (ii) the SD versus RVND directions, and (iii) the number of iterations, on the algorithm's performance. We noted that the generalized Or-opt and the cross-exchange operators are the most influential intra-route and inter-route operators, respectively. We also found that RVND is a faster and marginally better alternative to SD and that, time permitting, longer runs may converge to statistically significantly better outcomes.

Our algorithm has several notable advantages. In a short amount of computational time, the algorithm produces feasible solutions characterized by (i) the minimal number of buses for most problem instances, (ii) moderately short routes (with 80% of the generated routes under 22.5 minutes), and (iii) high capacity utilization rates (ranging from 69.44% to 100%). Moreover, it is capable of producing a series of solutions, not simply one solution. Different solutions with equivalent or near-equivalent quality can therefore be compared to identify the most appropriate one from the operators' perspectives. The obvious drawback of our algorithm, as with any other (meta-)heuristic technique, is that it does not guarantee the production of optimal solutions, nor does it indicate optimality when a global optimum happens to have been found. As it stands, it also treats travel times as deterministic parameters. This assumption is unrealistic in real-world transport networks due to factors such as weather conditions, traffic congestion, road closures, and traffic accidents. It is this downside that motivates the exploration of stochastic travel times in the next chapter.

# Chapter 5

## Modelling Stochastic Travel Times

### 5.1 Introduction

It is unrealistic to assume that everything follows a predetermined schedule in a real-life environment. Random disturbances in real-world operations can significantly affect the routing and scheduling of vehicles and, in turn, the service reliability and operational efficiency of transportation systems. In classical formulations of the VRP, the parameters representing the presence of customers, customer demands, service times, and travel times are assumed to be deterministic (known precisely in advance). However, in reality, any of these parameters are prone to fluctuations.

Ignoring the inherent randomness and unpredictability of these parameters may lead to solutions that are far from optimal or even infeasible when implemented in actual operations. For instance, an inaccurate travel time prediction can result in a late arrival, leading to customer dissatisfaction with the service, employee overtime, reduced quality of life, and cost overruns. In such a case, we say that the route ‘fails’ when exercised in a real situation. Given this, it is worthwhile to investigate stochastic VRPs (SVRPs) which are VRPs with at least one of the aforementioned parameters treated as stochastic. In this research, we only address disturbances arising from stochastic travel times. The reason why we do not consider other stochastic parameters, such as the demands, is that school bus drivers are not usually prenotified about absent students and are expected to perform their routes as planned. Therefore, we will henceforth focus on variations in travel times.

Travel time is a key factor in assessing the performance and quality of service offered by a transportation system. In the context of school bus routing, the

travel time from the bus garage to the school on a Monday morning is likely to be different from that on a Tuesday morning or the following Monday morning. This is because travel times tend to fluctuate owing to uncertain factors such as inclement weather, traffic congestion, traffic accidents, vehicle breakdowns, road maintenance works, construction zones, traffic signal controls, pedestrian crossings, and special events. For example, a traffic accident on a road leading to a bus stop can cause the travel time on a path between two bus stops to become unusually higher than the expected travel time. In this case, the actual arrival time at the school will not accord with the projected arrival time. Thus, the feasibility of a solution cannot be reasonably ensured unless the optimization methodology accounts for the uncertainty in travel times. Travel times are not only likely to vary from day to day but also across different hours of the same day. In this research, we do not tackle this time-dependent aspect of the VRP since the bus routes are always performed in the same morning hour before school starts.

To reflect real-world traffic conditions and reduce the negative effect of stochastic disturbances, in this chapter we establish a framework for optimal routing in traffic networks with travel time uncertainties. We therefore model the travel time along any arc in the road network as a random variable whose probability distribution is assumed to be known. We also assume that the routes are planned before the realizations of the travel times and that the drivers must follow the a priori routing plans (no reoptimization is performed en route). The stochastic nature of our proposed framework allows us to design routes that respond to travel time volatility and remain feasible even for travel time realizations that differ substantially from the expectations. It also helps us better understand the negative consequences of poor routing decisions made without accounting for travel time uncertainties. For this particular chapter, for tractability, we also assume that the travel time probability distributions are independent. This means that a travel time variation on a particular arc does not cause variations on other arcs. In the next chapter, we drop this assumption by considering correlations between travel times.

This chapter is structured as follows. Section 5.2 presents two primary approaches for modelling a VRP with stochastic travel times. Given that we assume known probability distributions for the random travel times, we discuss different distributions previously used to fit or model travel time data in Section 5.3. In Section 5.4, we introduce travel time reliability metrics that can measure the consistency of routes performed several times (for example, daily). Combining ideas discussed in the three previous sections, we propose a framework for addressing our SBRP (or, equivalently, a time-constrained SDVRP) under stochastic and

independent travel times in Section 5.5. This framework is applied to the twenty instances summarized in Tables 3.1 and 3.2 in Section 5.6. The results are then compared with those of the deterministic framework presented in Chapter 4.

## 5.2 Modelling a VRP with Stochastic Travel Times

The definition of an SVRP solution differs from that of its deterministic counterpart and the solution methodologies are considerably more complex (Gendreau et al., 1996). This section outlines two main paradigms for modelling SVRPs, namely *stochastic programming* and *robust optimization*.

Stochastic programming (SP) requires a complete knowledge of the probability distributions of the random parameters (in our case, travel times). Estimating these distributions and their related parameters requires a large amount of data and can present statistical and computational difficulties. A stochastic program is typically modelled as a *stochastic program with recourse* (SPR) or as a *chance-constrained program* (CCP). The former approach involves two stages. In the first stage, an a priori solution is determined, followed by the realizations of the random parameters. In the second stage, predetermined corrective actions, called *recourse*, are applied if failure occurs when exercising the first-stage solution. The recourse usually has an associated cost that must be considered when designing the first-stage solution. The objective of an SPR is to minimize the cost of the first-stage solution plus the expected cost of the recourse. For example, in a VRP with stochastic pick-up demands, the accumulated load along a preplanned route may exceed the vehicle's capacity when the actual demands are revealed. In this case, the driver must return to the depot to unload and resume the service at the last visited customer node. This is known as a *detour-to-depot*. In a VRPTW, servicing a customer outside their time window may be allowed at a penalty for either earliness or lateness. Some VRP works featuring stochastic travel times and an SPR model include those of Russell and Urban (2008), Li et al. (2010a), Taş et al. (2013), Zhang et al. (2013), Yan et al. (2014), Chu et al. (2017), Oyola (2019), Rajabi-Bahaabadi et al. (2021), Messaoud (2022), and Iklassov et al. (2024). An SPR model is not applicable in a school bus routing context like ours since the bus drivers must follow the same routes every morning and the route duration limits must be treated as hard constraints.

CCPs do not incorporate recourse or its expected cost, but instead constrain the probability of route failure (for example, load exceeding capacity or arrival time exceeding due time) to remain below a specified threshold. Regard-

ing stochastic travel times, travellers cannot precisely estimate the arrival time at their destination. Thus, they may be concerned with the likelihood that their trip is completed successfully within a given deadline (*on-time arrival*). Laporte et al. (1992) were the first to include stochastic travel times in a VRP model explicitly. They considered route duration constraints and presented a CCP model, minimizing the routing costs, and an SPR model, penalizing the expected tardiness costs. They modelled travel times using up to five discrete scenarios and applied a B&C algorithm to solve instances with up to twenty vehicles. In this study, we model stochastic travel times via a CCP model in Section 5.5. Here, we summarize some key articles that employ a similar approach.

Li et al. (2010a) studied a VRPTW with stochastic travel and service times. They considered a chance constraint for each customer's time window and the driver's duration and minimized the number of vehicles used and, secondarily, the travel distance. Their TS-based heuristic algorithm combined with Monte Carlo simulation for estimating the probabilities produced good quality solutions but was computationally expensive. Ehmke et al. (2015) also addressed the VRPTW with stochastic travel times via a TS algorithm. They included a chance constraint for each customer's service deadline and minimized the number of vehicles and, secondarily, the total duration of the routes. They approximated the arrival time and start-service time distributions using extreme value theory. Bomboi et al. (2022) extended the work of Ehmke et al. (2015) and presented a model with single chance constraints (risk of failure to adhere to the time window is checked for each customer independently) and another model with joint chance constraints (risk of failure is aggregated for all customers in an entire route). They devised and compared several algorithms based on closed formulas or sampling to assess the feasibility of routes for correlated and time-dependent travel times.

For a VRPTW with stochastic travel and service times, Zhang et al. (2013) proposed a CCP model that guarantees a minimum on-time arrival probability at each customer location and the depot. They minimized the number of vehicles and, secondarily, the total expected travel time (operating cost) and the weighted expected earliness, tardiness, and excess route duration (penalty cost). They developed an iterated TS algorithm that includes a route reduction mechanism to decrease the number of required vehicles. Moreover, they adapted Miller-Hooks and Mahmassani (1998)'s  $\alpha$ -discrete approximation method to estimate the arrival time distributions in the presence of time windows. Miranda and Conceição (2016) applied an ILS algorithm for the VRPTW with stochastic travel and service times and a minimum probability for arriving at each customer before the end of their time window (service level). They aimed to minimize the vehicles'

fixed costs and the total expected travel time. In addition, they proposed a statistical method to estimate the cumulative probability function of the arrival time and, consequently, the service level of each customer.

Hou and Zhou (2010) presented a CCP model for a VRP with simultaneous pickup and delivery (VRPSPD). Considering stochastic demands and travel times, they allowed small chances of route overload and overtime and minimized the total distance via a GA. Zhang et al. (2012) formulated a VRPSPD as a CCP model that minimizes the expected cost of all routes and the cost of vehicle departures. They transformed their chance constraints on the routes' durations into deterministic constraints and proposed a scatter search algorithm to optimize the model. Chen et al. (2014a) used similar transformations of the chance constraints for a variant of the capacitated arc routing problem with stochastic travel and service times. Their objective was to minimize the fixed cost of the vehicle fleet and the total dead-heading travel cost (i.e., the cost of travelling over arcs without servicing them). They initially employed a B&C algorithm, followed by an adaptive large neighbourhood search (ALNS) algorithm.

More recently, Chen et al. (2018a) aimed to find reliable paths by minimizing the travel time budget (threshold) under a predetermined on-time arrival probability. They proposed a moment-matching-based GA combined with local search for intensification. The moment-matching method, discussed further in Section 5.5.1.2, was employed to approximate the paths' travel time distributions. Specifically for the SBRP, Babaei and Rajabi-Bahaabadi (2019) claimed that their work was the first to incorporate on-time arrival chance constraints. They modelled travel times by discrete random variables that vary across different time-of-day intervals and enforced, for each bus, a minimum on-time arrival probability at the school. Using ACO, they sought to minimize an objective function modelling the costs imposed on students for using public instead of private transportation and the operating costs of buses. They stated that a possible extension of their work would be to consider split deliveries. Other examples of works with stochastic travel times and on-time arrival chance constraints are those of Tavakkoli-Moghaddam et al. (2012) (CVRP), Sarmiento et al. (2015) (VRP), Binart et al. (2016) (multidepot VRPTW), Gómez et al. (2016) (CVRP), Gutierrez et al. (2016) (VRPTW), and Vareias et al. (2019) (VRP with self-imposed TWs).

In contrast to SP, robust optimization (RO) does not characterize a stochastic parameter by its probability distribution. Instead, it assumes that the stochastic parameter belongs to a set of plausible values known as an *uncertainty set*. This methodology finds an optimal solution with respect to a robustness criterion such as *worst-case criterion*, whereby an optimal solution is found for the worst realization, leading to a min-max objective form (for a minimization problem). RO

is practical when historical data on uncertain parameters is insufficient to construct accurate probability distributions or when challenging statistical and/or computational issues arise when modelling the distributions. In such cases, an uncertainty set for each stochastic parameter can be estimated from the data. For instance, one can estimate that travelling between locations A and B takes between 20 (best-case) and 40 (worst-case) minutes. The uncertainty sets corresponding to different parameters can be independent or correlated. According to Ordóñez (2010), work in various applications has shown that an RO model is only modestly more complex than the original deterministic model. However, RO solutions tend to be overly conservative because they prioritize feasibility under the worst-case scenario within the uncertainty set, often at the expense of optimality in more probable scenarios. For more details on the application of RO for VRPs with stochastic travel times, the reader is referred to works such as those of Lee et al. (2012), Agra et al. (2013), Solano-Charris et al. (2015), Adulyasak and Jaillet (2016), Jaillet et al. (2016), Braaten et al. (2017), Samadi-Dana et al. (2017), Hu et al. (2018), Sun et al. (2018), Eufinger et al. (2020), and Yin et al. (2023).

It is also possible to combine the SP and RO frameworks in one model. For instance, Han et al. (2014) modelled travel time uncertainty on each arc by assigning multiple uncertainty intervals with associated probabilities – for example, 10–15 minutes (normal traffic) with a probability of 0.3, and 25–30 minutes (heavy traffic) with a probability of 0.7. Subsequently, they defined scenarios whereby each arc travel time takes a value in a single uncertainty interval. They determined the optimal solution that minimizes the total travel cost plus the expected worst-case penalty for excess route durations over all scenarios. Apart from the above works, other researchers have sought to minimize, for example, the probability of route failure (Dai et al., 2018, Srinivasan et al., 2014), a weighted sum of the expectation and variance of the total travel time (Lecluyse et al., 2009, Rostami et al., 2021), or the expected completion time of all the routes, i.e., makespan (Bakach et al., 2021, Kenyon and Morton, 2003).

Comprehensive summaries of the literature on SVRPs are available in the surveys of Gendreau et al. (1996), Berhan et al. (2014), Ritzinger et al. (2016), and Oyola et al. (2017, 2018). Oyola et al. (2018) observed that only 19% of the SVRP articles published between 1996 and 2016 deal with stochastic travel times, while 49% deal with stochastic demands. We argue that the modelling of travel time variability deserves more attention in the literature, especially concerning SDVRPs. An extremely minimal number of articles focusing on an SDVRP under stochastic travel times are available: Sun and Wang (2015) (RO), Chu et al. (2017) (SPR), and Li and Chung (2019) (RO). To our knowledge, our

proposed CCP framework is the first of its kind for SDVRPs.

### 5.3 Travel Time Distributions

Distributions of travel time variability have been a topic of interest for over seventy years, with one of the earliest studies being that of Berry and Belmont (1951). They found that vehicle speed distributions are approximately normal except when the traffic volume exceeds the capacity, in which case speeds tend to be right-skewed. Travel times for free-moving traffic, taken as reciprocals of speeds, were observed to be heavily right-skewed, with the skewness decreasing as delays increase. During the same time, Turner and Wardrop (1951) discovered that journey times in Central London did not follow a normal distribution, as was commonly assumed at the time. A year later, Wardrop (1952) identified that travel times follow a skewed distribution. This was confirmed by Herman and Lam (1974), who analysed urban travel time data in a study of work trip journey times in Detroit. They observed significant skewness to the right in the observations and suggested either the lognormal or the gamma distribution to model travel time variations. Richardson and Taylor (1978) found that travel times for trips to and from central Melbourne during the morning and evening rush hours were more accurately represented by a lognormal distribution than a normal distribution. Meanwhile, Polus (1979) analysed travel time data collected in Chicago and found that the gamma distribution fitted the data better than the normal and lognormal distributions. A lognormal or gamma distribution was also proposed by Dandy and McBean (1984).

In this section, we will summarize some key works featuring continuous travel time distributions that followed these early studies and classify them by the type of distribution, starting with the normal distribution. Several researchers, even recently, have assumed that travel times follow a normal distribution. These include Kenyon and Morton (2003), Li et al. (2010a), Seshadri and Srinivasan (2010), Chen et al. (2012a), Zhang et al. (2012), Ehmke et al. (2015), Adulyasak and Jaillet (2016), Miranda and Conceição (2016), Caceres et al. (2017), Wang and Haghani (2020), Bomboi et al. (2022), and Guan and Wang (2024). While assuming normality provides several analytical and computational benefits, it imposes unreasonable restrictions such as symmetry, no minimum, and non-zero probability for negative travel times. Zhang et al. (2013) stated that routing methods relying on normality assumptions may lead to large estimation errors. Some existing research has also attempted to sample travel times from a truncated normal distribution (for example, with a cut-off at zero), such as the works of Thompson et al. (2011), Lee et al. (2012), Yan et al. (2014), Zou et al. (2014),

Yan et al. (2015), and Babaei and Rajabi-Bahaabadi (2019).

A distribution that is appropriate for modelling travel time data is the lognormal distribution. This distribution is defined only for positive values and is right-skewed. Taniguchi et al. (2001) argued that a travel time can be accurately represented by a lognormal distribution since there is a minimum travel time (corresponding to free-flow speed), after which the time increases rapidly to a maximum and then gradually decreases with a long upper tail. Using real-life detector data in Orlando, Florida, Emam and Al-Deek (2006) concluded that the lognormal distribution is superior to the normal, Weibull, and exponential distributions. Pu (2011) showed that the four different shapes of travel time distributions depicted by van Lint et al. (2008) (free-flow, congestion onset, congestion, congestion dissolve) all resemble a lognormal distribution. Zou et al. (2014) claimed that a lognormal distribution is useful for analysing extreme values since it can possess a heavy tail. They also concluded that it gives lower prediction errors and is thus more desirable than a truncated normal distribution. In addition, Kieu et al. (2015) analysed data from eight bus routes in Brisbane and recommended the lognormal distribution (over the Burr, gamma, normal, and Weibull) for describing travel times on urban corridors. Chen et al. (2017a) compared the normal, gamma, and lognormal distributions and found that the lognormal distribution outperforms the other distributions and represents 81.64% variability in the travel times observed in Wuhan, China. Other researchers who have modelled travel times by lognormal distributions include Arroyo and Kornhauser (2005), Kaparias et al. (2008), Lecluyse et al. (2009), Rakha et al. (2010), Ji et al. (2011), Zhang et al. (2013), Li et al. (2017), Chen et al. (2018b), and Babaei and Rajabi-Bahaabadi (2019).

The two-parameter lognormal distribution is bounded below by zero, resulting in unrealistic free-flow speeds. This drawback of the lognormal distribution has redirected focus to the three-parameter *shifted lognormal* (SLN) distribution, which includes a positive shift parameter to allow the distribution to start from a non-zero value. For example, Srinivasan et al. (2014) fitted different distributions (normal, lognormal, SLN, beta, Burr, gamma, and log-logistic distributions) to the travel times along arterial roads belonging to four facility types in the city of Chennai. They found that the SLN distribution is reasonable for all facility types (estimates have an average error of less than 1%) and is superior to the normal and lognormal distributions for almost all facility types. Moreover, they showed that other distributions, such as the Burr and log-logistic, can also capture well the positive skew of observed travel times. However, they stated that they selected the SLN distribution since it enables a straightforward representation of correlations across links and the typically observed codependence between the

mean and variance of travel times. Based on Srinivasan et al. (2014)'s work, Dai et al. (2018) employed the SLN distribution while studying ten links in Hangzhou, China. They observed through the Kolmogorov-Smirnov goodness-of-fit test that the SLN distribution fits all links well. Rajabi-Bahaabadi et al. (2021) fitted different distributions (normal, lognormal, SLN, gamma, shifted gamma, Burr, and Weibull distributions) to travel time data in Seattle, Washington, and San Diego, California. They concluded using the Akaike Information Criterion that the SLN and Burr distributions outperform the others for both study sites. Barahimi et al. (2022) also showed the capability of the SLN distribution for modelling travel times under several assumptions, such as correlations between link travel times, uncertain traffic flow across links, and uncertain route selection of transit users.

As noted above, the Burr distribution is an alternative distribution that provides a reasonable fit to travel time data. Taylor and Susilawati (2012) stated that this distribution has a flexible shape and is well-behaved algebraically since its cumulative distribution function has a closed-form expression. Susilawati et al. (2013) analysed two urban arterial road corridors in Adelaide, Australia, and tested the normal, lognormal, Burr, generalized Pareto, Weibull, and gamma distributions. They showed that the Burr distribution (as well as the generalized Pareto, to a lower extent) can represent heavily right-skewed travel times and that the lognormal, Weibull, and gamma distributions, although having positive skewness, do not fully represent the very long upper tails in the observed data. Chen and Fan (2020) tested the lognormal, Weibull, gamma, and Burr distributions on data collected on an interstate highway in Charlotte, North Carolina. Their chi-square test results indicate that the Burr distribution can provide the highest goodness-of-fit when considering different times of day, days of week, and weather conditions. In addition, Low et al. (2021) used automatic vehicle location data of ten bus routes in Klang Valley, Malaysia, to examine the normal, lognormal, Weibull, gamma, generalized Pareto, and Burr distributions. They found that the Burr distribution is the best-fitting for both weekday and weekend travel time observations, with the lognormal distribution emerging as the second best-fitting distribution. The Burr distribution also features in the works of Guessous et al. (2014), Gómez et al. (2016), Ma et al. (2016), Taylor (2017a), and Khoo et al. (2021). A disadvantage of this distribution is the complexity involved in handling the sum of Burr-distributed random variables, especially when they are correlated. The sum does not follow any known standard distribution, making analytical treatment challenging. The Burr distribution's heavy tail and the nature of its probability density function, which can exhibit power-law-like behaviour, further complicate approximation efforts.

For example, moment-based approximations become computationally intensive because accurately modelling the heavy-tailed and highly skewed characteristics requires consideration of higher-order moments such as skewness and kurtosis.

As noted by Gómez et al. (2016) and suggested in this review, there is no universally accepted distribution for modelling travel times. Other less commonly applied distributions include the (shifted) gamma (Oyola, 2019, Russell and Urban, 2008, Sun et al., 2021, Taş et al., 2013, Wang and Lin, 2017, Wu and Geistefeldt, 2014), generalized extreme value (Chepuri et al., 2018, Harsha and Mulangi, 2022, Zhang et al., 2019), exponential (Bouyahia et al., 2018, Sarmiento et al., 2015), Weibull (Al-Deek and Emam, 2006), generalized Pareto (Lei et al., 2014), and log-logistic (Durán-Hormazábal and Tirachini, 2016, Xue et al., 2011) distributions. Furthermore, multimodality in travel time distributions has also been observed in certain studies (Chen et al., 2014b, Dong and Mahmassani, 2009, Guo et al., 2010). Susilawati et al. (2013) stated that for urban arterial roads, short link length and the presence of traffic signals may be factors that contribute to bimodality. Mixture distributions are better suited to model multimodal travel times than single distributions. Examples include using a mixture of normal (Yang and Wu, 2016), lognormal (Kazagli and Koutsopoulos, 2013), gamma (Kim and Mahmassani, 2015), generalized extreme value (Ansari Esfeh et al., 2020), and Burr (Cheok et al., 2024) distributions.

In this research, we have opted to use the SLN distribution for three reasons: (i) the SLN distribution generalizes the lognormal distribution which has domain  $(0, \infty)$ , by including a positive shift parameter  $\gamma$  to achieve domain  $(\gamma, \infty)$  (travel times that are very close to zero are unlikely); (ii) as seen above, existing research indicates that the SLN distribution fits travel times well; and (iii) the SLN distribution is algebraically tractable, especially when it comes to expressing and interrelating the mean and standard deviation, calculating percentiles, summing SLN distributions, and capturing correlations between SLN distributions.

## 5.4 Travel Time Reliability Metrics

The travel time experienced while navigating a road network can be divided into two components: *free-flow time*, the time spent encountering negligible traffic, and *additional time*, the extra time due to variations in the traffic conditions (Carrion and Levinson, 2012). The nature of such variations may be predictable or unpredictable, giving rise to what are called *recurrent* and *non-recurrent congestion*. According to the Federal Highway Administration (FHWA) in the United States Department of Transportation, recurrent (excepted) congestion accounts for around half of the congestion encountered by transit users. Causes of recurrent

congestion include insufficient infrastructure to support the increasing number of vehicles, routine traffic volumes during peak hours, and poorly designed traffic control measures such as speed bumps, stop signs, and train crossings. Non-recurrent congestion is caused by unexpected circumstances such as a temporary reduction in road capacity (e.g., a blocked lane due to a crash, vehicle breakdown, or short-term construction work), obscured drivers' visibility (due to fog, snow, torrential rain, extreme wind gusts, or direct sunlight), or a sudden spike in demand (due to a special event such as a conference, concert, or sports).

*Travel time reliability* (TTR) is associated with unpredictable variations, i.e., delays resulting from non-recurrent congestion. It measures “the consistency or dependability in travel times, as measured from day-to-day and/or across different times of the day” (Taylor, 2017b). Suppose a person drives to work every morning. If they have a consistent 45-minute commute, then their commute is reliable. However, if they are unsure from day to day whether it will take 30 minutes or an hour to arrive at work, then the commute is unreliable. They lack TTR, which refers to how consistently one arrives at work within an expected time frame.

Herman and Lam (1974) and Sterman and Schofer (1976) were the first to acknowledge the necessity of modelling TTR. Since then, several studies have concluded that although travel time influences travellers' route choice behaviours, TTR can play a more significant role (e.g., the works of Lam and Small (2001), Hollander (2006), Asensio and Matas (2008), Li et al. (2010b), and Carrion and Levinson (2013)). Generally, travellers are not only interested in how long a route will take them but are even more concerned with the reliability of their travel time predictions. Travellers are more likely to recall the few days they experienced excessive delays rather than their experiences of the typical travel times. Therefore, they tend to be risk-averse and prefer a longer route on average over a faster and less reliable route (van Lint and van Zuylen, 2005). TTR interests all stakeholders of a transportation system, not only the travellers but also the planners and service providers. Planners must create a reliable transportation system that attracts users and is resilient and sustainable. For service providers, reliability serves as an indicator for evaluating the performance of the transportation system, for example, to assess whether the service is staying on schedule and meeting users' expectations, as well as user retention. An unreliable transportation system negatively impacts social and economic development, environmental quality, safety, and public health. Some effects include increased wear and tear and depreciation of vehicles, fuel consumption, pollutant emissions, and lower concentration due to exhaustion, irritation, and stress.

Significant research has concentrated on establishing suitable metrics for quantifying TTR. Most of these metrics relate to characteristics of the day-to-

day travel time distributions and hence provide a more accurate representation of a commuter’s experience than merely using average (“typical”) travel time. According to Banik et al. (2021), TTR metrics can be broadly categorized as (i) statistical range measures, (ii) buffer time measures, (iii) tardy trip measures, (iv) probabilistic measures, and (v) congestion-based measures. Here, we provide a brief overview of each category; for more details, the reader is referred to Lomax et al. (2003) and Taylor (2017b). Zang et al. (2022) also provided a recent summary of the methodological developments of modelling TTR in transportation networks.

Statistical range measures include the *standard deviation* and *percent variation*. For example, to hedge against travel time variability, a travel time window can be defined as the average travel time plus or minus a factor multiplied by the standard deviation. However, such a window implicitly assumes that travel times are symmetrically distributed and treats early and late arrivals with equal weights. In practice, transit users care much more about late arrivals. The percent variation is the coefficient of variation (standard deviation divided by mean) expressed as a percentage. It is a measure of dispersion that indicates the variation in travel times relative to the mean. Statistical range measures can be challenging for a non-technical audience to understand. Moreover, travel time distributions are generally skewed, implying that measures based on mean and standard deviation will likely be biased and sensitive to outliers.

Buffer time measures can be communicated more clearly and effectively with the general public than statistical range measures. These include the *buffer time index* (BTI) and *planning time index* (PTI). The BTI is calculated as

$$\text{BTI} = \frac{\text{95th Percentile Travel Time} - \text{Average Travel Time}}{\text{Average Travel Time}}.$$

The numerator called *buffer time*, represents the extra time a transit user must budget over the average travel time to reach their destination on time 95% of the time (equivalent to being late once every twenty working days). A high buffer time indicates frequent delays, i.e., the travel time distribution has a long upper tail. The California Department of Transportation (Caltrans) classifies <20% BTI as reliable travel, 20-40% BTI as moderately unreliable travel, and >40% BTI as unreliable travel. Pu (2011) recommends using the median rather than the average travel time in the BTI definition because the average-based BTI underestimates TTR, especially for heavily skewed travel time distributions. Elefteriadou (2014) claims that the BTI may give unstable indications of reliability as it can move in a direction opposed to the percentile and average travel times. Suppose the percentile and average travel times for one route are 920.77s and 668s (giving

BTI = 0.38) and for another, 443.67s and 260s (giving BTI = 0.71). The BTI is much lower for the former route even though the travel times are significantly higher.

While the BTI indicates the *extra* time necessary to ensure on-time arrival in 95% of the trips, the PTI indicates the *total* time necessary to ensure on-time arrival in 95% of the trips. The PTI is defined as

$$\text{PTI} = \frac{\text{95th Percentile Travel Time}}{\text{Free-Flow Travel Time}},$$

where the free-flow travel time is calculated as the road length divided by the maximum speed limit. Similar to the BTI, a higher PTI indicates more unreliable conditions. The difference between the BTI and PTI is that the former captures the variation between the average and near-worst case travel times (i.e., concerns late arrivals), while the latter captures the variation between the lowest and near-worst case travel times (i.e., concerns both early and late arrivals). Lyman and Bertini (2008) observed that the BTI and PTI appear to dilute and exaggerate trends, respectively. The BTI and PTI are used by entities such as the FHWA, Georgia Department of Transportation, Florida Department of Transportation, Southern California Association of Governments, Maryland State Highway Administration, Minnesota Department of Transportation, and Washington State Department of Transportation. Note that one can set different reliability levels (apart from 95%) in the BTI and PTI definitions.

The third category involves tardy trip measures that attempt to answer the question “How frequently will a traveller arrive unacceptably late?”. Two examples of tardy trip measures are the *on-time arrival percentage* and *misery index* (MI). The former is the number of trips arriving on time relative to a pre-established threshold travel time (e.g., 110% of the average or median travel time). The MI focuses on the extra travel time experienced during the worst (most delayed) trips. It is calculated as

$$\text{MI} = \frac{\text{Average Travel Time of the 20\% Worst Trips} - \text{Average Travel Time}}{\text{Average Travel Time}}.$$

Therefore, the MI measures how bad the worst trips are. Note that 20% signifies the worst day of the working week; however, one can decrease this percentage as deemed fit. The MI suffers from the same disadvantages as the BTI.

Similar to the on-time arrival percentage, probabilistic measures use a threshold to differentiate between reliable and unreliable trips. Thus, their effectiveness relies heavily on this parameter, which varies by application and context. For example, Asakura (1999) evaluates reliability as the probability that a trip is suc-

cessful within a specified time interval. Such a probabilistic measure has similar implications to the on-time arrival percentage. The final category, congestion-based measures, includes metrics such as the *frequency of congestion* (FOC) and *percentage of travel under congestion* (PTC). The FOC measures the percentage of days or times when congestion exceeds some threshold (e.g., the percentage of days or times when the average route speed falls below a limit). The PTC measures the percentage of the vehicle miles travelled (i.e., the total number of miles driven by all vehicles within a specific area over a given period) under congested conditions.

Banik et al. (2021) compared TTR metrics across all categories and noted that different metrics provide different indications of reliability. Srinivasan et al. (2014) claims that the “choice of metric depends on [the] context of study, ease of measure and computation, ease of understanding by user and operator, relevance to decision making, and the ability to represent users['] attitude towards risk”. In our case, we opt to use a percentile travel time for the TTR metric since this is straightforward to the non-technical community, captures the skewed nature of travel times, and is useful for trips that require on-time arrivals (Lomax et al., 2003). Students must arrive on time at school since this allows them to settle in, connect with their peers, focus academically, and develop positive habits such as punctuality, commitment, and respect. Here, we focus primarily on the 99th percentile travel time to reflect extreme but plausible delays. In a 180-day academic year, the 99th percentile corresponds to approximately two school days with the worst travel delays. In cases where achieving this reliability level is infeasible due to the network structure, we instead consider the 95th percentile travel time, corresponding to approximately nine school days per year. This provides a practical alternative reliability benchmark while still capturing high-delay conditions. Lomax et al. (2003) classifies percentile travel times under buffer time measures, and Chen et al. (2003b) states that a percentile travel time provides a meaningful combination of the average travel time and its variability into a single figure. Given that evidence has shown the asymmetric characteristic of travel time distributions, analysis of extreme values is critical for evaluating TTR (Zang et al., 2022). The 95th percentile travel time is, in fact, one of the four TTR metrics recommended by Taylor (2017b).

## 5.5 Proposed Framework

As discussed in Section 5.1, it is desirable to construct robust school bus routes that perform well under various traffic density levels. This ensures that the transport service is reliable and attractive to the public. To account for travel time

variability, we require amendments to both the problem definition presented in Chapter 3 and the heuristic algorithm put forward in Chapter 4. Before presenting these amendments in Section 5.5.2, we first discuss the distributional assumptions underlying our proposed framework.

### 5.5.1 Distributional Assumptions

In this work, we propose incorporating a travel time distribution for each arc in the set  $E_1$  (as defined in Section 3.2.1). Specifically, we let the travelling time  $T_e$  along each arc  $e = (u, v) \in E_1$  be a random variable, rather than being fixed at some prespecified shortest driving time  $t(u, v)$ , as was the case previously.

#### 5.5.1.1 Arc Distributional Assumptions

The random travel time  $T_e$  along an arc  $e \in E_1$  is assumed to follow an SLN distribution

$$T_e \sim \mathcal{LN}(\mu_e, \sigma_e^2, \gamma_e)$$

or, equivalently,

$$\ln(T_e - \gamma_e) \sim \mathcal{N}(\mu_e, \sigma_e^2),$$

where  $\gamma_e > 0$  is the shift parameter,  $\mathcal{N}$  denotes the normal distribution, and  $\mu_e \in \mathbb{R}$  and  $\sigma_e > 0$  are the mean (location) and standard deviation (scale) of the random variable  $\ln(T_e - \gamma_e)$ , respectively. The probability density function (PDF) of random variable  $T_e$  is given by

$$\frac{1}{(t_e - \gamma_e)\sigma_e\sqrt{2\pi}} \exp\left[-\frac{(\ln(t_e - \gamma_e) - \mu_e)^2}{2\sigma_e^2}\right], \quad t_e > \gamma_e$$

and  $T_e$  can also be expressed as

$$T_e = \gamma_e + \exp[\mu_e + \sigma_e Z_e],$$

where  $Z_e \sim \mathcal{N}(0, 1)$  is the standard normal variable. In our context, the shift parameter  $\gamma_e$  represents the free-flow component of  $T_e$ , while the other component of  $T_e$  represents the excess travel time,

$$M_e := \exp[\mu_e + \sigma_e Z_e] \sim \mathcal{LN}(\mu_e, \sigma_e^2, 0).$$

The mean and variance of travel time  $T_e$  are given by:

$$\mathbb{E}[T_e] = \gamma_e + \exp\left[\mu_e + \frac{1}{2}\sigma_e^2\right] = \gamma_e + \mathbb{E}[M_e],$$

$$\text{Var}[T_e] = \exp [2\mu_e + \sigma_e^2] (\exp [\sigma_e^2] - 1),$$

and thus

$$\gamma_e = \mathbb{E}[T_e] - \mathbb{E}[M_e], \quad (5.1)$$

$$\sigma_e = \sqrt{\ln \left( 1 + \frac{\text{Var}[T_e]}{\mathbb{E}[M_e]^2} \right)}, \quad \text{and} \quad (5.2)$$

$$\mu_e = \ln(\mathbb{E}[M_e]) - \frac{1}{2}\sigma_e^2. \quad (5.3)$$

Here, we assume that

$$\begin{aligned} \mathbb{E}[M_e] &= \alpha \mathbb{E}[T_e], \\ \sqrt{\text{Var}[T_e]} &= \beta \mathbb{E}[T_e], \end{aligned}$$

with  $\alpha$  representing the portion of the mean travel time that is attributed to congested conditions and  $\beta$  representing the coefficient of variation (CoV) of the travel time distribution. Hence, it follows from Equations (5.1)-(5.3) that

$$\gamma_e = (1 - \alpha) \mathbb{E}[T_e], \quad (5.4)$$

$$\sigma_e = \sqrt{\ln \left( 1 + \frac{\beta^2}{\alpha^2} \right)}, \quad \text{and} \quad (5.5)$$

$$\mu_e = \ln(\alpha \mathbb{E}[T_e]) - \frac{1}{2}\sigma_e^2. \quad (5.6)$$

Figure 5.1 displays the PDFs corresponding to five different combinations of the parameters  $\alpha$  and  $\beta$  for an arc  $e$  with  $\mathbb{E}[T_e] = 1000\text{s}$ . From this figure, it is clear that, for fixed  $\beta$ , increasing  $\alpha$  causes a decrease in the shift parameter (green, blue, and red minimum travel times are 900s, 800s, and 700s, respectively) and an increase in variability (evidenced by a heavier tail). Furthermore, increasing  $\beta$  while keeping  $\alpha$  constant also leads to a distribution with higher variability.

### 5.5.1.2 Route Distributional Assumptions

Since the travel time along any arc is assumed to follow an SLN distribution, the total travel time distribution of a route with at least two arcs (i.e., at least two visited bus stops) is the sum of the individual SLN distributions of the arcs. Unfortunately, no known closed-form expression exists for the exact PDF of a sum of lognormal random variables. However, as stated by Zhang (2022), it is widely accepted that a sum of lognormal random variables can be well approximated

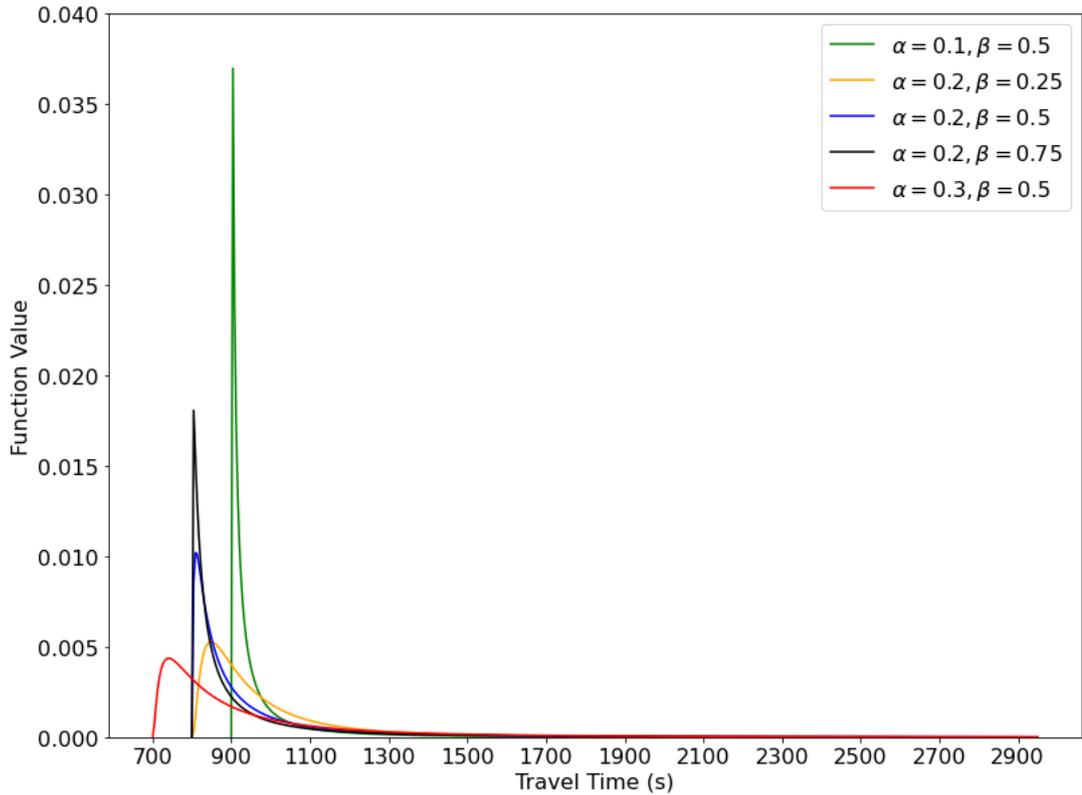


Figure 5.1: Different probability density functions (varying parameters  $\alpha$  and  $\beta$ ) for an arc with mean travel time 1000s.

by a single lognormal random variable. Even though no explicit mathematical justification has been provided for such an approximation, it has attracted considerable attention, with different methods developed to approximate the resultant lognormal distribution. These include the methods by Fenton (1960), Schwartz and Yeh (1982), Barakat (1976), Beaulieu and Xie (2004), and Wu et al. (2005), which we next describe briefly. Note that the original proposals of these methods were concerned with sums of *independent* lognormal random variables.

Fenton (1960) proposed a moment-matching technique called the Fenton-Wilkinson method that matches the first and second central moments (mean and variance) of the lognormal sum and the approximating lognormal random variable. Schwartz and Yeh (1982) proposed a similar procedure called the Schwartz-Yeh method that matches the mean and variance of the logarithms of the lognormal sum and the approximating lognormal random variable. Their approach is iterative, considering two summands at a time and assuming that each successive sum is lognormal. Beaulieu et al. (1995) and Wu et al. (2005) observed that the Fenton-Wilkinson method accurately tracks the tail portion of the lognormal sum's PDF (but not the head), while the Schwartz-Yeh method is better at tracking the head portion (but not the tail).

To achieve a trade-off in accurately approximating different regions of the log-

normal sum's PDF, Wu et al. (2005) suggested matching the moment-generating function (MGF) of the lognormal sum with that of the approximating lognormal random variable at two real and positive points  $s_1$  and  $s_2$ . They exploited the property that the MGF of a sum of independent random variables is the product of the MGFs of the individual random variables. Moreover, Wu et al. (2005) viewed the MGF as a weighted integral of the PDF and declared that the weight function can be adjusted, as required, to focus on different regions of the PDF. For selected  $s_1$  and  $s_2$ , their results indicate that the proposed method is more accurate at matching the head portion than both the Fenton-Wilkinson and the Schwartz-Yeh methods and is comparable to the Fenton-Wilkinson method in approximating the tail portion.

Similar to Wu et al. (2005), Barakat (1976) utilized the property that the characteristic function (CF, which is a special case of the MGF) of a sum of independent random variables is the product of the CFs of the individual random variables. He computed the individual CFs numerically using a Taylor series expansion and then applied an inverse Fourier transform to their product to approximate the lognormal sum's PDF. Mehta et al. (2007) claim that Barakat (1976)'s method is difficult (because of the oscillatory nature of the Fourier integrand and the gradual decay of the lognormal PDF's tail), inaccurate, and does not provide analytical expressions of the approximate distribution.

Beaulieu and Xie (2004) also worked with CFs and an inverse Fourier transform but employed a modified version of the Clenshaw-Curtis method (Clenshaw and Curtis, 1960) for the numerical computation of the CFs and the lognormal sum's cumulative distribution function (CDF). The latter was then plotted on lognormal probability paper, on which a lognormal CDF appears as a straight line. To best fit the lognormal sum's CDF, they applied the minimax approximation that minimizes the maximum absolute deviation between the approximate and true functions over a given interval. This minimization significantly weighs the relative error of the approximation in the distribution's tails. Mehta et al. (2007) state that "while the method is optimal in the minimax sense on lognormal paper, this does not imply optimality in directly matching the probability distribution".

Extensions of the former three methods mentioned above were proposed for sums of *correlated* lognormal random variables: the Fenton-Wilkinson method (Abu-Dayya and Beaulieu, 1994, Graziosi et al., 2001), the Schwartz-Yeh method (Safak, 1993), and Wu et al. (2005)'s MGF-matching method (Mehta et al., 2006). To our knowledge, the CF-based methods by Barakat (1976) and Beaulieu and Xie (2004) are limited to the case of independent lognormal random variables. We also highlight that several other methods exist for approximating the distribution

of a lognormal sum. Some use a single approximating lognormal random variable similar to the methods described above, some use a single non-lognormal random variable, and others are based on bound analysis. For instance, Schleher (1977) proposed a method based on cumulants matching that partitions the lognormal sum’s CDF range into three segments, with each segment approximated by a different lognormal random variable (Abu-Dayya and Beaulieu (1994) provided the correlated extension). Bounds on the lognormal sum’s CDF have also been derived; for example, Farley’s approximation (Beaulieu et al., 1995) for independent and identically distributed summands, Slimane (2001)’s bounds for independent summands, and Berggren and Slimane (2004)’s bound that can handle correlated summands. However, such bounds do not directly describe the shape of the approximating PDF and “can be quite loose for certain typical parameters” (Wu et al., 2005). For a survey of the various methods to approximate the distribution of a sum of lognormal random variables, the reader is referred to the work of Zhang (2022).

Beaulieu et al. (1995) compared four methods (Fenton-Wilkinson, Schwartz-Yeh, Schleher (1977)’s cumulants matching, and Farley’s methods) and concluded that all methods have their advantages and disadvantages and that no method is superior to the others. Hcine and Bouallegue (2015) state that a method’s accuracy depends heavily on the region of the lognormal sum’s distribution being examined. They state that, for example, the Fenton-Wilkinson method achieves high accuracy for CDF values larger than 0.9.

In this work, we opted to employ the Fenton-Wilkinson method for four reasons: (i) our route TTR metric is a high percentile, i.e., a value located far in the tail of the distribution, (ii) the method has an extension that can handle correlated summands, (iii) the method and its correlated extension have been shown to accurately estimate large values of the CDF, i.e., high percentiles (for example, by Abu-Dayya and Beaulieu (1994), Song and Kim (2022), and Wei et al. (2024)), and (iv) the method offers a computationally simple way of deriving closed-form expressions for the parameters of the approximating PDF. These parameters are required to compute statistical measures such as the mean, standard deviation, and different percentiles. For route travel time modelling, the Fenton-Wilkinson method has also been adopted by Srinivasan et al. (2014), Chen et al. (2018a), Dai et al. (2018), Rajabi-Bahaabadi et al. (2021), and Barahimi et al. (2022).

Therefore, for each route  $R \in \mathcal{R}$ , we use the Fenton-Wilkinson method to approximate the distribution  $\mathcal{LN}(\mu_R, \sigma_R^2, \gamma_R)$  of the total *travel* time. Based on this approach, the parameters of the approximating distribution satisfy:

$$\gamma_R = \sum_{e \in R} \gamma_e, \tag{5.7}$$

$$\exp \left[ \mu_R + \frac{1}{2} \sigma_R^2 \right] = \sum_{e \in R} \exp \left[ \mu_e + \frac{1}{2} \sigma_e^2 \right], \text{ and} \quad (5.8)$$

$$\begin{aligned} \exp [2\mu_R + \sigma_R^2] (\exp [\sigma_R^2] - 1) &= \sum_{e_1, e_2 \in R} \left( \rho_{e_1 e_2} \sqrt{\exp [2\mu_{e_1} + \sigma_{e_1}^2] (\exp [\sigma_{e_1}^2] - 1)} \right. \\ &\quad \left. \times \sqrt{\exp [2\mu_{e_2} + \sigma_{e_2}^2] (\exp [\sigma_{e_2}^2] - 1)} \right), \end{aligned} \quad (5.9)$$

where  $\rho_{e_1 e_2}$  denotes the correlation between the travel times on arcs  $e_1, e_2 \in E_1$ . Note that correlations  $\rho_{ee} = 1, \forall e \in E_1$ , and, in this chapter, it is assumed that  $\rho_{e_1 e_2} = 0, \forall e_1, e_2 \in E_1, e_1 \neq e_2$  (independent arcs). By Equations (5.4)-(5.9), the parameters  $\mu_R, \sigma_R^2, \gamma_R$  are, in our case, given by

$$\gamma_R = (1 - \alpha) \sum_{e \in R} \mathbb{E}[T_e], \quad (5.10)$$

$$\sigma_R^2 = \ln \left( 1 + \frac{\beta^2}{\alpha^2} \cdot \frac{\sum_{e_1, e_2 \in R} \rho_{e_1 e_2} \mathbb{E}[T_{e_1}] \mathbb{E}[T_{e_2}]}{(\sum_{e \in R} \mathbb{E}[T_e])^2} \right), \text{ and} \quad (5.11)$$

$$\mu_R = \ln \left( \alpha \sum_{e \in R} \mathbb{E}[T_e] \right) - \frac{1}{2} \sigma_R^2. \quad (5.12)$$

To incorporate bus dwell times (as defined in the second parenthesis in Equation (3.5)) and get a *journey* time distribution for route  $R \in \mathcal{R}$ , we can shift the total travel time distribution to the right by the total bus dwell time. We can do this since the service times here are assumed to be deterministic, and because a horizontal shift of a distribution does not alter its shape. Thus, given a route  $R = (v_1, v_2, \dots, v_j, v_0)$ , the approximate journey time distribution is

$$T(R) \sim \mathcal{LN} \left( \mu_R, \sigma_R^2, \gamma_R + \sum_{i=1}^j (15 + 5s(v_i, R)) \right). \quad (5.13)$$

### 5.5.2 Problem and Algorithm Modifications

As remarked at the beginning of Section 5.5, modifications in the problem definition and heuristic algorithm are needed to allow arc travelling times to be described by random variables instead of fixed values.

The first modification is to Constraints (3.4). Recall that these ensure that the fixed total journey time  $t(R)$  of each route  $R \in \mathcal{R}$  is at most the maximum journey time  $m_t$ . To capture travel time variability and reliability, we replace the constraint

$$t(R) \leq m_t$$

by an on-time arrival chance constraint of the form

$$\mathbb{P}[T(R) \leq m_t] \geq \kappa, \quad (5.14)$$

where  $\kappa \in (0, 1)$  is a desired reliability level (typically  $\kappa \geq 0.9$ ). Constraint (5.14) can be rewritten as

$$P_{100\kappa}(R) \leq m_t, \quad (5.15)$$

where  $P_{100\kappa}(R)$  denotes the  $(100\kappa)^{\text{th}}$  percentile of the journey time distribution of route  $R$ . Such a constraint models users' aversion to late arrivals by ensuring that in  $(100\kappa)\%$  of the cases, route  $R$  should not take longer than  $m_t$ . For the SLN distribution given in Equation (5.13), the  $(100\kappa)^{\text{th}}$  percentile is given by

$$P_{100\kappa}(R) = \gamma_R + \sum_{i=1}^j (15 + 5s(v_i, R)) + \exp[\mu_R + \Phi^{-1}(\kappa)\sigma_R],$$

where  $\Phi$  denotes the CDF of the standard normal distribution. We note that Constraint (5.15) imposes a hard cut-off on the allowable lateness, which may lead to infeasibility when no solution reaches the specified reliability level  $\kappa$ . Indeed, a limitation of this formulation is that it does not distinguish between routes that marginally violate the threshold and those that perform substantially worse. A softer modelling approach therefore represents a natural direction for future work, as discussed in Section 7.3.6.

The second modification is to the cost function, which is altered to shift the focus towards improving service reliability. In particular, the fixed total journey time  $t(R)$  is replaced everywhere by  $P_{100\kappa}(R)$  in Equations (3.6) and (3.7). Therefore, instead of minimizing the expected total journey time (which assumes risk neutrality), we minimize the percentile as a worst-case representative of the entire journey time distribution.

We now proceed to the modifications required in Algorithm 1. A key modification here is in the selection of the visited bus stops in Steps 2 and 7. Before this selection is carried out, we first need to consider each bus stop  $v \in V_1$  and check whether the journey time distribution of the route  $R = (v, 0)$ , with one student boarding at stop  $v$  (i.e.,  $s(v, R) = 1$ ), has  $P_{100\kappa}(R)$  exceeding  $m_t$ . If this is the case, then bus stop  $v$  should not be considered for selection since no route visiting  $v$  is feasible with respect to Constraint (5.15), not even the route that visits  $v$  and collects a single student. Note that, for selected parameters  $\alpha, \beta, \kappa$ , this process may cause infeasibility in the sense that at least one student would not have a bus stop that can be feasibly visited and yet is within maximum walking distance  $m_w$ . For example, assume that a student can only walk to stop  $v$  and that arc  $e = (v, 0)$  has mean travel time  $\mathbb{E}[T_e] = 1000\text{s}$  (refer to Figure 5.1). Table 5.1 shows that for the last three selections of  $\alpha, \beta, \kappa$ , the percentile journey time of

Table 5.1: Different percentile journey times in seconds (varying parameters  $\alpha$ ,  $\beta$ , and  $\kappa$ ) for a route containing a single arc with mean travel time 1000s and one boarding student.

$\alpha$	$\beta$	$\kappa$	$P_{100\kappa}(R)$
0.1	0.5	0.95	1301.88
0.2	0.25	0.95	1436.10
0.2	0.5	0.95	1572.14
0.2	0.75	0.95	1593.62
0.3	0.5	0.95	1748.17
0.1	0.5	0.99	2226.66
0.2	0.25	0.99	2013.31
0.2	0.5	0.99	2782.76
0.2	0.75	0.99	3196.60
0.3	0.5	0.99	2975.71

route  $R = (v, 0)$  with  $s(v, R) = 1$  will exceed  $m_t = 2700$  seconds (45 minutes). In practice, this student would still be entitled to school transport, even though a proposed route involving stop  $v$  sometimes exceeds  $m_t$ . In such a case, two possible remedies exist: either opting for a lower  $\kappa$  in the algorithmic framework, or attempting to add a new bus stop to  $V_1$  that is closer to the school but still within the student's walking distance  $m_w$ , such that the percentile journey time does not exceed  $m_t$ . For example, for  $\alpha = 0.2, \beta = 0.75$ , the mean travel time along arc  $e = (v, 0)$  must be at most 843.67s to guarantee reliability level  $\kappa = 0.99$ .

Two final modifications in the algorithm are the following. If multiple moves give the best improvement in Step 6 of Algorithm 1, the one yielding the smallest discrepancy between the longest and shortest  $(100\kappa)^{\text{th}}$  percentile journey times is performed. Finally, in the solution repair procedure in Step 7, new stop occurrences are added at the position that causes the least increase in the  $(100\kappa)^{\text{th}}$  percentile journey time.

## 5.6 Computational Results

In this section, and throughout the remainder of the thesis, we will consider three different scenarios, which we refer to as the *deterministic*, *independent*, and *correlated* scenarios. In the deterministic scenario, each travel time is assumed to be constant and fixed at the driving time provided in the instance data, as seen in Chapters 3 and 4. In the independent scenario, which is the focus of this chapter, we assume that travel times follow an SLN distribution and are independent. Consequently,  $\rho_{e_1 e_2} = 0, \forall e_1, e_2 \in E_1, e_1 \neq e_2$ . In the correlated scenario, the independence assumption is dropped; instead, correlations between arcs are allowed, as will be discussed in Chapter 6. For comparative purposes,

we set each mean travel time in the independent and correlated scenarios to the driving time provided in the instance data.

Let us shift focus to the independent scenario, the C++ source code for which can be downloaded from (Sciortino, 2025c). Recall from Section 5.5.1.1 that parameter  $\alpha$  is the ratio of the excess travel time to the mean travel time, and parameter  $\beta$  is the CoV of the travel time of each arc. We tested nine different configurations for  $\alpha$  and  $\beta$ , with  $\alpha \in \{0.1, 0.2, 0.3\}$  and  $\beta \in \{0.25, 0.5, 0.75\}$ . The values for  $\alpha$  were chosen such that, for each arc, the deviation between the mean and minimal travel times is 10%, 20%, or 30% of the mean travel time. A similar setting (with 20%) was considered by Hu et al. (2018). In addition, our selections for  $\beta$  were informed by the work of Srinivasan et al. (2014), where CoV values were randomly drawn from the range (0,1) with average CoVs 0.25, 0.5, and 0.75, representing low, medium and high variability levels, respectively. We start by presenting and analysing the results for the configuration  $\alpha = 0.2, \beta = 0.5$ , which represents the moderate case among the tested settings and thus serves as a natural baseline for interpretation. We then present the results for the remaining configurations and highlight noteworthy differences observed across them.

### 5.6.1 Baseline Configuration

Table 5.2 displays the results for parameters  $\alpha = 0.2, \beta = 0.5$ . Recall from Chapter 4 that, for the deterministic scenario, the only instance requiring one additional route beyond the lower bound  $\underline{k}$  was the Bridgend instance. Having said that, observe from Column 4 of Table 5.2 that, for eleven instances, some of the solutions achieved for the deterministic scenario are infeasible with respect to the on-time arrival chance constraints (5.15) ( $\kappa = 0.99$ ). This means that at least one route in these solutions has a 99<sup>th</sup> percentile journey time larger than 45 minutes. For this reason, we do not present the average sum of the routes' 99<sup>th</sup> percentile journey times under Column 5 of Table 5.2 for these instances and exclude them from any future comparisons.

The last three columns of Table 5.2 correspond to the independent scenario. The settings  $\alpha = 0.2, \beta = 0.5, \kappa = 0.99$  yield infeasibility for the Suffolk, Canberra, Bridgend, and Brisbane instances. This issue, discussed in Section 5.5.2, can be resolved by lowering  $\kappa$  to 0.95, for which we have confirmed that feasible solutions exist (refer to Appendix C.1). Moreover, note that the average number of routes in the independent scenario matches that in the deterministic scenario for all instances except Porthcawl, Cardiff, and Victoria. For the former two respective instances, one and two additional routes beyond  $\underline{k}$  were required in all runs for each route's 99<sup>th</sup> percentile journey time to be at most 45 minutes.

Table 5.2: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.2, \beta = 0.5$ . TMT and TPT stand for Total Mean Time and Total 99<sup>th</sup> Percentile Time (minutes), presented as mean  $\pm$  standard deviation. Comparative results for the deterministic scenario are also presented, with the column IR indicating the number of runs for which the deterministic solution is infeasible with respect to the on-time arrival chance constraints. A hyphen indicates infeasibility.

Location	Deterministic Scenario				Independent Scenario		
	$k$	TMT	IR	TPT	$k$	TMT	TPT
Mġarr	4	54.10 $\pm$ 0.00	0	100.48 $\pm$ 0.35	4.00	54.91 $\pm$ 0.26	97.31 $\pm$ 0.22
Mellieħa	4	56.30 $\pm$ 0.00	0	102.92 $\pm$ 0.00	4.00	57.06 $\pm$ 0.28	99.73 $\pm$ 0.21
Porthcawl	1	27.83 $\pm$ 0.43	25	-	2.00	24.32 $\pm$ 0.11	52.70 $\pm$ 0.26
Qrendi	5	75.73 $\pm$ 0.00	0	136.08 $\pm$ 0.04	5.00	77.37 $\pm$ 1.50	135.84 $\pm$ 0.63
Suffolk	3	116.61 $\pm$ 0.90	25	-	-	-	-
Senglea	6	72.01 $\pm$ 0.54	0	141.76 $\pm$ 2.95	6.00	74.50 $\pm$ 1.26	138.23 $\pm$ 1.87
Victoria	4	92.96 $\pm$ 0.41	19	-	4.32	93.94 $\pm$ 1.69	155.13 $\pm$ 3.38
Pembroke	7	103.49 $\pm$ 0.68	0	202.95 $\pm$ 1.98	7.00	106.26 $\pm$ 1.24	195.23 $\pm$ 3.13
Canberra	7	179.11 $\pm$ 0.53	25	-	-	-	-
Ħandaq	6	96.61 $\pm$ 0.57	0	176.48 $\pm$ 2.28	6.00	98.31 $\pm$ 1.01	171.62 $\pm$ 2.99
Valetta	6	104.50 $\pm$ 0.78	18	-	6.00	110.11 $\pm$ 2.12	207.25 $\pm$ 4.27
Birkirkara	6	97.49 $\pm$ 2.64	3	-	6.00	97.34 $\pm$ 2.50	170.99 $\pm$ 5.74
Ħamrun	7	99.31 $\pm$ 1.40	0	186.15 $\pm$ 3.38	7.00	100.62 $\pm$ 1.26	177.80 $\pm$ 2.59
Cardiff	2	67.12 $\pm$ 0.00	25	-	4.00	67.92 $\pm$ 0.00	150.53 $\pm$ 0.00
M. Keynes	4	59.88 $\pm$ 0.57	0	116.03 $\pm$ 1.51	4.00	59.36 $\pm$ 0.56	112.93 $\pm$ 1.34
Bridgend	6	170.55 $\pm$ 0.53	25	-	-	-	-
Edinburgh-2	4	59.13 $\pm$ 0.00	0	110.48 $\pm$ 0.00	4.00	60.11 $\pm$ 0.04	110.06 $\pm$ 0.19
Edinburgh-1	9	143.71 $\pm$ 1.56	5	-	9.00	145.93 $\pm$ 2.80	266.41 $\pm$ 4.97
Adelaide	8	128.18 $\pm$ 1.45	2	-	8.00	129.49 $\pm$ 1.62	239.88 $\pm$ 2.75
Brisbane	10	213.41 $\pm$ 4.31	25	-	-	-	-

For the Victoria instance, eight runs required one additional route beyond  $k$ . More details on the computational times and feasibility rates of the independent scenario’s runs with  $\alpha = 0.2, \beta = 0.5$  can be found in Appendix C.2.

We now compare the average sums of the routes’ mean journey times (TMTs) and the average sums of the routes’ 99<sup>th</sup> percentile journey times (TPTs) of the deterministic and independent scenarios. For eight of the nine comparable instances, we found that the average TPT of the independent scenario is lower at the expense of having a higher average TMT. This trend was anticipated given that the objective of the independent scenario is to minimize the TPT, while that of the deterministic scenario is to minimize the total journey time (equivalent to TMT). For the Milton Keynes instance, the average TPT and TMT are both lower for the independent scenario. The relative deterioration in the average TMT from the deterministic to the independent scenario ranges between  $-0.87\%$  to  $3.46\%$ , whereas the relative improvement in the average TPT ranges between  $0.18\%$  to  $4.49\%$ . The trends observed from the averages across the 25 runs were confirmed by Mann-Whitney tests performed on the 25 pairs of TMTs and TPTs. The results of these tests are presented in Table 5.3, with all tests being significant

Table 5.3: Mann-Whitney test results comparing the TMTs and TPTs of the deterministic and independent scenarios. The values under the Det. and Ind. columns represent the mean ranks, and  $U$  represents the test statistic.

Location	TMT				TPT			
	$U$	$p$ -value	Det.	Ind.	$U$	$p$ -value	Det.	Ind.
Mgarr	625	< 0.001	13.00	38.00	0	< 0.001	38.00	13.00
Mellieħa	625	< 0.001	13.00	38.00	0	< 0.001	38.00	13.00
Qrendi	500	< 0.001	18.00	33.00	134	< 0.001	32.64	18.36
Senglea	625	< 0.001	13.00	38.00	129	< 0.001	32.84	18.16
Pembroke	618	< 0.001	13.28	37.72	12	< 0.001	37.52	13.48
ħandaq	599	< 0.001	14.04	36.96	64	< 0.001	35.44	15.56
ħamrun	466	0.003	19.36	31.64	13	< 0.001	37.48	13.52
Milton Keynes	146	< 0.001	32.16	18.84	25	< 0.001	37.00	14.00
Edinburgh-2	625	< 0.001	13.00	38.00	0	< 0.001	38.00	13.00

at the 0.01 level. From the mean ranks in Table 5.3, we see that the routes of the deterministic scenario are statistically significantly shorter since they have lower TMT ranks (except for the Milton Keynes instance). Meanwhile, the routes of the independent scenario are statistically significantly more reliable, as evidenced by the lower TPT ranks. This shows that reliability loss occurs when travel time variability is ignored during the route generation phase.

The best solutions for the independent scenario with  $\alpha = 0.2, \beta = 0.5$  are displayed in Table 5.4 and visualized in Appendix C.3. These solutions have the lowest TPT and, in the case of ties, TMT and, in turn, average walking time. Some of these solutions were achieved in multiple runs (Mellieħa: 3, Porthcawl: 22, Qrendi: 6, Suffolk: 8, Cardiff: 13, Milton Keynes: 14, Edinburgh-2: 24 runs). Moreover, the Qrendi, Suffolk, Victoria, and Cardiff instances have equivalent alternative solutions (in terms of TPT, TMT, and average walking time) to the ones presented in Table 5.4. Upon inspection of the solutions in this table, we have observed that 90% of the routes have a capacity utilization higher than 90%, approximately 82% have a mean journey time that is less than half the maximum journey time  $m_t = 45$  minutes, and approximately 74% have a percentile journey time that is less than 35 minutes.

Table 5.4: Best solutions for the independent scenario with parameters  $\alpha = 0.2, \beta = 0.5$ . Cap. is short for Capacity, MT stands for Mean Time (minutes), and PT stands for Percentile Time (minutes). The routes are presented in terms of the indices of the vertices in  $V_1$ , and the subscripts represent the number of boarding students.

Location	Routes	Load/Cap.	MT,PT	TMT, TPT
----------	--------	-----------	-------	-------------

Mġarr	$38_{18} \rightarrow 21_{32} \rightarrow 45_1 \rightarrow 0$	51/53	13.13,22.51	
	$52_{20} \rightarrow 57_{15} \rightarrow 35_{12} \rightarrow 15_2 \rightarrow 0$	49/53	11.52,19.29	54.50,
	$42_{10} \rightarrow 19_{11} \rightarrow 45_{20} \rightarrow 0$	41/44	15.40,30.51	97.97
	$48_1 \rightarrow 33_9 \rightarrow 11_{10} \rightarrow 12_{29} \rightarrow 0$	49/53	14.45,25.66	
Mellieħa	$30_8 \rightarrow 18_4 \rightarrow 20_{14} \rightarrow 27_6 \rightarrow 0$	32/36	14.93,26.77	
	$66_5 \rightarrow 83_9 \rightarrow 80_{14} \rightarrow 5_{21} \rightarrow 7_1 \rightarrow 0$	50/53	16.10,26.92	56.75,
	$33_7 \rightarrow 7_4 \rightarrow 9_{28} \rightarrow 0$	39/44	9.90,17.18	99.75
	$43_{14} \rightarrow 45_3 \rightarrow 46_{14} \rightarrow 59_{19} \rightarrow 0$	50/53	15.82,28.88	
Porthcawl	$79_{34} \rightarrow 38_{27} \rightarrow 0$	61/61	15.77,30.43	24.28,
	$61_5 \rightarrow 0$	5/8	8.52,22.35	52.79
Qrendi	$29_{43} \rightarrow 76_8 \rightarrow 0$	51/53	13.58,25.84	
	$124_{15} \rightarrow 126_{25} \rightarrow 155_4 \rightarrow 152_9 \rightarrow 0$	53/53	20.08,36.47	75.73,
	$89_{27} \rightarrow 76_{26} \rightarrow 0$	53/53	13.40,25.27	136.04
	$50_{15} \rightarrow 99_{19} \rightarrow 146_1 \rightarrow 54_{10} \rightarrow 0$	45/53	11.95,19.80	
	$67_3 \rightarrow 68_{18} \rightarrow 63_{20} \rightarrow 137_{12} \rightarrow 0$	53/53	16.72,28.67	
Suffolk <sup>a</sup>	$4_{34} \rightarrow 116_{18} \rightarrow 172_2 \rightarrow 5_{22} \rightarrow 98_4 \rightarrow 0$	80/80	31.08,40.43	
	$102_2 \rightarrow 67_4 \rightarrow 9_8 \rightarrow 65_{12} \rightarrow 0$	26/27	21.17,29.30	
	$78_9 \rightarrow 155_9 \rightarrow 111_{17} \rightarrow (\text{cont.})$			115.43,
	$140_8 \rightarrow 43_2 \rightarrow 0$	45/45	29.20,38.67	153.22
	$171_{29} \rightarrow 130_6 \rightarrow 173_1 \rightarrow (\text{cont.})$			
	$50_5 \rightarrow 122_{12} \rightarrow 3_5 \rightarrow 0$	58/61	33.98,44.82	
Senglea	$49_{53} \rightarrow 0$	53/53	9.80,18.85	
	$167_{39} \rightarrow 59_{14} \rightarrow 0$	53/53	13.77,25.91	
	$123_{21} \rightarrow 81_6 \rightarrow 150_{22} \rightarrow 0$	49/53	18.60,36.41	73.07,
	$182_{17} \rightarrow 120_{23} \rightarrow 59_{12} \rightarrow 0$	52/53	15.68,28.02	138.65
	$55_{11} \rightarrow 0$	11/14	2.98,6.19	
	$140_{36} \rightarrow 49_{12} \rightarrow 0$	48/53	12.23,23.28	
Victoria	$143_{13} \rightarrow 141_8 \rightarrow 152_{11} \rightarrow (\text{cont.})$			
	$121_6 \rightarrow 177_5 \rightarrow 172_9 \rightarrow 168_1 \rightarrow 0$	53/53	26.62,43.07	
	$199_2 \rightarrow 294_{11} \rightarrow 271_6 \rightarrow (\text{cont.})$			93.05,
	$134_9 \rightarrow 137_9 \rightarrow 12_6 \rightarrow 23_1 \rightarrow 0$	44/44	23.17,36.88	152.85
	$79_7 \rightarrow 41_5 \rightarrow 38_4 \rightarrow 52_6 \rightarrow (\text{cont.})$			
	$29_{12} \rightarrow 59_2 \rightarrow 73_5 \rightarrow 234_8 \rightarrow 0$	49/53	28.10,43.41	
	$204_9 \rightarrow 182_{11} \rightarrow 111_5 \rightarrow 0$	25/36	15.17,29.49	
Pembroke	$88_{21} \rightarrow 84_{23} \rightarrow 22_9 \rightarrow 0$	53/53	19.28,36.34	
	$44_{33} \rightarrow 17_9 \rightarrow 0$	42/44	13.10,25.59	
	$132_{15} \rightarrow 15_{17} \rightarrow 0$	32/36	9.90,19.16	
	$66_{27} \rightarrow 68_{15} \rightarrow 22_1 \rightarrow 13_6 \rightarrow 0$	49/53	14.38,24.63	104.55,
	$50_{15} \rightarrow 193_{14} \rightarrow 256_{15} \rightarrow (\text{cont.})$			191.12
	$206_6 \rightarrow 207_3 \rightarrow 0$	53/53	20.20,37.68	
	$22_{47} \rightarrow 293_6 \rightarrow 0$	53/53	10.55,18.27	
	$62_{23} \rightarrow 246_8 \rightarrow 181_{21} \rightarrow 13_1 \rightarrow 0$	53/53	17.13,29.44	

Canberra <sup>a</sup>	$327_{11} \rightarrow 85_4 \rightarrow 134_{13} \rightarrow 0$	28/28	28.07,40.29	
	$214_{48} \rightarrow 179_{23} \rightarrow 189_4 \rightarrow 0$	75/78	17.63,23.11	
	$329_{17} \rightarrow 328_{27} \rightarrow 321_3 \rightarrow 330_1 \rightarrow 0$	48/49	29.95,42.45	
	$268_{73} \rightarrow 285_2 \rightarrow 0$	75/78	18.32,24.07	188.82,
	$122_{47} \rightarrow 120_{15} \rightarrow 0$	62/65	16.77,22.81	258.10
	$35_9 \rightarrow 27_{18} \rightarrow 23_9 \rightarrow 14_5 \rightarrow 68_{38} \rightarrow 0$	79/80	27.47,36.73	
	$155_7 \rightarrow 169_{23} \rightarrow 206_{28} \rightarrow 0$	58/59	24.37,32.95	
	$246_{14} \rightarrow 227_{27} \rightarrow 66_{33} \rightarrow 0$	74/78	26.25,35.68	
Handaq	$160_{16} \rightarrow 20_{23} \rightarrow 53_{12} \rightarrow 0$	51/53	12.10,20.45	
	$375_3 \rightarrow 206_{13} \rightarrow 88_{21} \rightarrow 77_{10} \rightarrow 0$	47/53	16.33,28.40	
	$275_{18} \rightarrow 36_{16} \rightarrow 41_1 \rightarrow 0$	35/36	13.68,25.24	97.15,
	$66_{11} \rightarrow 58_{10} \rightarrow 69_{16} \rightarrow 41_9 \rightarrow 0$	46/53	14.28,24.49	170.71
	$2_{11} \rightarrow 6_9 \rightarrow 338_{14} \rightarrow 332_{19} \rightarrow 0$	53/53	21.05,37.13	
	$291_9 \rightarrow 210_{16} \rightarrow 86_{13} \rightarrow 113_{15} \rightarrow 0$	53/53	19.70,35.00	
Valletta	$122_{15} \rightarrow 319_{17} \rightarrow 334_9 \rightarrow (\text{cont.})$			
	$401_7 \rightarrow 259_1 \rightarrow 0$	49/53	22.83,39.67	
	$283_{12} \rightarrow 0$	12/14	5.72,13.59	
	$427_{19} \rightarrow 151_3 \rightarrow 86_{31} \rightarrow 0$	53/53	19.95,39.80	106.03,
	$268_{10} \rightarrow 329_{23} \rightarrow 401_{14} \rightarrow 259_1 \rightarrow 0$	48/53	19.08,33.64	204.34
	$46_{13} \rightarrow 213_9 \rightarrow 58_{17} \rightarrow 61_{14} \rightarrow 0$	53/53	20.35,39.76	
Birkirkara	$28_{17} \rightarrow 32_{36} \rightarrow 0$	53/53	18.10,37.87	
	$371_{17} \rightarrow 370_{19} \rightarrow 185_{10} \rightarrow 0$	46/53	11.93,21.36	
	$397_{14} \rightarrow 293_{18} \rightarrow 358_{14} \rightarrow 447_7 \rightarrow 0$	53/53	16.62,28.04	
	$466_3 \rightarrow 251_5 \rightarrow 316_{12} \rightarrow (\text{cont.})$			
	$209_{14} \rightarrow 5_{12} \rightarrow 170_6 \rightarrow 0$	52/53	19.15,31.73	92.65,
	$434_{29} \rightarrow 104_{10} \rightarrow 419_9 \rightarrow 170_5 \rightarrow 0$	53/53	18.80,32.92	161.08
	$152_{25} \rightarrow 154_{15} \rightarrow 335_{10} \rightarrow 0$	50/53	12.78,22.35	
Hamrun	$393_{13} \rightarrow 120_{21} \rightarrow 235_{18} \rightarrow 0$	52/53	13.37,24.66	
	$501_{15} \rightarrow 437_8 \rightarrow 43_6 \rightarrow 47_4 \rightarrow 19_{20} \rightarrow 0$	53/53	15.32,25.69	
	$243_{17} \rightarrow 244_3 \rightarrow 74_4 \rightarrow 76_1 \rightarrow 0$	25/36	10.28,18.87	
	$451_{25} \rightarrow 3_{13} \rightarrow 6_{13} \rightarrow 0$	51/53	14.03,25.03	98.98,
	$262_{16} \rightarrow 386_{13} \rightarrow 385_{23} \rightarrow 0$	52/53	16.13,29.27	175.31
	$123_{23} \rightarrow 210_{16} \rightarrow 12_{10} \rightarrow 0$	49/53	12.02,20.56	
	$497_{20} \rightarrow 164_{13} \rightarrow 365_{17} \rightarrow 0$	50/53	15.08,28.07	
Cardiff	$509_6 \rightarrow 239_2 \rightarrow 95_{24} \rightarrow 74_1 \rightarrow 76_8 \rightarrow 0$	41/44	16.12,27.82	
	$281_{23} \rightarrow 116_{12} \rightarrow 0$	35/37	19.53,41.62	
	$456_{47} \rightarrow 0$	47/49	14.70,33.27	67.92,
	$534_{45} \rightarrow 362_2 \rightarrow 0$	47/49	18.15,37.14	150.53
M. Keynes	$442_{27} \rightarrow 0$	27/27	15.53,38.51	
	$293_{43} \rightarrow 373_{37} \rightarrow 0$	80/80	15.20,27.61	
	$182_{40} \rightarrow 253_{27} \rightarrow 0$	67/70	14.95,28.94	59.03,
	$165_{36} \rightarrow 365_{14} \rightarrow 0$	50/51	11.38,22.39	111.88
	$199_{25} \rightarrow 354_{52} \rightarrow 0$	77/78	17.50,32.94	

	$57_{80} \rightarrow 0$	80/80	23.73,33.02	
	$57_{80} \rightarrow 0$	80/80	23.73,33.02	
	$178_{23} \rightarrow 385_{29} \rightarrow 92_7 \rightarrow 0$	59/61	31.37,43.50	
Bridgend <sup>a</sup>	$567_{72} \rightarrow 0$	72/74	24.18,34.09	177.72,
	$363_8 \rightarrow 393_4 \rightarrow 225_1 \rightarrow (\text{cont.})$			247.78
	$74_8 \rightarrow 130_{19} \rightarrow 0$	40/43	32.75,44.16	
	$247_2 \rightarrow 432_{21} \rightarrow 0$	23/23	21.05,30.88	
	$507_{18} \rightarrow 225_3 \rightarrow 389_6 \rightarrow 0$	27/27	20.90,29.12	
	$142_{80} \rightarrow 0$	80/80	13.40,24.83	
Edinburgh-2	$142_{48} \rightarrow 301_{32} \rightarrow 0$	80/80	17.75,33.29	60.10,
	$215_{75} \rightarrow 301_5 \rightarrow 0$	80/80	19.03,36.77	110.10
	$301_{80} \rightarrow 0$	80/80	9.92,15.20	
	$70_{73} \rightarrow 910_3 \rightarrow 0$	76/78	18.08,33.57	
	$823_{67} \rightarrow 0$	67/70	10.62,19.05	
	$92_{73} \rightarrow 0$	73/74	14.67,29.36	
Edinburgh-1	$581_{10} \rightarrow 129_{19} \rightarrow 725_{51} \rightarrow 0$	80/80	16.42,27.64	140.55,
	$88_{47} \rightarrow 635_{20} \rightarrow 0$	67/70	10.95,17.85	256.56
	$706_{46} \rightarrow 297_{30} \rightarrow 910_4 \rightarrow 0$	80/80	23.25,41.86	
	$92_{80} \rightarrow 0$	80/80	15.25,29.94	
	$337_{56} \rightarrow 791_{21} \rightarrow 0$	77/78	17.47,32.21	
	$99_{46} \rightarrow 910_{34} \rightarrow 0$	80/80	13.85,25.09	
	$1172_{80} \rightarrow 0$	80/80	17.08,35.00	
	$1152_{70} \rightarrow 0$	70/70	9.75,16.21	
	$675_{38} \rightarrow 1086_{37} \rightarrow 0$	75/78	19.40,38.53	
Adelaide	$146_{62} \rightarrow 1152_{16} \rightarrow 0$	78/78	16.48,29.75	127.20,
	$1094_{29} \rightarrow 1050_{20} \rightarrow 508_{29} \rightarrow 0$	78/78	18.17,33.08	242.59
	$604_{27} \rightarrow 786_{43} \rightarrow 106_{10} \rightarrow 0$	80/80	22.45,40.23	
	$633_{31} \rightarrow 0$	31/33	13.87,33.32	
	$1152_{73} \rightarrow 0$	73/78	10.00,16.46	
	$915_{36} \rightarrow 1657_{44} \rightarrow 0$	80/80	18.97,24.99	
	$811_{53} \rightarrow 425_{26} \rightarrow 0$	79/80	20.23,26.90	
	$1529_{37} \rightarrow 252_{41} \rightarrow 0$	78/78	16.67,21.86	
	$434_{80} \rightarrow 0$	80/80	19.57,26.55	
Brisbane <sup>a</sup>	$726_{75} \rightarrow 0$	75/78	19.12,26.08	209.12,
	$647_{34} \rightarrow 1462_4 \rightarrow 1243_{22} \rightarrow 1249_4 \rightarrow 0$	64/65	23.93,31.37	279.73
	$1373_{74} \rightarrow 0$	74/78	24.70,34.79	
	$1025_{29} \rightarrow 1035_4 \rightarrow 1592_7 \rightarrow 1507_{27} \rightarrow 0$	67/70	19.43,25.30	
	$1395_{28} \rightarrow 1764_{24} \rightarrow 602_{27} \rightarrow 811_1 \rightarrow 0$	80/80	26.00,34.60	
	$1088_{68} \rightarrow 1035_{12} \rightarrow 0$	80/80	20.50,27.28	

<sup>a</sup>For this problem instance, the reliability level  $\kappa = 0.95$ .

### 5.6.2 Other Configurations

The results from the other configurations of  $\alpha$  and  $\beta$  are presented in Tables 5.5 to 5.12. As we might expect, we see from these tables that the higher the CoV parameter  $\beta$  ( $\alpha$  kept fixed), the higher the likelihood of experiencing infeasibility when analysing the deterministic solutions under SLN and independent travel times and when attempting to generate routes for the independent scenario. The same observation holds vice-versa; i.e., when  $\beta$  is fixed and  $\alpha$  is varied.

For a fixed  $\alpha$ , the higher  $\beta$  is, the larger the dispersion in the routes' journey times and the higher the routes' percentile journey times. For most cases, we found that keeping  $\alpha$  fixed and increasing  $\beta$  did not alter the number  $k$  of routes generated for the independent scenario or altered it by one. However, there are six cases which stand out. For the Victoria instance and  $\alpha \in \{0.2, 0.3\}$ ,  $k$  increased from 4 for a low variability level ( $\beta = 0.25$ ) to 6 for a high variability level ( $\beta = 0.75$ ). Similarly, for the Suffolk instance and  $\alpha = 0.1$ , the average  $k$  increased from 4.76 to 7. The other three cases are the following: (i) Canberra with  $\alpha = 0.1$ , for which the average  $k$  increased from 8 to 12.04 (the case  $\beta = 0.75$  is infeasible), (ii) Adelaide with  $\alpha = 0.3$ , for which the average  $k$  increased from 8 to 9.72, and (iii) Brisbane with  $\alpha = 0.1$ , for which the average  $k$  increased from 10 to 15.80.

For fixed  $\beta = 0.75$ , we observed similar overall trends as to those described

Table 5.5: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.2, \beta = 0.25$ .

Location	Deterministic Scenario				Independent Scenario		
	$k$	TMT	IR	TPT	$k$	TMT	TPT
Mgarr	4	54.10 $\pm$ 0.00	0	76.81 $\pm$ 0.25	4	54.51 $\pm$ 0.07	75.21 $\pm$ 0.05
Mellieħa	4	56.30 $\pm$ 0.00	0	78.32 $\pm$ 0.00	4	56.85 $\pm$ 0.20	77.06 $\pm$ 0.32
Porthcawl	1	27.83 $\pm$ 0.43	0	40.60 $\pm$ 1.10	1	27.40 $\pm$ 0.53	39.02 $\pm$ 0.86
Qrendi	5	75.73 $\pm$ 0.00	0	104.35 $\pm$ 0.02	5	75.97 $\pm$ 0.72	104.35 $\pm$ 0.18
Suffolk	3	116.61 $\pm$ 0.90	25	-	4	116.44 $\pm$ 0.65	156.44 $\pm$ 0.93
Senglea	6	72.01 $\pm$ 0.54	0	108.53 $\pm$ 2.12	6	74.22 $\pm$ 1.05	105.73 $\pm$ 1.70
Victoria	4	92.96 $\pm$ 0.41	0	120.82 $\pm$ 0.90	4	93.52 $\pm$ 0.81	119.72 $\pm$ 1.19
Pembroke	7	103.49 $\pm$ 0.68	0	153.50 $\pm$ 1.35	7	105.98 $\pm$ 1.25	148.86 $\pm$ 2.37
Canberra	7	179.11 $\pm$ 0.53	25	-	8	193.20 $\pm$ 1.04	288.35 $\pm$ 1.83
Handaq	6	96.61 $\pm$ 0.57	0	134.29 $\pm$ 1.52	6	97.60 $\pm$ 0.62	131.83 $\pm$ 1.47
Valletta	6	104.50 $\pm$ 0.78	0	160.01 $\pm$ 2.99	6	107.87 $\pm$ 1.54	155.98 $\pm$ 2.96
Birkirkara	6	97.49 $\pm$ 2.64	0	136.92 $\pm$ 6.33	6	98.40 $\pm$ 2.07	134.01 $\pm$ 3.44
Hamrun	7	99.31 $\pm$ 1.40	0	141.49 $\pm$ 2.36	7	100.72 $\pm$ 1.38	137.99 $\pm$ 2.32
Cardiff	2	67.12 $\pm$ 0.00	25	-	3	57.53 $\pm$ 0.00	90.34 $\pm$ 0.00
M. Keynes	4	59.88 $\pm$ 0.57	0	89.80 $\pm$ 1.09	4	59.57 $\pm$ 0.68	87.67 $\pm$ 1.13
Bridgend	6	170.55 $\pm$ 0.53	25	-	7	179.73 $\pm$ 1.29	275.54 $\pm$ 1.28
Edinburgh-2	4	59.13 $\pm$ 0.00	0	87.04 $\pm$ 0.00	4	60.10 $\pm$ 0.00	86.56 $\pm$ 0.00
Edinburgh-1	9	143.71 $\pm$ 1.56	0	210.92 $\pm$ 3.55	9	146.96 $\pm$ 2.34	209.59 $\pm$ 3.63
Adelaide	8	128.18 $\pm$ 1.45	0	190.38 $\pm$ 2.49	8	129.26 $\pm$ 1.09	186.68 $\pm$ 1.46
Brisbane	10	213.41 $\pm$ 4.31	1	-	10	220.87 $\pm$ 4.09	320.22 $\pm$ 6.45

Table 5.6: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.2, \beta = 0.75$ .

Location	Deterministic Scenario			Independent Scenario			
	$k$	TMT	IR	TPT	$k$	TMT	TPT
Mġarr	4	54.10 ± 0.00	0	116.47 ± 0.33	4.00	54.91 ± 0.16	113.39 ± 0.21
Mellieħa	4	56.30 ± 0.00	0	120.60 ± 0.00	4.00	57.29 ± 0.26	117.02 ± 0.12
Porthcawl	1	27.83 ± 0.43	25	-	2.00	24.28 ± 0.00	60.65 ± 0.00
Qrendi	5	75.73 ± 0.00	0	159.01 ± 0.05	5.00	76.82 ± 1.35	159.01 ± 0.31
Suffolk	3	116.61 ± 0.90	25	-	-	-	-
Senglea	6	72.01 ± 0.54	0	162.39 ± 3.15	6.00	74.51 ± 1.10	158.14 ± 3.16
Victoria	4	92.96 ± 0.41	25	-	6.00	96.32 ± 0.68	203.79 ± 1.81
Pembroke	7	103.49 ± 0.68	18	-	7.00	105.77 ± 0.83	226.00 ± 3.25
Canberra	7	179.11 ± 0.53	25	-	-	-	-
Ħandaq	6	96.61 ± 0.57	11	-	6.00	97.89 ± 0.85	203.47 ± 2.57
Valetta	6	104.50 ± 0.78	25	-	7.00	119.66 ± 2.66	272.40 ± 10.64
Birkirkara	6	97.49 ± 2.64	21	-	6.00	96.33 ± 2.17	197.27 ± 5.16
Ħamrun	7	99.31 ± 1.40	3	-	7.00	100.78 ± 1.07	210.23 ± 3.54
Cardiff	2	67.12 ± 0.00	25	-	-	-	-
M. Keynes	4	59.88 ± 0.57	0	131.44 ± 1.74	4.00	59.79 ± 0.77	129.68 ± 2.13
Bridgend	6	170.55 ± 0.53	25	-	-	-	-
Edinburgh-2	4	59.13 ± 0.00	25	-	4.00	60.09 ± 0.08	123.98 ± 0.51
Edinburgh-1	9	143.71 ± 1.56	21	-	9.00	145.25 ± 1.72	302.77 ± 5.26
Adelaide	8	128.18 ± 1.45	25	-	8.64	133.86 ± 2.29	289.13 ± 7.01
Brisbane	10	213.41 ± 4.31	25	-	-	-	-

Table 5.7: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.1, \beta = 0.25$ .

Location	Deterministic Scenario			Independent Scenario			
	$k$	TMT	IR	TPT	$k$	TMT	TPT
Mġarr	4	54.10 ± 0.00	0	77.29 ± 0.17	4.00	54.51 ± 0.05	76.22 ± 0.06
Mellieħa	4	56.30 ± 0.00	0	79.61 ± 0.00	4.00	56.74 ± 0.17	78.48 ± 0.18
Porthcawl	1	27.83 ± 0.43	0	41.23 ± 1.04	1.00	27.38 ± 0.54	39.84 ± 1.06
Qrendi	5	75.73 ± 0.00	0	105.91 ± 0.02	5.00	75.87 ± 0.46	105.91 ± 0.15
Suffolk	3	116.61 ± 0.90	25	-	4.76	115.88 ± 0.99	164.86 ± 1.06
Senglea	6	72.01 ± 0.54	0	106.88 ± 1.74	6.00	73.78 ± 1.12	106.05 ± 1.12
Victoria	4	92.96 ± 0.41	0	124.56 ± 0.91	4.00	93.31 ± 0.39	123.00 ± 0.49
Pembroke	7	103.49 ± 0.68	0	153.22 ± 1.30	7.00	105.19 ± 1.11	149.88 ± 1.69
Canberra	7	179.11 ± 0.53	25	-	8.00	193.43 ± 0.89	290.10 ± 1.69
Ħandaq	6	96.61 ± 0.57	0	136.55 ± 1.37	6.00	97.44 ± 0.63	135.00 ± 1.29
Valetta	6	104.50 ± 0.78	0	159.50 ± 2.43	6.00	106.67 ± 1.37	156.77 ± 1.98
Birkirkara	6	97.49 ± 2.64	0	138.65 ± 5.93	6.00	97.81 ± 1.74	135.91 ± 3.02
Ħamrun	7	99.31 ± 1.40	0	142.73 ± 2.27	7.00	99.98 ± 1.06	139.47 ± 1.81
Cardiff	2	67.12 ± 0.00	25	-	3.00	57.53 ± 0.00	89.36 ± 0.00
M. Keynes	4	59.88 ± 0.57	0	87.95 ± 1.03	4.00	59.65 ± 0.50	87.12 ± 0.63
Bridgend	6	170.55 ± 0.53	25	-	8.00	194.29 ± 1.44	305.59 ± 1.73
Edinburgh-2	4	59.13 ± 0.00	0	84.81 ± 0.00	4.00	59.13 ± 0.00	84.81 ± 0.00
Edinburgh-1	9	143.71 ± 1.56	0	207.32 ± 3.01	9.00	144.93 ± 2.42	206.02 ± 3.44
Adelaide	8	128.18 ± 1.45	0	187.20 ± 2.03	8.00	128.46 ± 1.43	184.71 ± 2.48
Brisbane	10	213.41 ± 4.31	1	-	10.00	218.45 ± 4.53	318.70 ± 7.28

Table 5.8: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.1, \beta = 0.5$ .

Location	Deterministic Scenario				Independent Scenario		
	$k$	TMT	IR	TPT	$k$	TMT	TPT
Mgarr	4	54.10 ± 0.00	0	90.51 ± 0.14	4.00	54.60 ± 0.14	89.58 ± 0.15
Mellieħa	4	56.30 ± 0.00	0	94.45 ± 0.00	4.00	56.68 ± 0.18	93.35 ± 0.16
Porthcawl	1	27.83 ± 0.43	25	-	2.00	24.28 ± 0.00	44.83 ± 0.10
Qrendi	5	75.73 ± 0.00	0	125.21 ± 0.02	5.00	76.01 ± 0.84	125.49 ± 1.07
Suffolk	3	116.61 ± 0.90	25	-	6.00	122.28 ± 0.32	219.69 ± 0.73
Senglea	6	72.01 ± 0.54	0	123.68 ± 1.84	6.00	72.97 ± 0.65	123.47 ± 1.44
Victoria	4	92.96 ± 0.41	1	-	4.00	93.91 ± 1.61	150.71 ± 2.95
Pembroke	7	103.49 ± 0.68	0	179.59 ± 1.58	7.00	104.65 ± 0.87	176.82 ± 1.94
Canberra	7	179.11 ± 0.53	25	-	12.04	239.72 ± 2.78	457.52 ± 5.94
Handaq	6	96.61 ± 0.57	0	162.53 ± 1.58	6.00	97.58 ± 0.66	161.01 ± 1.11
Valletta	6	104.50 ± 0.78	0	188.40 ± 2.56	6.00	106.56 ± 0.99	186.10 ± 1.20
Birkirkara	6	97.49 ± 2.64	0	164.14 ± 6.88	6.00	97.23 ± 1.90	160.58 ± 4.44
Hamrun	7	99.31 ± 1.40	0	168.24 ± 2.75	7.00	100.00 ± 1.24	165.52 ± 2.36
Cardiff	2	67.12 ± 0.00	25	-	3.00	57.53 ± 0.00	105.07 ± 0.00
M. Keynes	4	59.88 ± 0.57	0	100.31 ± 1.22	4.00	59.80 ± 0.64	99.76 ± 1.42
Bridgend	6	170.55 ± 0.53	25	-	-	-	-
Edinburgh-2	4	59.13 ± 0.00	0	95.50 ± 0.00	4.00	59.13 ± 0.00	95.50 ± 0.00
Edinburgh-1	9	143.71 ± 1.56	0	236.74 ± 3.39	9.00	144.74 ± 1.94	236.38 ± 3.92
Adelaide	8	128.18 ± 1.45	0	214.76 ± 2.43	8.00	128.65 ± 1.43	213.34 ± 2.65
Brisbane	10	213.41 ± 4.31	25	-	11.04	222.46 ± 2.98	386.64 ± 6.61

Table 5.9: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.1, \beta = 0.75$ .

Location	Deterministic Scenario				Independent Scenario		
	$k$	TMT	IR	TPT	$k$	TMT	TPT
Mgarr	4	54.10 ± 0.00	0	96.53 ± 0.10	4.00	54.20 ± 0.28	95.90 ± 0.16
Mellieħa	4	56.30 ± 0.00	0	101.57 ± 0.00	4.00	56.58 ± 0.10	100.80 ± 0.11
Porthcawl	1	27.83 ± 0.43	25	-	2.00	24.28 ± 0.00	47.37 ± 0.10
Qrendi	5	75.73 ± 0.00	0	134.59 ± 0.02	5.00	75.76 ± 0.12	134.59 ± 0.10
Suffolk	3	116.61 ± 0.90	25	-	7.00	126.96 ± 0.32	249.48 ± 0.62
Senglea	6	72.01 ± 0.54	0	130.82 ± 1.78	6.00	72.79 ± 0.81	131.11 ± 1.60
Victoria	4	92.96 ± 0.41	25	-	5.00	91.88 ± 0.35	167.05 ± 0.68
Pembroke	7	103.49 ± 0.68	0	191.19 ± 1.69	7.00	104.53 ± 0.91	189.34 ± 2.06
Canberra	7	179.11 ± 0.53	25	-	-	-	-
Handaq	6	96.61 ± 0.57	0	175.27 ± 1.58	6.00	97.36 ± 0.60	174.38 ± 1.71
Valletta	6	104.50 ± 0.78	14	-	6.00	106.59 ± 1.45	199.90 ± 2.87
Birkirkara	6	97.49 ± 2.64	3	-	6.00	95.38 ± 1.45	168.89 ± 3.90
Hamrun	7	99.31 ± 1.40	0	180.05 ± 2.99	7.00	99.61 ± 1.37	178.24 ± 2.60
Cardiff	2	67.12 ± 0.00	25	-	4.00	63.02 ± 0.00	126.33 ± 0.00
M. Keynes	4	59.88 ± 0.57	0	105.23 ± 1.31	4.00	59.92 ± 0.73	105.00 ± 1.54
Bridgend	6	170.55 ± 0.53	25	-	-	-	-
Edinburgh-2	4	59.13 ± 0.00	0	99.64 ± 0.00	4.00	59.13 ± 0.00	99.64 ± 0.00
Edinburgh-1	9	143.71 ± 1.56	1	-	9.00	143.93 ± 2.10	248.27 ± 4.40
Adelaide	8	128.18 ± 1.45	0	226.23 ± 2.72	8.00	128.27 ± 0.98	224.48 ± 1.87
Brisbane	10	213.41 ± 4.31	25	-	15.80	264.60 ± 15.25	513.80 ± 38.05

Table 5.10: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.3, \beta = 0.25$ .

Location	Deterministic Scenario				Independent Scenario		
	$k$	TMT	IR	TPT	$k$	TMT	TPT
Mġarr	4	54.10 ± 0.00	0	74.86 ± 0.24	4	54.50 ± 0.00	73.35 ± 0.00
Mellieħa	4	56.30 ± 0.00	0	76.30 ± 0.00	4	56.85 ± 0.24	75.25 ± 0.30
Porthcawl	1	27.83 ± 0.43	0	39.44 ± 1.05	1	27.54 ± 0.51	38.23 ± 0.92
Qrendi	5	75.73 ± 0.00	0	101.78 ± 0.02	5	76.03 ± 0.81	101.80 ± 0.15
Suffolk	3	116.61 ± 0.90	25	-	4	115.69 ± 0.05	151.38 ± 0.03
Senglea	6	72.01 ± 0.54	0	106.17 ± 2.10	6	74.22 ± 0.91	103.40 ± 1.48
Victoria	4	92.96 ± 0.41	0	118.21 ± 0.85	4	93.68 ± 0.99	117.48 ± 1.31
Pembroke	7	103.49 ± 0.68	0	149.56 ± 1.30	7	105.69 ± 1.16	144.82 ± 2.07
Canberra	7	179.11 ± 0.53	25	-	8	193.39 ± 1.10	279.24 ± 1.43
Ħandaq	6	96.61 ± 0.57	0	130.88 ± 1.46	6	97.55 ± 0.61	128.34 ± 1.04
Valetta	6	104.50 ± 0.78	0	155.71 ± 2.96	6	106.90 ± 1.09	150.96 ± 1.45
Birkirkara	6	97.49 ± 2.64	0	133.43 ± 6.08	6	97.93 ± 2.39	129.92 ± 3.75
Ħamrun	7	99.31 ± 1.40	0	137.84 ± 2.28	7	100.82 ± 1.05	134.95 ± 1.82
Cardiff	2	67.12 ± 0.00	25	-	3	57.53 ± 0.00	87.98 ± 0.00
M. Keynes	4	59.88 ± 0.57	0	87.89 ± 1.06	4	59.34 ± 0.57	85.22 ± 0.85
Bridgend	6	170.55 ± 0.53	25	-	7	180.49 ± 1.04	270.06 ± 1.90
Edinburgh-2	4	59.13 ± 0.00	0	85.44 ± 0.00	4	60.10 ± 0.00	84.83 ± 0.00
Edinburgh-1	9	143.71 ± 1.56	0	206.64 ± 3.56	9	146.59 ± 3.11	203.25 ± 3.19
Adelaide	8	128.18 ± 1.45	0	186.37 ± 2.51	8	129.24 ± 1.28	182.59 ± 1.59
Brisbane	10	213.41 ± 4.31	0	316.10 ± 7.12	10	220.13 ± 5.37	312.73 ± 5.84

Table 5.11: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.3, \beta = 0.5$ .

Location	Deterministic Scenario				Independent Scenario		
	$k$	TMT	IR	TPT	$k$	TMT	TPT
Mġarr	4	54.10 ± 0.00	0	101.20 ± 0.46	4	55.11 ± 0.49	97.03 ± 0.31
Mellieħa	4	56.30 ± 0.00	0	102.53 ± 0.01	4	57.20 ± 0.32	98.61 ± 0.14
Porthcawl	1	27.83 ± 0.43	25	-	2	224.33 ± 0.12	54.69 ± 0.45
Qrendi	5	75.73 ± 0.00	0	135.68 ± 0.05	5	78.63 ± 0.88	134.47 ± 0.95
Suffolk	3	116.61 ± 0.90	25	-	-	-	-
Senglea	6	72.01 ± 0.54	0	145.83 ± 3.51	6	75.22 ± 1.18	138.82 ± 2.68
Victoria	4	92.96 ± 0.41	1	-	4	94.70 ± 1.86	149.62 ± 2.89
Pembroke	7	103.49 ± 0.68	0	206.21 ± 2.12	7	107.40 ± 1.91	195.38 ± 2.92
Canberra	7	179.11 ± 0.53	25	-	-	-	-
Ħandaq	6	96.61 ± 0.57	0	175.68 ± 2.58	6	98.57 ± 1.16	170.52 ± 3.15
Valetta	6	104.50 ± 0.78	18	-	6	110.82 ± 2.62	205.50 ± 5.19
Birkirkara	6	97.49 ± 2.64	3	-	6	98.38 ± 3.12	169.65 ± 6.74
Ħamrun	7	99.31 ± 1.40	0	187.03 ± 3.67	7	101.56 ± 1.57	178.38 ± 3.03
Cardiff	2	67.12 ± 0.00	25	-	-	-	-
M. Keynes	4	59.88 ± 0.57	0	120.08 ± 1.63	4	59.62 ± 0.99	116.00 ± 1.42
Bridgend	6	170.55 ± 0.53	25	-	-	-	-
Edinburgh-2	4	59.13 ± 0.00	25	-	4	60.09 ± 0.07	113.30 ± 0.46
Edinburgh-1	9	143.71 ± 1.56	5	-	9	148.41 ± 2.54	272.42 ± 5.85
Adelaide	8	128.18 ± 1.45	2	-	8	130.57 ± 2.31	242.24 ± 3.48
Brisbane	10	213.41 ± 4.31	25	-	-	-	-

Table 5.12: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.3, \beta = 0.75$ .

Location	Deterministic Scenario				Independent Scenario			
	$k$	TMT	IR	TPT	$k$	TMT	TPT	
Mġarr	4	54.10 ± 0.00	0	123.67 ± 0.52	4.00	55.07 ± 0.48	118.28 ± 0.07	
Mellieħa	4	56.30 ± 0.00	0	126.22 ± 0.01	4.00	57.71 ± 0.42	120.56 ± 0.27	
Porthcawl	1	27.83 ± 0.43	25	-	2.00	24.31 ± 0.10	66.92 ± 0.44	
Qrendi	5	75.73 ± 0.00	0	166.25 ± 0.07	5.00	78.44 ± 0.56	164.48 ± 1.04	
Suffolk	3	116.61 ± 0.90	25	-	-	-	-	
Senglea	6	72.01 ± 0.54	22	-	6.00	75.85 ± 1.36	167.55 ± 2.39	
Victoria	4	92.96 ± 0.41	25	-	6.00	96.35 ± 0.84	207.72 ± 3.26	
Pembroke	7	103.49 ± 0.68	25	-	-	-	-	
Canberra	7	179.11 ± 0.53	25	-	-	-	-	
Handaq	6	96.61 ± 0.57	14	-	6.00	98.76 ± 1.29	210.77 ± 3.33	
Valetta	6	104.50 ± 0.78	25	-	-	-	-	
Birkirkara	6	97.49 ± 2.64	22	-	-	-	-	
Hamrun	7	99.31 ± 1.40	4	-	7.00	101.21 ± 1.46	216.53 ± 4.02	
Cardiff	2	67.12 ± 0.00	25	-	-	-	-	
M. Keynes	4	59.88 ± 0.57	0	144.10 ± 2.01	4.00	59.24 ± 0.67	138.49 ± 1.20	
Bridgend	6	170.55 ± 0.53	25	-	-	-	-	
Edinburgh-2	4	59.13 ± 0.00	25	-	4.96	54.86 ± 1.11	125.60 ± 1.60	
Edinburgh-1	9	143.71 ± 1.56	25	-	-	-	-	
Adelaide	8	128.18 ± 1.45	25	-	9.72	139.70 ± 2.73	326.48 ± 11.83	
Brisbane	10	213.41 ± 4.31	25	-	-	-	-	

above when increasing  $\alpha$ . Regarding the average  $k$ , an instance that draws attention is Adelaide, which saw an increase from 8 ( $\alpha = 0.1$ ) to 9.72 ( $\alpha = 0.3$ ) routes. For a fixed medium variability level ( $\beta = 0.5$ ), the routes' percentile journey times increased with an increase in  $\alpha$  for most instances. Nonetheless, there are some instances for which an increase in  $\alpha$  from 0.2 to 0.3 yielded a lower average TPT. For example, the average TPTs for  $\alpha \in \{0.1, 0.2, 0.3\}$  when analysing Victoria's deterministic solutions under SLN and independent travel times are 150.71, 155.13, and 149.62 minutes, respectively. This occurs because, as  $\alpha$  increases, the rate of decrease in a route's  $\gamma_R$  and  $\sigma_R^2$  parameters exceeds the rate of increase in  $\mu_R$  (refer to Equations (5.10)-(5.12)). The same pattern is visible across all the instances when  $\beta$  is fixed to 0.25 and  $\alpha$  increases from 0.2 to 0.3. It also emerges for most instances for fixed  $\beta = 0.25$  and an increase in  $\alpha$  from 0.1 to 0.2.

## 5.7 Summary

The quality of public transportation services, such as school bus transportation, plays a central role in influencing citizens' decisions to choose public over private transportation. In a world with limited land resources, increasing traffic demand,

and environmental challenges, quantifying, modelling, and optimizing TTR is crucial when evaluating transportation systems. For a long time, surveys of transportation system users have indicated that TTR is one of the key service quality indicators. For instance, Bates et al. (2001) observed excessive disutility due to late arrivals and concluded that punctuality is highly valued by users.

In this chapter, we have incorporated stochastic travel times modelled by SLN distributions for a more realistic assessment of vehicle arrivals at the planning stage. Our primary goal was not to fit travel time distributions to empirical data but to develop a framework for designing a priori routing strategies insensitive to travel time uncertainties. Deterministically generated solutions, such as those based on average travel times, may seem promising from a planning point of view; however, their effectiveness can diminish in practice due to disruptions caused by stochastic travel times. For instance, unforeseen disturbances may require the addition of extra routes. Using stochastic travel times allows for a comprehensive set of metrics to evaluate TTR and to identify areas that should be avoided due to potential adverse traffic conditions. In this way, unanticipated consequences are reduced when implementing plans in the real world.

As a TTR metric, we have opted for a percentile travel time. This metric conveys travel time variability in a way that allows users to better relate to their past travel experiences, understand downside risks, and reflect attitudes towards delays. Being a “logical measure of ‘being late’” (Taylor, 2017b), this metric also enables operators to control the acceptance level of delays, especially in systems with strict punctuality requirements.

Our computational results indicate that a stochastic modelling framework offers a significant improvement over a deterministic one. It is preferable to have more predictability in the routing strategies rather than relying on faster but less reliable routes. In this chapter, we have assumed that travel times on different arcs are independent. While this assumption simplifies the proposed framework, it may not always reflect real-world conditions where travel times can be correlated. Our framework is flexible enough to handle not only different TTR requirements, as needed for other applications, but also correlations between travel times. This extension will be the subject of the next chapter.

# Chapter 6

## Modelling Travel Time Correlations

### 6.1 Introduction

Due to the strong interconnections between the various constituents of a road network, traffic conditions on one road section can be significantly affected by the traffic conditions on both adjacent and nearby non-adjacent road sections (Bakach et al., 2021). For example, congestion on roads *a* and *b* in Figure 6.1 can propagate to road *c*, and subsequently impact roads *d* and *e*, highlighting the cascading effects of localized traffic build-up within a road network. Investigating these inherent interactions between road sections is essential for accurately predicting traffic congestion (Guo et al., 2019a). Understanding such dependencies can help to reveal how traffic flow propagates in the network and quantify the impact of specific road sections on others.

This chapter investigates whether incorporating travel time dependencies into

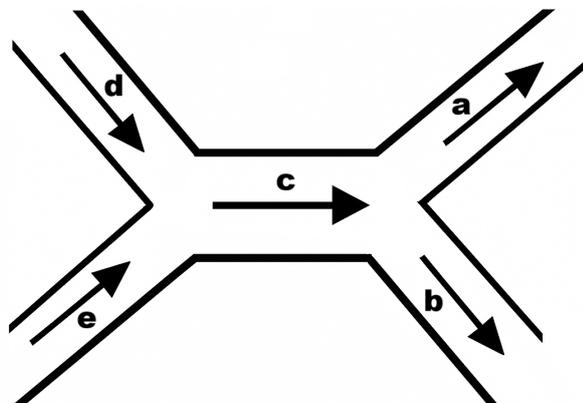


Figure 6.1: Road network diagram illustrating the propagation of traffic conditions – congestion on roads *a* and *b* can impact roads *c*, *d*, and *e*.

the framework proposed in Chapter 5 yields routing strategies with improved reliability. Importantly, the objective function and constraints of the CCP framework in Chapter 5 remain exactly the same. The extension studied in this chapter is the introduction of non-zero correlations  $\rho_{e_1 e_2}$  between the travel times on selected arc pairs  $e_1, e_2 \in E_1$  (see Equation (5.11)).

In this chapter, the terms ‘arc’ and ‘link’ will be used interchangeably and refer to a physical path connecting two nodes in a network. We will also be distinguishing between two types of correlations that can exist between arc travel times. These are *spatial correlation* and *temporal correlation*. The former refers to the correlation between the travel times on two different arcs in a single time period. Meanwhile, the latter refers to the correlation between the travel times on a single arc in two different time periods. In addition, by *spatial-temporal correlation*, we mean the correlation between the travel times on two different arcs in two different time periods. Since our routes are always performed in the morning and during a narrow time interval (45 minutes), our research will focus solely on spatial correlations. Some works incorporating temporal correlations alongside spatial correlations are referred to in Section 6.2 for completeness.

Although it is acknowledged in the literature that independence between links does not exist in practice, several researchers still assume it for the sake of simplicity and to facilitate tractability (Adulyasak and Jaillet, 2016, Binart et al., 2016, Chen et al., 2014a, Chen and Fan, 2020, Chu et al., 2017, Ehmke et al., 2015, Kenyon and Morton, 2003, Li et al., 2010a, Miranda and Conceição, 2016, Nie and Wu, 2009b, Russell and Urban, 2008, Taş et al., 2013, Tavakkoli-Moghaddam et al., 2012, Taylor, 2017a, Taylor and Susilawati, 2012, Thompson et al., 2011, Vareias et al., 2019, Wang and Lin, 2017, Wu and Geistefeldt, 2015, Yan et al., 2014, Zhang et al., 2013). Sen et al. (1997) state that travel times on consecutive links are correlated since, during a specific time interval, many of the vehicles on them are the same ones. They emphasize that while ignoring such correlations is convenient, it yields a biased estimate of the travel time variation. Cheng et al. (2011) also accentuate that oversimplifying the relationships between links can lead to poor forecasts, especially in non-recurrent traffic conditions. Bakach et al. (2021) claim that it is crucial to understand how correlations impact VRP solutions since the distribution of a route’s travel time is affected by correlations.

A travel time correlation indicates how one link’s travel time varies with another link’s travel time. A positive correlation signifies that when the travel time on one link increases, so does that on another. For instance, a traffic accident on a major urban road may cause significant delays on other surrounding roads. A lane drop may also cause propagation of traffic congestion in the upstream links. Furthermore, heavy snowfall, rain, or fog is likely to affect all the links in

a local area simultaneously and increase total travel times (for example, by more than 45% in the work of Jiang and Mahmassani (2014)). A negative correlation represents that the longer the travel time on one link, the shorter the travel time on another. Such a correlation can occur when a bottleneck in one link restricts the flow in the downstream links and leads to increased speeds on the downstream links (Nicholson, 2015). Parallel links can also be negatively correlated in peak periods, as noted by Ermagun et al. (2017).

Correlations between travel times imply that the travel time variance of a route is non-additive in the variances of the individual links making up the route. As stated by Laporte et al. (1992), the travel time variance of a route  $R$  is given by

$$\sum_{e \in R} \text{Var}[T_e] + \sum_{\substack{e_1, e_2 \in R \\ e_1 \neq e_2}} \rho_{e_1 e_2} \sqrt{\text{Var}[T_{e_1}] \text{Var}[T_{e_2}]},$$

where  $\text{Var}[T_e]$  denotes the variance of the travel time on arc  $e \in E_1$ . Thus, positive (negative) correlations cause a route's travel time variance to be higher (lower) than the sum of the travel time variances of the individual links. If a network has a mixture of positive and negative correlations, then some of them will cancel out the effect of the others to some extent (Nicholson, 2015). A route with a higher travel time variance has a more dispersed probability distribution and, consequently, the probability that it is completed within a specified time interval is lower (Sarmiento et al., 2015). The same holds vice-versa. Therefore, route evaluators that assume independence can overestimate the on-time arrival probability for positive correlations and underestimate it for negative correlations. This further emphasizes the importance of including correlations to avoid compromising solution quality.

This chapter is organized as follows. Section 6.2 reviews key VRP studies that have either demonstrated the existence of, or incorporated, travel time correlations. Building on insights from these works, in Section 6.3, we propose a strategy for modelling spatial correlations within the framework introduced in Chapter 5. In Section 6.4, we revisit the twenty instances summarized in Tables 3.1 and 3.2, allowing for correlations between selected pairs of arcs. As shown in this chapter, modelling travel time correlations is computationally intensive, particularly in terms of memory requirements due to the large number of arc pairs. Section 6.4 also highlights improvements in the computational results compared to those presented in Chapters 4 and 5. Finally, Section 6.5 presents a Monte Carlo simulation approach for generating SLN travel times based on the correlation structure proposed in Section 6.3, and applies it to the solutions from Chapters 4 to 6.

## 6.2 Related Works

This section will review works of two types. Section 6.2.1 covers studies that have proven the existence of travel time correlations using historical real-world traffic data. Then, Section 6.2.2 discusses studies that have modelled these correlations and demonstrated the positive impact of incorporating them in routing decisions. Note that while comprehensive, this review does not claim to be exhaustive.

### 6.2.1 Existence of Travel Time Correlations

For decades, the literature has consistently established that traffic links are correlated. Nearly three decades ago, a study by MVA Consultancy (1996) identified both positive and negative travel time correlations. This study involved two datasets: one for a congested urban network and another for three groups of links (motorway, inter-urban, and urban). Correlations between 0.2 and 0.4 were generally found for the first dataset, with the highest correlations observed for spatially close links. For the second dataset, correlations between adjacent links were typically positive, reaching up to 0.4, while correlations between non-adjacent links were inconsistent but mostly positive. This study confirmed a generally positive correlation structure, with the travel time variances of routes almost always being higher than the sum of the individual links' travel time variances.

Munakata (2007) analysed data collected in one-minute intervals over 93 days for 39 links on the Tokyo Metropolitan Expressway Central Circular Route, calculating correlations between all pairs of links. For weekdays, the 06:00–06:01 interval – selected as an illustrative example of conditions during the morning peak period – exhibited correlations ranging from  $-0.094$  to  $0.987$ , with an average value of  $0.362$ . Munakata (2007) noted that the higher correlations were generally found between spatially close links, and correlations typically decreased as the distance between links increased. Additionally, the sum of the links' travel time variances accounted for only about 9% of the total variance. Using similar data, Nicholson (2015) revealed strong and statistically significant correlations, even between non-adjacent links. They emphasized that while correlations complicate route choice modelling, they should not be ignored, even when they are not statistically significant individually. This is because “the large number of correlation-related terms contributing to the total travel time variance means that they are collectively ‘scientifically significant’.”

El Esawey and Sayed (2011) suggested that nearby links are subject to similar traffic conditions, making their travel times likely to be positively correlated. They generated travel time data for 27 links in downtown Vancouver, simulat-

ing five demand levels over two hours. They defined neighbours according to criteria such as a Pearson correlation coefficient greater than 0.5. The strongest correlations were typically observed among adjacent links, with significant correlations also identified for parallel links, particularly those on parallel one-way streets. Furthermore, they noted low correlations among spatially distant links, even those on the same street.

Cheng et al. (2011) argued that a link is only directly influenced by its adjacent links (at spatial order one). However, they acknowledged that two links flowing into or out of the same link in the same direction can be indirectly related. For instance, heavy traffic on an arc  $(u_1, v)$  is likely to cause congestion downstream, which can accumulate upstream on an arc  $(u_2, v)$ . These two links are not directly related but are indirectly related via another link in between (links at spatial order two). Cheng et al. (2011) analysed 152,064 journey time observations from the London Congestion Analysis Project network, divided into three time periods (morning peak, interpeak, and evening peak). They examined spatial and temporal changes in the autocorrelation structure, finding that at temporal order zero and spatial order one, cross-correlations ranged from approximately zero to 0.5. This suggests that some adjacent links were highly correlated, while others were not. The same phenomenon was observed for space-time autocorrelations and at spatial order two (with overall lower correlations).

El Esawey and Sayed (2012) attributed the correlation between travel times on nearby links (adjacent, parallel, or intersecting) to three factors: correlation in traffic demand, similarity in traffic control, and queue spillback. They studied link travel times on a 9 km route in downtown Vancouver, collecting data over multiple business days under clear weather conditions, with drivers travelling back and forth during three-hour morning or evening peak periods. Nearby links, referred to as *link neighbours*, were defined by a correlation threshold, initially set at 0.4 and adjusted during sensitivity analysis. The null hypothesis of zero correlation was rejected at the 0.05 level for all pairs of link neighbours. Each link had between one and seventeen neighbours.

To predict travel times along freeway corridors, Zou et al. (2014) estimated spatial-temporal cross-correlations using eight months of data collected along a segment of US-290 in Houston, Texas. They found that travel time changes at upstream links tend to cause changes at downstream links, with a lead time between five and sixty minutes. Additionally, they observed a decreasing cross-correlation trend as the distance between links increased (from a maximum of 0.757 for adjacent links to 0.529 for links 7.1 miles apart).

Ermagun et al. (2017) analysed data across 140 links in the Minneapolis-St. Paul freeway network. Traffic flow data was collected on all Tuesdays in 2015

during three-hour periods (morning rush, non-rush, and evening rush hours). To ensure reliable correlations, they detrended the data by removing the time-of-day trend, then the day-of-week trend, and finally the total system travel by time-of-day trend. They observed both positive and negative correlations in their grid-like network, with strong positive correlations between downstream and upstream links being prevalent regardless of the time or day. These positive correlations diminished with distance, varying between 0.0002 and 0.5413 for the morning rush hour (mean  $0.0732 \pm 0.0821$ ). Strong negative correlations were mostly observed for parallel links and upstream links that were spatially distant, where drivers could choose alternative routes after receiving information about downstream congestion. For the morning rush hour, 39.9% of links were negatively correlated (mean  $-0.0921 \pm 0.0847$ ), though only 2.3% had an absolute value greater than 0.2.

In other work, Guo et al. (2019a) analysed data for 4,530 major road sections in Beijing, China, across three 3-hour periods (morning, early afternoon, and evening). They collected real-time speed records with a resolution of one minute over five weekdays. For each pair of road sections, Guo et al. (2019a) considered their interaction by calculating the cross-correlation between the two time series, with a time lag. They defined the *link weight* as the normalized maximal cross-correlation and the *time delay* as the time lag at which this maximal cross-correlation occurred. Additionally, they shuffled all fifteen hours of the two time series independently and recalculated the link weight and time delay. The distance between two road sections was defined as the smaller of the great circle distance between the endpoints of the road sections. Observing that original link weights were larger within shorter distances than shuffled ones, they considered only those road section pairs with a certain minimal link weight and maximal distance (4.2 and 0.91 km, respectively, for the morning rush hour). Most strong correlations were found to (i) occur between links at spatial orders one, two, and three; (ii) have an absolute time delay of at most two minutes; and (iii) rarely be influenced by the day of the week.

Guo et al. (2020) analysed temporal, spatial, and spatial-temporal correlations of travel speeds in an urban road network with 5,943 links in Chengdu, China (110-day dataset), and a freeway network with 438 links in California (102-day dataset). For spatial correlations, Guo et al. (2020) defined two classes of links: *directly adjacent* and *indirectly adjacent*. Directly adjacent links share a node and have the same traffic-flow direction, while indirectly adjacent links share a node but have different traffic-flow directions or have the same direction but another link in between. For the morning peak period and the urban network, Guo et al. (2020) observed significantly positive (negative) spatial correlations

in 48.95% (0.33%) of directly adjacent pairs and 17.61% (1.05%) of indirectly adjacent pairs. For the freeway network, these proportions were higher at 57.44% (0.65%) and 40.67% (1.27%), respectively.

Most recently, Chen et al. (2024) proposed a Bayesian Gaussian model to estimate correlations between travel times along a bus route. Using weekday in-out-stop record data collected in Guangzhou, China, they found that most link pairs did not exhibit strong correlations (i.e., correlations of at least 0.5). Strong positive correlations were primarily observed between geographically proximate links, while the few negative correlations were weak to moderate. Notably, a small number of strong correlations were also identified between link pairs separated by long distances.

The above review has shown that although several researchers have investigated travel time correlations in transportation networks, their findings are not always consistent due to differences in network types, data sources, and methodological approaches. Nonetheless, the following key points emerge across the literature. Travel time correlations (i) can be both positive and negative (Eliasson, 2007); (ii) are predominantly positive (Figliozzi et al., 2007); (iii) can reach values as high as 0.75 (Park and Rilett, 1999); and (iv) typically decrease with increased spatial distance. As an example of the latter, Zou et al. (2010) found that spatial autocorrelation during the morning peak hour approached zero beyond an average distance of approximately 1.85 km.

## 6.2.2 Modelling and Analysing the Impact of Travel Time Correlations

When data is scarce or unavailable, spatial correlations are often modelled in a limited or simplified manner. For instance, Shahabi et al. (2013) and Prakash and Srinivasan (2018) generated a correlation matrix randomly from the space of all valid correlation matrices (i.e., symmetric positive semi-definite matrices with diagonal entries equal to one and non-diagonal entries ranging from  $-1$  to  $1$ ). Zhang et al. (2017) on the other hand, employed the Toeplitz correlation method to generate correlation matrices with entries lying in  $[-\rho, \rho]$  for a prescribed value of  $\rho \in [0, 1]$ . Furthermore, other studies limit correlations to adjacent or neighbouring arcs (arcs within a certain distance threshold of each other). These simplifications are often necessary due to the computational challenges of specifying all pairwise correlations in a network. Recall from Chapter 3 that a network with  $n$  bus stops yields vertex set  $V_1$  of cardinality  $(n + 1)$  and edge set  $E_1$  of cardinality  $n(n + 1)$ . This implies that such a network contains  $\binom{n(n+1)}{2}$  pairs of distinct arcs. To illustrate the high dimensionality involved, consider the

smallest problem instance in Table 3.1, where the network has  $n = 59$  bus stops, resulting in  $59 \cdot 60 = 3,540$  arcs and, therefore,  $\binom{3540}{2} \approx 6.26$  million pairs of distinct arcs.

If correlations are assumed only between adjacent links, the dependence structure can be modelled as a Markov chain. For example, Fan et al. (2005) assume that each link has two possible states – congested or uncongested (if the travel time is within an a priori bound) – and define conditional distributions to model the transitions between these states, capturing the spatial correlations between adjacent links. Other researchers who have adopted a Markovian framework include Nie and Wu (2009a), Dong et al. (2012), Ramezani and Geroliminis (2012), and Ma et al. (2017). Additionally, sampling-based methods can also be employed to account for correlations between travel speeds or times. For example, Xing and Zhou (2011) directly use morning travel time samples from a freeway corridor in California for path travel time computations, noting that this “considerably complicates the path search process”. Similarly, Yang and Zhou (2017) and Zhang et al. (2021) adopt sampling-based travel times and scenarios sampled from the empirical distribution of travel speeds, respectively.

Correlations can also be handled using copula-based methods, although these methods are challenging to apply to large-scale networks (Zang et al., 2022). A copula-based approach was previously applied by Chen et al. (2017b) on data from two arterials, one in Shanghai and one in Los Angeles. They observed that neglecting correlations led to overestimated path TTR evaluations. Likewise, copula functions were used in the works of Zhao et al. (2014), Chen et al. (2019), Qin et al. (2020), and Zhang et al. (2021). Another way to address correlations between travel times is by accounting for time dependency (Babaei and Rajabi-Bahaabadi, 2019). For the simultaneous bus routing and scheduling problem, their findings indicate that the optimal routes differ, and that scheduled departure times are later when time dependency is incorporated.

The remainder of this section now highlights important findings regarding the impact of travel time correlations on routing solutions. Ji et al. (2011) acknowledged that in realistic transportation networks, “correlation among link travel times may exist not only [between] adjacent links, but also non-adjacent links within close vicinity”. To examine the impact of correlations on path travel time budgets in their multi-objective GA, they simulated correlated normal and log-normal travel times. Ji et al. (2011) assumed that the covariance between links  $e_1$  and  $e_2$  is given by  $0.8\sqrt{\text{Var}[T_{e_1}]\text{Var}[T_{e_2}]}$  under the normal distribution. Their findings demonstrate that assuming independent travel times results in lower path travel time variances and substantial bias in travel time budgets.

For the spatially dependent reliable shortest path problem (SPP), Chen et al.

(2012a) represented spatial travel time correlations using variance-covariance matrices. They assumed that only links within a local impact area, where any two links are at most  $k$  links apart, are correlated. They extended the work of Nie and Wu (2009b) and Nikolova (2009), who assumed correlations only between adjacent links. Through a case study using data collected during the morning peak hour in Hong Kong, they confirmed that strong correlations exist within local impact areas and that increasing  $k$  significantly decreases these correlations. For example, the mean absolute correlation coefficient decreases from 0.29 to 0.04 as  $k$  increases from 1 to 4. This finding supports those of other researchers, such as Gajewski and Rilett (2005). Chen et al. (2012a) also demonstrated that the path travel time standard deviation is underestimated by 26.5% when correlations are ignored, is 87.1% accurate when correlations among adjacent links are assumed, and is 99.1% accurate when  $k = 4$ .

Zockaie et al. (2013) demonstrated that correlations between link travel times significantly impact the optimal solutions of the SPP when considering on-time arrival reliability. Using a large-scale network in Chicago, they developed a Monte Carlo simulation algorithm, assuming that travel times follow a multivariate normal distribution for simplicity. A non-zero correlation was applied only to adjacent links, based on their travel time variances and a constant base level of correlation. They highlighted that the Pareto frontier of travel time budgets against the probability of arriving on time stretches as the correlation level increases. Additionally, they observed that as the origin-destination distance grows, the discrepancy between the Pareto frontiers of the independent and correlated cases becomes more pronounced. The impact of correlations between links at spatial order two was also examined in the work of Zockaie et al. (2016), which extends the study by Zockaie et al. (2013).

Srinivasan et al. (2014) analysed data from 3,644 links in the Chennai network. Since travel times were available for only a few links, the authors could not use maximum likelihood to estimate the variance-covariance matrix. As a result, they assumed correlations only between adjacent links. They randomly sampled fifty origin-destination pairs and, for each pair, identified the most reliable path under SLN travel times, which were either independent, randomly positively correlated (values between 0 and 0.92), or randomly positively and negatively correlated (values between  $-0.59$  and 0.86). Their findings demonstrated that the independence assumption leads to suboptimal solutions in 30% of the cases, an average inaccuracy in the reliability objective of 15%, and reliability compromises of up to 8.5%.

Sarmiento et al. (2015) examined the impact of correlation on the reliability of a route composed of four edges, each following a lognormal distribution. For each

edge, they generated 1,000 realizations and then incorporated a fixed correlation among each pair of edges using the method of Iman and Conover (1982). They observed that as the correlation factor increases, the percentage of total duration realizations that do not exceed forty time units (termed successful realizations) monotonically decreases. Sarmiento et al. (2015) employed the same approach to generate four families of fifty instances each (city, suburban, municipal, and state), with each family having different numbers of nodes and clusters of nodes. They assumed that edges within the same cluster have a correlation of  $\rho$  (either 0 or randomly chosen in the range  $[0.64, 0.96]$ ), while edges between different clusters have a correlation of  $\frac{\rho}{2}$ . They evaluated forty randomized routes per instance and found that the route evaluator assuming independence had an overall average of only 33.20% successful evaluations in scenarios with correlated travel times.

For the CVRP, Rostami et al. (2017) tested two correlation structures: one with correlations only between adjacent arcs (up to 75 customers) and another with correlations between all arcs (up to 32 customers). They generated random covariance matrices based on mean travel times and uniform random CoVs in the range  $[0.01, 0.2]$ , allowing negative covariances with a probability of 5%. Rostami et al. (2017) demonstrated that incorporating correlations results in routes with significantly less travel time variance (up to 70% for some instances), without considerably increasing the total expected travel time.

Bauer et al. (2019) state that since routing procedures can construct routes that include any pair of links, estimates of all covariances between pairs of link travel times are needed. They emphasize that this is impractical for large networks and, therefore, a model is required to estimate the correlation between two links. Their feed-forward neural network for correlation estimation is based on two factors: the driving distance between the links and their joint usage of roads. Using a floating taxi dataset in Vienna, the latter factor was found to have a significant impact. Moreover, correlations were almost always positive and high only for small distances. In particular, correlations were smaller than 0.4 for distances greater than 0.5 km. In a few cases, high correlations were observed on road segments in the opposite direction, suggesting common disturbances affecting both directions. Furthermore, empirical correlations demonstrated that the assumption of constant correlation over time-of-day intervals (e.g., one hour) is reasonable. Finally, Bauer et al. (2019) concluded that while including correlations in the route travel time variance calculations leads to improved results, “explicit modelling of the correlation only leads to minor performance enhancements compared to simple models using sample correlations for nearby links and setting the correlation to zero for distances larger than 1 km”.

Guo et al. (2019b) assumed that travel times on a link in one time period are correlated with those on the same link in the next and previous periods (with a correlation of 0.6). They also assumed that travel times on links sharing a node are correlated in the same period (correlation 0.7) and the next period (correlation 0.42). Guo et al. (2019b)'s largest instance of the Beijing network consists of 418 links and 60 time periods, resulting in over 310 million correlations (with 264,447 correlations treated as non-zero based on their assumptions). They highlight that even if data is available to estimate a correlation matrix, its high dimensionality still poses numerical challenges. They state that if data is not available, one can guess a correlation matrix, though this is almost certainly going to violate the positive semi-definite property. In that case, one can search for a positive semi-definite matrix close to the guessed matrix. However, they stress that this approach may lead to unexpected non-zero correlations. Alternatively, to evaluate the objective function of given solutions, Guo et al. (2019b) employed Kaut and Wallace (2007)'s method to generate scenario sets with specified marginal distributions and correlations. They claim that this method still results in some spurious correlations but showed that such correlations if they appear, do not greatly affect the route travel time distribution's estimations.

Zhang and Guo (2019) demonstrated the importance of considering stochastic and correlated travel speeds in SPPs with a sustainable objective function. Using 102 days of real-speed data from the Los Angeles freeway network and analysing 200 randomly selected origin-destination pairs, the study found that incorporating travel speed correlations can reduce the mean carbon emissions by 0.02% (approximately 2,630 kg in one month).

Using ideas from extreme value theory, Bakach et al. (2021) embedded correlations within a VRP heuristic and conducted tests using two patterns and different levels of correlation. In the first pattern, they defined three random incident locations in the Euclidean plane and assumed positive correlations between all arcs within five unit blocks. In the second pattern, they designated a central block and assumed that all arcs inside the block are positively correlated, all arcs outside are positively correlated, and the arcs between the two groups are negatively correlated. Bakach et al. (2021) observed significant impacts of correlations on the distribution of the makespan, even for small instances with twenty customers. Compared to the scenario assuming no correlations, savings of up to 13.76% were found.

Rajabi-Bahaabadi et al. (2021) were the first to address the VRP with soft TWs and correlated travel times. They represented the correlations using a variance-covariance matrix and developed a hybrid max-min ant colony system combined with TS. To generate the correlations, they employed a method similar

to that of Zockaie et al. (2013), where the travel time correlations are based on travel time variances and random number generations between  $-1$  and  $1$ . Using modified versions of Solomon’s VRP datasets with 100 customers, they demonstrated that neglecting travel time correlations results in suboptimal solutions. This occurs because the penalty cost estimates associated with earliness and tardiness are biased due to inaccurate arrival time distributions. Furthermore, real travel time data from two case study sites (Seattle, Washington, and downtown San Diego, California) revealed significant positive correlations between travel times, which followed the SLN distribution.

For a network with 21 links, Barahimi et al. (2022) used Monte Carlo simulation to generate a correlation matrix. They compared the travel times between different origin-destination pairs with those of Srinivasan et al. (2014), who assumed correlations only between adjacent links. They observed that the longer the route or the more links it contains, the greater the discrepancy in the total travel time, due to the inclusion of more correlations.

Bomboi et al. (2022) assumed that travel times are jointly normally distributed with covariances. They argue that these covariances improve the accuracy of customer arrival time estimations. To estimate the travel time covariances, they used traffic data collected over 25 consecutive days in Duisburg, Germany. To ensure positive definiteness, Bomboi et al. (2022) added  $10^{-4}$  to all diagonal entries of the covariance matrix (where necessary). They reported correlations ranging from  $-0.79$  to  $1$  for adjacent arcs, and from  $-0.81$  to  $1$  for non-adjacent arcs, with average correlations of  $0.64$  and  $0.60$ , respectively. Additionally, through sampling, they found that their chance constraints are more likely to be satisfied when correlations are considered.

In summary, the reviewed literature underscores the significant impact of travel time correlations on routing decisions and reliability assessments across various transportation applications. While numerous modelling approaches have been proposed – from simplifying assumptions based on adjacency or spatial proximity, to sampling- and copula-based methods – the scalability and realism of these methods remain limited, particularly in large-scale networks. These insights have directly informed the present study, motivating the development of a correlation structure that is both computationally tractable and behaviourally realistic. By explicitly incorporating this structure into our algorithmic framework, we aim to more accurately capture the influence of correlated travel times on route selection and TTR. As the literature suggests, such incorporation is likely to yield more reliable routing solutions compared to those derived under the common, but restrictive, assumption of independence.

## 6.3 Correlation Structure

For a transportation network with millions of arc pairs, fully modelling the correlation structure is computationally challenging, particularly due to memory constraints. Consequently, it is necessary to simplify the structure to enhance both memory efficiency and computational performance. We propose modelling relationships between travel times by first classifying pairs of arcs into seven categories in Section 6.3.1. Some of these categories are informed by prior research, while others are novel. A key contribution of our work is the consideration of pairs of arcs that overlap across some road segments in the network. To our knowledge, this is the first study to address correlations for such overlapping arc pairs. In Section 6.3.2, we then introduce a strategy for assigning spatial correlations to arc pairs within each category.

### 6.3.1 Classification of Arc Pairs

The classification of arc pairs is primarily based on studies demonstrating the spatial distance between arcs as a major predictor of travel time correlations. One such study is the work of Rachtan et al. (2013), who fitted regression models for correlation estimation using morning, midday, and afternoon data along an urban freeway in Los Angeles, California. They considered factors such as spatial distance, temporal distance, traffic flow, and the number of lanes, and found that spatial distance was the primary factor.

We begin to discuss our procedure for estimating the arc travel time correlations  $\rho_{e_1 e_2}$ , where  $e_1, e_2 \in E_1$  such that  $e_1 \neq e_2$ , featuring in Equation (5.11). Recall that  $E_1$  is the set of all arcs between the school and each bus stop, and between each pair of distinct bus stops. We start by classifying pairs of distinct arcs,  $e_1 = (u_1, v_1)$  and  $e_2 = (u_2, v_2)$ , into seven categories:

- (a) *First-order adjacent arcs*: pairs of arcs that share one vertex and have the same traffic flow direction;
- (b) *Second-order adjacent arcs*: pairs of arcs that share one vertex and have different traffic flow directions;
- (c) *Opposing arcs*: pairs of arcs that share two vertices and thus form a cycle;
- (d) *First-order nearby arcs*: pairs of arcs whose minimal distance between their endpoints is in the range  $(0, 0.5]$  km. That is,

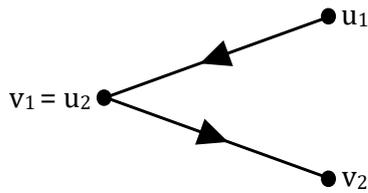
$$\min\{d(u_1, u_2), d(u_1, v_2), d(v_1, u_2), d(v_1, v_2), \\ d(u_2, u_1), d(u_2, v_1), d(v_2, u_1), d(v_2, v_1)\} \in (0, 0.5];$$

- (e) *Second-order nearby arcs*: pairs of arcs whose minimal distance between

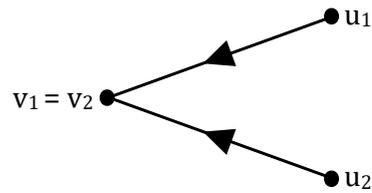
their endpoints is in the range  $(0.5, 1]$  km;

- (f) *Overlapping arcs*: pairs of arcs whose minimal distance between their endpoints is greater than 1 km and that share one or more common road segments;
- (g) *Unrelated arcs*: all other arc pairs that do not fall under (a)-(f).

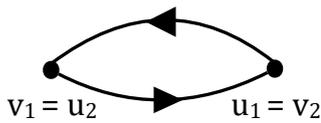
Figure 6.2 illustrates an example of an arc pair from each category. Categories



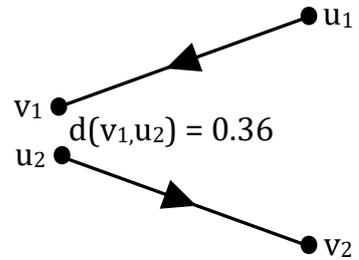
(a) First-Order Adjacent Arcs



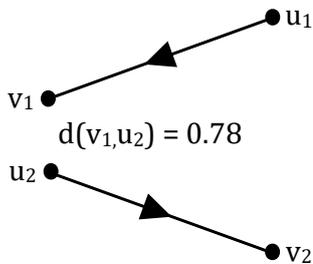
(b) Second-Order Adjacent Arcs



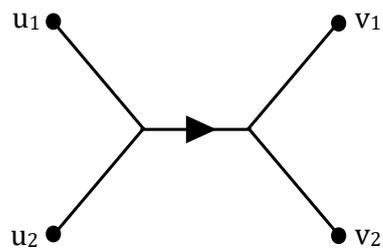
(c) Opposing Arcs



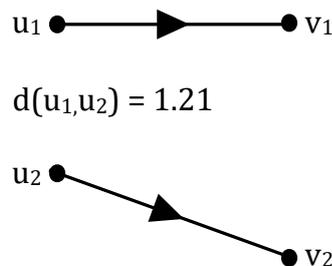
(d) First-Order Nearby Arcs



(e) Second-Order Nearby Arcs



(f) Overlapping Arcs



(g) Unrelated Arcs (all other arc pairs)

Figure 6.2: Classification of arc pairs for correlation structure.

(a) and (b) were motivated by the work of Guo et al. (2020) and by findings from several studies reporting strong correlations in travel times on adjacent arcs. The motivation for Category (c) is that a traffic incident influences traffic flow not only in the direction of the incident but also in the opposite direction. This is due to the likelihood of “rubbernecking”, where drivers in the opposite direction slow down out of curiosity or concern, often to look at the incident. This behaviour is discussed in more detail in Section 6.3.2.2. The motivation for Category (f) is that in our road network, the vertices in  $V_1$  represent bus stops (and the school) and the arcs in  $E_1$  represent the shortest paths between them. Thus, two arcs may overlap across some road segments, and this overlap should be accounted for when estimating the correlation.

The other categories are related to the spatial distance between the arcs. Correlation of arc travel times exists not only between adjacent arcs but also between non-adjacent arcs within close proximity. As explained above, we define nearby arcs according to the minimal distance (in km) between their endpoints. We use distance rather than travel time since this metric is fixed. The 1 km threshold in Category (e) is similar to that adopted by Bauer et al. (2019) and Guo et al. (2020), where nearby links are defined as those links whose midpoints are at most 1 km apart. Bauer et al. (2019) estimated non-zero correlations for pairs of links with distances smaller than 1 km, and set zero correlations for larger distances. As highlighted in Section 6.2.2, they observed that correlation tends to zero as distance increases, with correlations smaller than 0.4 for distances of 0.5 km. This latter observation motivated our selection of the 0.5 km threshold in Category (d).

### 6.3.2 Correlation Estimates

As discussed in this section, we assign the same (or a similar estimate, in the case of Category (f)) spatial correlation to arc pairs within the same category. We note that an arc pair in Categories (a) to (e) may still overlap across road segments. In such cases, the correlation estimate assigned to the arc pair is taken as the maximum of the estimates assigned to arc pairs in the respective category and those in Category (f).

#### 6.3.2.1 Adjacent Arcs

We begin by discussing the estimation of correlations between adjacent arcs. Evidence of strong correlations between travel times on adjacent links has been extensively reported in the literature (for example, Chan et al., 2009, Chen et al., 2012a, Gajewski and Rilett, 2005, He et al., 2006). As mentioned in Section 6.2.2,

Bomboi et al. (2022) found that the average correlation between adjacent arcs is 0.64. Similarly, Isukapati et al. (2013) reported correlations between 0.6 and 0.9 for adjacent arcs. However, some studies (for example, Cheng et al., 2011, Guo et al., 2020) have concluded that adjacent arcs may not be significantly correlated. Recall from Section 6.2.1 that Guo et al. (2020) found that, during morning hours, significant correlations were observed for only around 50% of the first-order adjacent arcs in the urban network. Based on these findings, we attempt to assign a correlation of 0.6 to a pair of first-order adjacent arcs with a probability of 0.5, and 0 otherwise. In this context, a correlation value of 0.6 indicates a strong positive correlation. The rationale for using the term ‘attempt to’ will be clarified in Section 6.3.3.

As claimed by Cheng et al. (2011) and highlighted in Section 6.2.1, arcs flowing into or out of the same vertex may indirectly influence each other. For this reason, we refer to such arcs as second-order adjacent arcs and attempt to assign a correlation of 0.4 between them with a probability of 0.5, and 0 otherwise. In this context, a 0.4 correlation value indicates a moderate positive correlation.

### 6.3.2.2 Opposing Arcs

If a driver is rubbernecking, their attention shifts from the direction of their driving to the incident scene in the opposite direction. This makes rubbernecking not only a traffic congestion issue (due to reduced vehicle speeds) but also a safety concern. A study conducted by Virginia Commonwealth University, which analysed more than 2,700 vehicle crashes, revealed that rubbernecking was the leading distraction in 16% of the crashes studied (Atkins, 2003).

Masinick et al. (2014) evaluated the effects of rubbernecking by studying data on 840 vehicle accidents in the Hampton Roads area in Virginia in 2000. They found that 12.1% of the accidents significantly impacted traffic occupancy (the percentage of time vehicles occupy a section of roadway) in the opposite direction. The data indicated that, on average, rubbernecking caused a 12.7% reduction in capacity (the maximum hourly rate at which vehicles traverse a section of roadway) and a delay (the difference between actual and free-flow travel times) of 107 vehicle hours.

More recently, Reina (2021) analysed 637 incidents along Interstate 5, California in 2016-2017 and found that 12.4% of them developed rubbernecking queues, corroborating Masinick et al. (2014)’s findings. The formation of these rubbernecking queues was confirmed by screening several daily speed datasets to ensure the queue resulted from non-recurrent congestion. Additionally, a survey study investigating drivers’ experiences and behaviours around rubbernecking events found that 58% of respondents reduced their speed when driving in the opposite

direction of an incident. The respondents attributed this behaviour to personal tendencies rather than factors related to the surroundings. Furthermore, 91% of respondents reported having experienced rubbernecking queues at least once.

Other researchers have also examined traffic disruptions related to rubbernecking. Knoop et al. (2008) collected microscopic traffic data from two incident locations in the Netherlands in 2007 and observed that road capacity in the opposite direction of an incident was reduced by half due to rubbernecking. In the same study, Knoop et al. (2009) noted an increase in the mean headway and reaction time, along with a decrease in travel speed and bottleneck discharge rate, which they attributed to rubbernecking. Chen et al. (2012b) analysed traffic oscillations from trajectory data collected between 2005 and 2006 as part of FHWA’s Next-Generation Simulation program. They suggested that rubbernecking likely caused the oscillations and reported significant relationships between the percentage of rubberneckerers, speed reduction, and the oscillation period.

These studies, along with others in the literature, support the premise that rubbernecking is likely to cause traffic delays. Rubbernecking can only occur if the opposing arcs share common road segments. If this is not the case, the opposing arcs are likely to have parallel road segments, which may also be correlated (El Esawey and Sayed, 2011). Therefore, we attempt to assign a correlation of 0.2 to a pair of opposing arcs with a probability of 0.5, and 0 otherwise. We treat a correlation value of 0.2 as representing a weak positive correlation.

### 6.3.2.3 Nearby and Unrelated Arcs

It is reasonable to assume that the more spatially distant links are, the lower the impact they will exert on each other. This phenomenon, highlighted by Chen et al. (2012a), Zou et al. (2014), and Ermagun et al. (2017) amongst others, is related to Tobler’s First Law of Geography. This law states that “everything is related to everything else, but near things are more related to each other”.

Chen et al. (2012a) state that correlations can only be modelled for arcs that are spatially very close. This is the strategy taken here, where arc pairs at a minimal distance greater than 1 km that do not overlap (termed unrelated arcs) are assumed to have zero correlation. Moreover, we attempt to assign a correlation of 0.4 (0.2) to a pair of first-order (second-order) nearby arcs with a probability of 0.5, and 0 otherwise.

### 6.3.2.4 Overlapping Arcs

To account for overlap across arcs, we utilize the manoeuvre data provided by the Bing Maps Routes application programming interface. For each arc, this data

contains a collection of sequential GPS coordinates that connect to form the arc. For each pair of distinct arcs  $e_1 = (u_1, v_1)$  and  $e_2 = (u_2, v_2)$ , we identify any common subcollections of coordinates, along with their corresponding shortest driving distances. We then set the total driving distance  $CD_{e_1e_2}$  as an estimate of the overlapping distance between the two arcs. It is important to note that this process is computationally demanding, especially when  $|V_1|$  exceeds 300 or so. For example, a network with 500 bus stops involves  $500 \cdot 501 = 250,500$  arcs, resulting in  $\binom{250,500}{2} \approx 31.37$  billion pairs of distinct arcs. Moreover, if each value  $CD_{e_1e_2}$  is stored as an 8-byte floating-point number, the total memory required would exceed 233 GB.

Finally, we estimate the correlation  $\rho_{e_1e_2}$  using the cosine similarity index given by:

$$\rho_{e_1e_2} = \frac{CD_{e_1e_2}}{\sqrt{d(u_1, v_1)d(u_2, v_2)}}.$$

Note that this index, also used by Chen et al. (2003a), assigns zero correlation if  $CD_{e_1e_2} = 0$  and a correlation in the range  $(0,1]$  if  $CD_{e_1e_2} > 0$ .

### 6.3.2.5 Summary

To summarize, correlation estimates are assigned according to the values specified in Table 6.1.

Table 6.1: Summary of assigned correlation estimates.

Arc Pair Category	Correlation
First-order adjacent arcs	0.6 with probability 0.5, 0 otherwise
Second-order adjacent arcs	0.4 with probability 0.5, 0 otherwise
Opposing arcs	0.2 with probability 0.5, 0 otherwise
First-order nearby arcs	0.4 with probability 0.5, 0 otherwise
Second-order nearby arcs	0.2 with probability 0.5, 0 otherwise
Overlapping arcs	Using $\rho_{e_1e_2} = \frac{CD_{e_1e_2}}{\sqrt{d(u_1, v_1)d(u_2, v_2)}}$ for $e_1, e_2 \in E_1, e_1 \neq e_2$
Unrelated arcs	0

### 6.3.3 Constructing the Correlation Matrix and Ensuring Positive Definiteness

In Section 6.3.2, we provided estimates for the arc travel time correlations  $\rho_{e_1e_2}$ , where  $e_1, e_2 \in E_1$  and  $e_1 \neq e_2$ , as summarized in Table 6.1. These correlations were adopted from values reported in the literature and specified pairwise, without imposing a joint correlation structure. For a matrix to be a valid correlation matrix, all of its eigenvalues must be non-negative. When correlations are specified independently for many arc pairs, there is no guarantee that this condition

will hold. In particular, small inconsistencies among subsets of correlations (for example, among three or more arcs) can accumulate as the dimension of the matrix increases. As a result, the assembled matrix  $\boldsymbol{\rho}$  may exhibit negative eigenvalues. Thus, even when all assumed correlations lie within the interval  $[-1, 1]$  and appear reasonable individually, the resulting matrix  $\boldsymbol{\rho}$  may fail to be positive semi-definite. For this reason, and as suggested by Guo et al. (2019b), we can search for a positive semi-definite matrix close to  $\boldsymbol{\rho}$ . A similar approach was adopted by Zhang and Khani (2019), who generated normally distributed samples for the reliable SPP for three street networks, Anaheim, Barcelona, and Chicago. Due to numerical issues in the computations, they noted that the variance-covariance matrix might not be positive semi-definite. Therefore, they found the nearest symmetric positive semi-definite matrix in the Frobenius norm to obtain a valid variance-covariance matrix.

Here, we employ the *scaled spectral method* to adjust  $\boldsymbol{\rho}$  so that it becomes positive semi-definite. The reader is referred to Marée (2012) for alternative methods such as the *iterative spectral method*. The scaled spectral method begins by performing a spectral decomposition of  $\boldsymbol{\rho}$  into its eigenvalues and eigenvectors:

$$\boldsymbol{\rho} = \mathbb{V}\boldsymbol{\Lambda}\mathbb{V}^T,$$

where  $\boldsymbol{\Lambda}$  is the diagonal matrix of eigenvalues, and  $\mathbb{V}$  is the matrix of eigenvectors. To make  $\boldsymbol{\rho}$  positive semi-definite (i.e., with all eigenvalues being non-negative), the method then replaces all the negative eigenvalues in  $\boldsymbol{\Lambda}$  by a minimal positive value (such as  $10^{-6}$ ). Call the resultant matrix  $\boldsymbol{\Lambda}^+$ . The nearest symmetric positive semi-definite matrix to  $\boldsymbol{\rho}$  in the Frobenius norm is given by

$$\boldsymbol{\rho}_{\text{PSD}} = \mathbb{V}\boldsymbol{\Lambda}^+\mathbb{V}^T.$$

Since we are dealing with correlation matrices, we note that the matrix  $\boldsymbol{\rho}_{\text{PSD}}$  yielded by the scaled spectral method may not have ones on the main diagonal. Therefore, for the correlation matrix to be valid, a diagonal scaling matrix  $\mathbb{S}$  should be defined with diagonal entries  $s_{i,i} = r_{i,i}^{-1}$ , where  $r_{i,i}$  is the  $i^{\text{th}}$  diagonal entry of  $\boldsymbol{\rho}_{\text{PSD}}$ . Finally,  $\boldsymbol{\rho}_{\text{PSD}}$  should be scaled as

$$\sqrt{\mathbb{S}}\boldsymbol{\rho}_{\text{PSD}}\sqrt{\mathbb{S}},$$

which results in a valid correlation matrix (Marée, 2012).

Recall from Algorithm 1 that our heuristic algorithm uses a varying subset of visited bus stops  $V_1'$ . The scaled spectral method is non-iterative and requires only a single spectral decomposition. However, it remains computationally expensive,

particularly in terms of memory requirements, when dealing with hundreds of thousands of arcs or more. For this reason, we limit the correlation matrix  $\boldsymbol{\rho}$  to only include arcs between pairs of vertices in  $V_1'$ , rather than the entire set  $V_1$ . This restriction results in a much smaller matrix  $\boldsymbol{\rho}$ , making it easier to handle in terms of memory. That said, note that  $\boldsymbol{\rho}$  must be updated in each iteration of our heuristic algorithm to reflect changes in  $V_1'$ .

Moreover, we observed from trial experiments that since our correlation matrices contain many zeros, modifying any negative eigenvalues tends to suppress the positive correlations. To address this, we replaced the values 0.6, 0.4, and 0.2 in the original  $\boldsymbol{\rho}$  matrix with 0.9, 0.6, and 0.3, corresponding to strong, moderate, and weak correlations, respectively. In this way, the final values in  $\boldsymbol{\rho}_{\text{PSD}}$  are closer to the intended values. The scaled spectral method may also introduce negative correlations in place of zero or very small correlations. This results in a mixture of positive and negative correlations, which is not problematic as long as the following observations about the lognormal distribution are respected.

A lognormal distribution is neither symmetric nor shape-invariant. Due to the absence of these characteristics, the correlation between two lognormally distributed random variables  $X_1 \sim \text{LN}(\mu_1, \sigma_1)$  and  $X_2 \sim \text{LN}(\mu_2, \sigma_2)$  is limited to a range  $[l, u]$  that may not be  $[-1, 1]$ . The limits  $l$  and  $u$  depend on the distribution parameters  $\sigma_1$  and  $\sigma_2$  and have been derived by Žerovnik et al. (2013) as:

$$\begin{aligned} l &= \frac{1}{\eta_1 \eta_2} \left( \exp[-\sigma_1 \sigma_2] - 1 \right), \\ u &= \frac{1}{\eta_1 \eta_2} \left( \exp[\sigma_1 \sigma_2] - 1 \right), \end{aligned} \tag{6.1}$$

where  $\eta_i = \sqrt{\exp[\sigma_i^2] - 1}$  for  $i \in \{1, 2\}$ . As noted by Žerovnik et al. (2013), it is seen from Equation (6.1), that for large values of  $\eta_1$  and  $\eta_2$ , strong negative correlations are prohibited. Recall that in our case,  $\sigma_e, e \in E_1$ , depends on  $\alpha$  and  $\beta$  as given in Equation (5.5). Therefore,  $l = -\frac{\alpha^2}{\alpha^2 + \beta^2}$  and  $u = 1$ . Since  $l$  is always strictly greater than  $-1$  in our case (otherwise  $\beta = 0$ ), we must ensure that any negative correlation is no smaller than  $l$ ; otherwise, it must be clipped to  $l$ .

To illustrate the outcome of the scaled spectral method, consider a small-scale example for the Mgarr instance with  $V_1' = \{11, 31, 39\}$ ,  $\alpha = 0.2$ ,  $\beta = 0.5$ , and  $l = -0.138$ . For three visited stops and the school, the number of arcs is 12, which are enumerated in Figure 6.3. In this figure, the lower triangular part of the symmetric matrix  $\boldsymbol{\rho}_{\text{PSD}}$  is displayed. The legend of the figure shows each entry's original value in the matrix  $\boldsymbol{\rho}$ . As discussed above, the values 0.9, 0.6, and 0.3 were used in place of the intended values 0.6, 0.4, and 0.2, respectively.

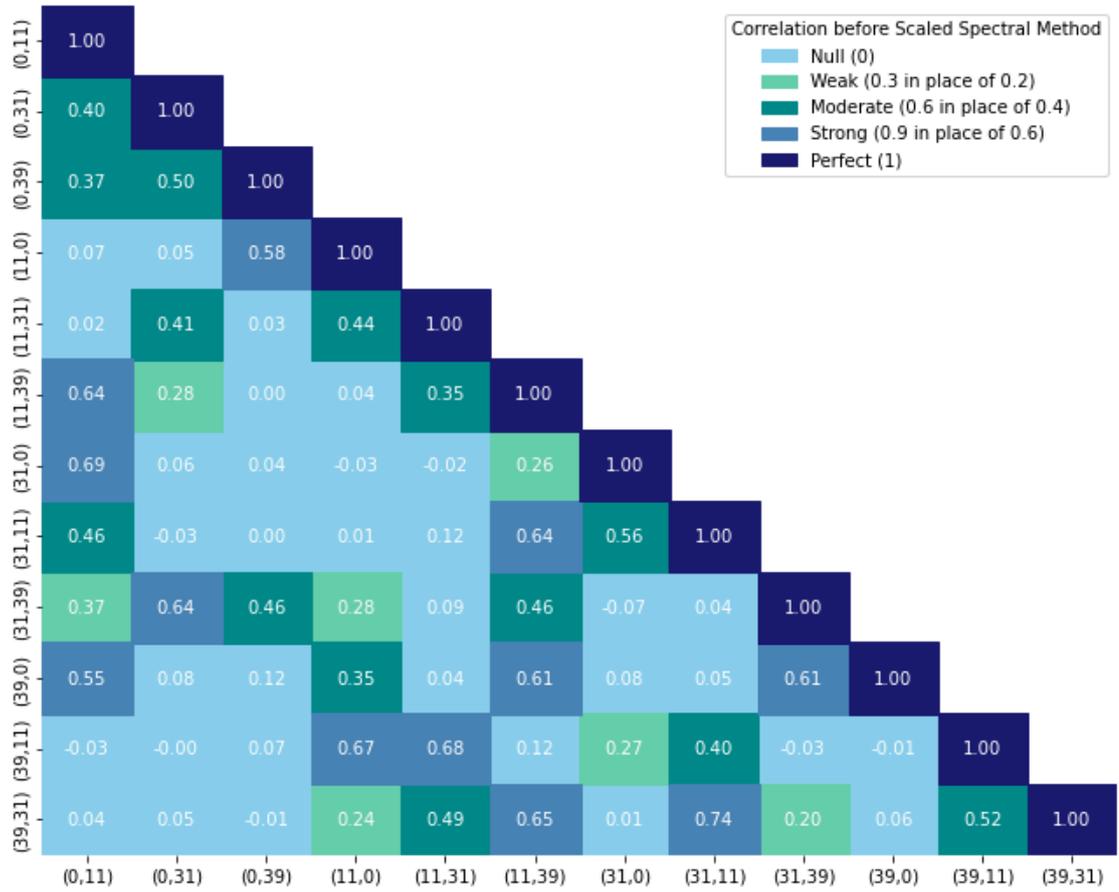


Figure 6.3: Example of a valid correlation matrix yielded by the scaled spectral method.

In Figure 6.3, we note, for example:

- the pair of first-order adjacent arcs (0, 11) and (11, 39) with a final correlation value of 0.64 (close to the 0.6 intended value);
- the pair of opposing arcs (31, 39) and (39, 31) with a final correlation value of 0.20 (approximately equal to the 0.2 intended value);
- the pair of second-order nearby arcs (11, 0) and (39, 31) (minimal distance of 0.989 km) with a final correlation value of 0.24 (close to the 0.2 intended value).

No pairs of overlapping arcs exist in this example. The Frobenius norm of the difference  $\boldsymbol{\rho} - \boldsymbol{\rho}_{\text{PSD}}$  was calculated to be 1.167. Note that most entries in  $\boldsymbol{\rho}_{\text{PSD}}$  align very closely with their intended values. However, some entries deviate marginally, such as the correlations between the arc pairs (31,0) and (31,11) (0.56 versus the intended value of 0.4), and (0,11) and (31,39) (0.37 versus the intended value of 0.2). The majority of these deviations are conservative, meaning they overestimate the correlations relative to the intended values. Consequently, the resulting TTR evaluations will be slightly more stringent. Figure 6.3 therefore not only illustrates the general agreement between the intended and resultant

correlations but also provides visual confirmation that these minor deviations do not compromise the assessment of TTR.

## 6.4 Computational Results

In this section, we present the results for the correlated scenario with parameters  $\alpha = 0.2$  and  $\beta = 0.5$ . The source code used for this scenario can be downloaded from (Sciortino, 2025b). As in the independent scenario, the parameters  $\alpha$  and  $\beta$ , as well as the reliability level  $\kappa$ , can be adjusted as needed. For comparative purposes, some results from Chapters 4 and 5 are also revisited.

We start with Table 6.2, which presents the proportions of arc pairs in each category defined in Section 6.3.1. The table summarizes how pairs of distinct arcs (PDA) are distributed across the seven categories, ranging from first-order adjacent to unrelated arcs. PDA exhibits rapid growth across locations with progressively larger  $|V_1|$ . The majority of arc pairs fall under category (g), whereas the proportions of first- and second-order adjacent arcs (categories (a) and (b)) remain relatively small. The nearby categories (d) and (e) and, to a lesser extent the overlapping category (f), account for a more substantial share, and are therefore critical when assessing spatial interaction and structural connectivity between arcs. If a broader inclusion of arcs outside category (g) is desired, the 1 km threshold used to define the nearby categories can be increased to capture additional pairs.

For each pair of distinct arcs, we verified, using Python version 3.10.9 and the code available at (Sciortino, 2025a), whether the overlap across them is strictly positive. In cases where an overlap exists, the corresponding correlation value  $\rho_{e_1e_2}$  was stored in a text file, provided that it exceeded the threshold specified in Table 6.1 for categories (a)-(e). For category (f), all such correlations were retained regardless of their magnitude. This preprocessing stage required substantial storage capacity and computational time, ranging from a few megabytes and minutes for the smallest instance up to approximately 1.5 terabytes and two weeks for the largest instance. During the algorithmic phase, the relevant correlation values associated with a selected subset of bus stops were then retrieved on demand by means of an unordered map implemented with Robin Hood hashing. In this way, only the data rows that were actually required were loaded into memory. This two-stage design – precomputing and storing overlap relationships rather than performing the overlap checks dynamically within the algorithm – offers a key advantage: modifications to algorithmic components, such as changes in the bus stop selection mechanism, do not necessitate recomputation of overlaps. Since the overlap analysis is carried out once in advance, subsequent experiments

Table 6.2: Proportions of arc pairs in each category (%): (a) first-order adjacent, (b) second-order adjacent, (c) opposing, (d) first-order nearby, (e) second-order nearby, (f) overlapping, and (g) unrelated. PDA stands for the number of pairs of distinct arcs (in millions).

Location	PDA	(a)	(b)	(c)	(d)	(e)	(f)	(g)
Mġarr	6.26	3.28	3.28	0.03	18.23	20.99	2.53	51.67
Mellieħa	26.71	2.30	2.30	0.01	14.94	15.54	4.03	60.88
Porthcawl	270.41	1.30	1.30	0.00	19.94	25.78	1.54	50.14
Qrendi	307.66	1.26	1.26	0.00	15.43	21.54	4.65	55.86
Suffolk	453.05	1.14	1.14	0.00	6.98	8.42	5.31	77.00
Senglea	592.01	1.07	1.07	0.00	13.33	22.41	4.61	57.51
Victoria	4,954.06	0.63	0.63	0.00	8.87	13.07	5.29	71.50
Pembroke	5,341.80	0.62	0.62	0.00	8.39	15.56	4.20	70.61
Canberra	5,965.54	0.60	0.60	0.00	6.27	9.54	5.70	77.29
Handaq	11,866.55	0.51	0.51	0.00	5.82	11.29	5.70	76.17
Valetta	19,518.83	0.45	0.45	0.00	7.69	12.86	3.43	75.12
Birkirkara	24,088.26	0.43	0.43	0.00	5.79	11.99	3.84	77.53
Hamrun	35,859.89	0.39	0.39	0.00	6.46	10.81	3.75	78.21
Cardiff	46,254.07	0.36	0.36	0.00	13.62	17.13	2.88	65.65
Milton Keynes	55,999.16	0.34	0.34	0.00	4.94	8.17	2.45	83.74
Bridgend	80,022.20	0.32	0.32	0.00	4.95	7.37	6.96	80.09
Edinburgh-2	352,776.06	0.22	0.22	0.00	9.76	19.91	2.87	67.03
Edinburgh-1	422,024.60	0.21	0.21	0.00	4.99	9.62	2.68	82.30
Adelaide	994,269.27	0.17	0.17	0.00	2.91	6.60	1.32	88.84
Brisbane	5,443,916.00	0.11	0.11	0.00	2.40	4.90	1.35	91.13

can simply reuse the stored correlation data as needed.

Table 6.3 presents the results for the correlated scenarios (Columns 6–8) alongside those for the deterministic (Columns 2–5) and independent (Columns 9–12) scenarios. The IR columns indicate the number of runs in which the deterministic or independent solution is infeasible with respect to the on-time arrival chance constraints (5.15). From Column 4, we observe that for eleven instances, at least one solution generated under the deterministic scenario becomes infeasible when correlations are considered, with several cases where all solutions are infeasible. Similarly, Column 11 shows that for nine instances, at least one solution generated under the independent scenario is infeasible. The average number  $k$  of routes across the 25 runs is the same for the independent and correlated scenarios, except for the Victoria instance, where all runs under the correlated scenarios include five routes.

For details on the computational times and feasibility rates of the correlated scenario runs, the reader is referred to Appendix D.1. Results show that the average computational time of the correlated scenario is higher than that of the deterministic and independent scenarios for most locations. Relative to the deterministic scenario, the increase ranges from approximately 4% up to 45,661%,

Table 6.3: Results averaged across 25 runs for the correlated scenario with parameters  $\alpha = 0.2, \beta = 0.5$ . TMT and TPT stand for Total Mean Time and Total 99<sup>th</sup> Percentile Time (minutes), presented as mean  $\pm$  standard deviation. Comparative results for the deterministic and independent scenarios are also presented, with the column IR indicating the number of runs for which the solution is infeasible with respect to the on-time arrival chance constraints. An asterisk signifies an instance for which  $\kappa$  was reduced to 0.95 due to infeasibility.

Location	Deterministic Scenario					Correlated Scenario					Independent Scenario				
	$k$	TMT	IR	TPT	$k$	TMT	TPT	$k$	TMT	IR	TPT	$k$	TMT	IR	TPT
Mgarr	4	54.10 $\pm$ 0.00	0	102.91 $\pm$ 0.38	4	54.80 $\pm$ 0.26	99.60 $\pm$ 0.20	4.00	54.91 $\pm$ 0.26	0	100.16 $\pm$ 0.82	4	54.91 $\pm$ 0.26	0	100.16 $\pm$ 0.82
Melieha	4	56.30 $\pm$ 0.00	0	107.55 $\pm$ 0.04	4	56.77 $\pm$ 0.17	103.33 $\pm$ 0.37	4.00	57.06 $\pm$ 0.28	0	104.97 $\pm$ 0.35	4	57.06 $\pm$ 0.28	0	104.97 $\pm$ 0.35
Porthcawl	1	27.83 $\pm$ 0.43	25	-	2	24.28 $\pm$ 0.00	52.66 $\pm$ 0.04	2.00	24.32 $\pm$ 0.11	0	52.86 $\pm$ 0.42	1	24.32 $\pm$ 0.11	0	52.86 $\pm$ 0.42
Qrendi	5	75.73 $\pm$ 0.00	0	141.10 $\pm$ 0.55	5	76.37 $\pm$ 1.07	137.64 $\pm$ 1.29	5.00	77.37 $\pm$ 1.50	0	141.74 $\pm$ 1.40	5	77.37 $\pm$ 1.50	0	141.74 $\pm$ 1.40
Suffolk*	3	116.61 $\pm$ 0.90	25	-	4	117.27 $\pm$ 1.12	157.94 $\pm$ 1.45	4.00	116.01 $\pm$ 0.87	17	-	3	116.01 $\pm$ 0.87	17	-
Senglea	6	72.01 $\pm$ 0.54	0	145.54 $\pm$ 2.33	6	75.41 $\pm$ 1.45	143.16 $\pm$ 2.39	6.00	74.50 $\pm$ 1.26	0	142.82 $\pm$ 2.39	6	74.50 $\pm$ 1.26	0	142.82 $\pm$ 2.39
Victoria	4	92.96 $\pm$ 0.41	25	-	5	92.56 $\pm$ 0.29	160.74 $\pm$ 0.81	4.32	93.94 $\pm$ 1.69	16	-	4	93.94 $\pm$ 1.69	16	-
Pembroke	7	103.49 $\pm$ 0.68	0	207.16 $\pm$ 2.13	7	106.85 $\pm$ 2.13	200.53 $\pm$ 3.14	7.00	106.26 $\pm$ 1.24	0	201.90 $\pm$ 2.85	7	106.26 $\pm$ 1.24	0	201.90 $\pm$ 2.85
Camberra*	7	179.11 $\pm$ 0.53	25	-	8	190.50 $\pm$ 0.82	261.37 $\pm$ 1.06	8.00	190.27 $\pm$ 0.96	1	-	7	190.27 $\pm$ 0.96	1	-
Handaq	6	96.61 $\pm$ 0.57	1	-	6	98.34 $\pm$ 1.33	177.05 $\pm$ 3.40	6.00	98.31 $\pm$ 1.01	3	-	6	98.31 $\pm$ 1.01	3	-
Valletta	6	104.50 $\pm$ 0.78	21	-	6	109.42 $\pm$ 2.02	210.55 $\pm$ 3.23	6.00	110.11 $\pm$ 2.12	20	-	6	110.11 $\pm$ 2.12	20	-
Birkirkara	6	97.49 $\pm$ 2.64	3	-	6	97.47 $\pm$ 2.22	176.14 $\pm$ 5.06	6.00	97.34 $\pm$ 2.50	1	-	6	97.34 $\pm$ 2.50	1	-
Hamrun	7	99.31 $\pm$ 1.40	0	192.32 $\pm$ 3.42	7	100.91 $\pm$ 1.35	183.01 $\pm$ 3.38	7.00	100.62 $\pm$ 1.26	0	183.79 $\pm$ 3.05	7	100.62 $\pm$ 1.26	0	183.79 $\pm$ 3.05
Cardiff	2	67.12 $\pm$ 0.00	25	-	4	69.05 $\pm$ 0.56	154.37 $\pm$ 0.06	4.00	67.92 $\pm$ 0.00	0	156.54 $\pm$ 0.26	2	67.92 $\pm$ 0.00	0	156.54 $\pm$ 0.26
M. Keynes	4	59.88 $\pm$ 0.57	0	117.69 $\pm$ 1.63	4	59.65 $\pm$ 0.57	114.79 $\pm$ 1.26	4.00	59.36 $\pm$ 0.56	0	114.49 $\pm$ 1.62	4	59.36 $\pm$ 0.56	0	114.49 $\pm$ 1.62
Bridgend*	6	170.55 $\pm$ 0.53	25	-	7	179.02 $\pm$ 0.56	249.87 $\pm$ 0.70	7.00	179.09 $\pm$ 0.76	2	-	6	179.09 $\pm$ 0.76	2	-
Edinburgh-2	4	59.13 $\pm$ 0.00	0	110.85 $\pm$ 0.00	4	59.13 $\pm$ 0.00	110.84 $\pm$ 0.03	4.00	60.11 $\pm$ 0.04	0	111.88 $\pm$ 0.10	4	60.11 $\pm$ 0.04	0	111.88 $\pm$ 0.10
Edinburgh-1	9	143.71 $\pm$ 1.56	5	-	9	146.30 $\pm$ 2.57	273.68 $\pm$ 5.29	9.00	145.93 $\pm$ 2.80	1	-	9	145.93 $\pm$ 2.80	1	-
Adelaide	8	128.18 $\pm$ 1.45	7	-	8	130.71 $\pm$ 2.06	245.03 $\pm$ 4.80	8.00	129.49 $\pm$ 1.62	1	-	8	129.49 $\pm$ 1.62	1	-
Brisbane*	10	213.41 $\pm$ 4.31	0	288.76 $\pm$ 6.37	10	212.42 $\pm$ 3.81	285.86 $\pm$ 5.26	10.00	213.42 $\pm$ 3.66	0	287.42 $\pm$ 4.96	10	213.42 $\pm$ 3.66	0	287.42 $\pm$ 4.96

with absolute differences between 31 seconds and 16 hours. Compared to the independent scenario, the correlated scenario shows increases from around 64% to 13,079%, corresponding to absolute differences between 11 minutes and 16 hours.

Mann-Whitney tests were applied to compare the 25 pairs of TMTs and TPTs under the deterministic and correlated scenarios. As shown in Table 6.4, across the nine comparable instances, the deterministic solutions yield significantly lower TMTs in six cases, whereas the correlated solutions achieve significantly lower TPTs in eight cases ( $p < 0.05$ ). The relative deterioration in the average TMT from the deterministic to the correlated scenario ranges between  $-0.47\%$  to  $4.72\%$ , while the relative improvement in the average TPT ranges between  $0.01\%$  to  $4.84\%$ . The analysis was also conducted between the independent and correlated scenarios. Results in Table 6.5 show that the independent solu-

Table 6.4: Mann-Whitney test results comparing the TMTs and TPTs of the deterministic and correlated scenarios. The Det. and Corr. columns report the mean ranks,  $U$  represents the test statistic, while bold  $p$ -values are less than 0.05.

Location	TMT				TPT			
	$U$	$p$ -value	Det.	Corr.	$U$	$p$ -value	Det.	Corr.
Mgarr	0.0	< <b>0.001</b>	13.00	38.00	0.0	< <b>0.001</b>	38.00	13.00
Mellieha	0.0	< <b>0.001</b>	13.00	38.00	0.0	< <b>0.001</b>	38.00	13.00
Qrendi	0.0	< <b>0.001</b>	13.00	38.00	32.0	< <b>0.001</b>	36.72	14.28
Senglea	0.0	< <b>0.001</b>	13.00	38.00	162.0	<b>0.003</b>	31.52	19.48
Pembroke	28.0	< <b>0.001</b>	14.12	36.88	30.0	< <b>0.001</b>	36.80	14.20
Hamrun	119.0	< <b>0.001</b>	17.76	33.24	6.0	< <b>0.001</b>	37.76	13.24
Milton Keynes	270.5	0.390	27.18	23.82	35.0	< <b>0.001</b>	36.60	14.40
Edinburgh-2	312.5	1.000	25.50	25.50	262.5	<b>0.039</b>	27.50	23.50
Brisbane	287.0	0.621	26.52	24.48	226.0	0.093	28.96	22.04

Table 6.5: Mann-Whitney test results comparing the TMTs and TPTs of the independent and correlated scenarios. The Ind. and Corr. columns report the mean ranks,  $U$  represents the test statistic, while bold  $p$ -values are less than 0.05.

Location	TMT				TPT			
	$U$	$p$ -value	Ind.	Corr.	$U$	$p$ -value	Ind.	Corr.
Mgarr	246.5	0.196	28.14	22.86	186.5	<b>0.014</b>	30.54	20.46
Mellieha	107.0	< <b>0.001</b>	33.72	17.28	0.0	< <b>0.001</b>	38.00	13.00
Porthcawl	275.0	0.077	27.00	24.00	66.0	< <b>0.001</b>	35.36	15.64
Qrendi	290.5	0.656	26.38	24.62	24.0	< <b>0.001</b>	37.04	13.96
Senglea	171.0	<b>0.006</b>	19.84	31.16	254.0	0.256	23.16	27.84
Pembroke	272.5	0.438	23.90	27.10	227.0	0.097	28.92	22.08
Hamrun	271.5	0.426	23.86	27.14	273.0	0.443	27.08	23.92
Cardiff	0.0	< <b>0.001</b>	13.00	38.00	0.0	< <b>0.001</b>	38.00	13.00
Milton Keynes	183.0	<b>0.009</b>	20.32	30.68	307.0	0.914	25.72	25.28
Edinburgh-2	0.0	< <b>0.001</b>	38.00	13.00	0.0	< <b>0.001</b>	38.00	13.00
Brisbane	259.0	0.299	27.64	23.36	256.0	0.273	27.76	23.24

tions have significantly lower TMTs in five of the eleven comparable instances, whereas the correlated solutions outperform in TPTs for six instances ( $p < 0.05$ ). The relative deterioration in the average TMT from the independent to the correlated scenario varies from  $-1.63\%$  to  $1.66\%$  (six of which are negative), and the relative improvement in the average TPT ranges between  $-0.26\%$  to  $2.89\%$ .

The solutions with the lowest TPT for the correlated scenario are presented in Table 6.6 and illustrated in Appendix D.2. Note that the Edinburgh-2 solution was obtained over three runs. Table 6.6 reveals that 89% of the routes exhibit capacity utilization exceeding 90%, 82% have a mean journey time less than half  $m_t = 45$  minutes, and 78% have a percentile journey time below 35 minutes.

Table 6.6: Best solutions for the correlated scenario with parameters  $\alpha = 0.2, \beta = 0.5$ . Cap. is short for Capacity, MT stands for Mean Time (minutes), and PT stands for Percentile Time (minutes). The routes are presented in terms of the indices of the vertices in  $V_1$ , and the subscripts represent the number of boarding students.

Location	Routes	Load/Cap.	MT,PT	TMT, TPT
Mġarr	$48_1 \rightarrow 33_9 \rightarrow 11_{10} \rightarrow 31_{29} \rightarrow 0$	49/53	14.45,26.07	
	$42_{10} \rightarrow 19_{10} \rightarrow 45_{20} \rightarrow 0$	40/44	15.32,30.63	54.73,
	$38_{18} \rightarrow 21_{32} \rightarrow 19_1 \rightarrow 45_1 \rightarrow 0$	52/53	13.45,22.95	99.28
	$52_{20} \rightarrow 57_{15} \rightarrow 35_{12} \rightarrow 15_2 \rightarrow 0$	49/53	11.52,19.63	
Mellieħa	$66_5 \rightarrow 83_9 \rightarrow 79_{14} \rightarrow 34_{23} \rightarrow 7_1 \rightarrow 0$	52/53	16.28,27.87	
	$6_5 \rightarrow 7_4 \rightarrow 9_{28} \rightarrow 0$	37/44	9.75,17.74	56.83,
	$43_{14} \rightarrow 45_3 \rightarrow 46_{14} \rightarrow 64_{19} \rightarrow 0$	50/53	15.87,29.34	102.81
Porthcawl	$30_8 \rightarrow 20_{14} \rightarrow 25_2 \rightarrow 27_8 \rightarrow 0$	32/36	14.93,27.85	
	$61_5 \rightarrow 0$	5/8	8.52,22.35	24.28,
Qrendi	$79_{34} \rightarrow 38_{27} \rightarrow 0$	61/61	15.77,30.25	52.61
	$124_{15} \rightarrow 126_{25} \rightarrow 155_4 \rightarrow 149_9 \rightarrow 0$	53/53	20.08,37.03	
	$17_{43} \rightarrow 76_9 \rightarrow 0$	52/53	13.72,25.78	75.92,
	$67_3 \rightarrow 68_{18} \rightarrow 63_{20} \rightarrow 137_{12} \rightarrow 0$	53/53	16.72,29.09	137.08
	$50_{18} \rightarrow 106_{16} \rightarrow 146_1 \rightarrow 54_{10} \rightarrow 0$	45/53	11.98,20.16	
Suffolk <sup>a</sup>	$112_{27} \rightarrow 76_{25} \rightarrow 0$	52/53	13.42,25.02	
	$43_4 \rightarrow 116_{18} \rightarrow 172_2 \rightarrow 5_{22} \rightarrow 98_4 \rightarrow 0$	80/80	31.08,40.80	
	$70_{29} \rightarrow 83_6 \rightarrow 150_5 \rightarrow 122_{12} \rightarrow 3_5 \rightarrow 0$	57/57	32.63,43.63	116.67,
	$173_1 \rightarrow 102_2 \rightarrow 67_4 \rightarrow 9_8 \rightarrow 65_{12} \rightarrow 0$	27/27	23.73,33.15	156.82
Senglea	$78_9 \rightarrow 158_9 \rightarrow 111_{17} \rightarrow 140_8 \rightarrow 43_2 \rightarrow 0$	45/45	29.22,39.25	
	$14_{30} \rightarrow 106_{16} \rightarrow 53_6 \rightarrow 0$	52/53	14.48,26.87	
	$74_3 \rightarrow 82_{21} \rightarrow 171_8 \rightarrow 115_{15} \rightarrow 54_1 \rightarrow 0$	48/53	17.48,31.55	
	$23_{52} \rightarrow 53_1 \rightarrow 0$	53/53	11.22,20.26	74.18,
	$49_{51} \rightarrow 0$	51/53	9.63,18.68	134.78
	$34_8 \rightarrow 101_8 \rightarrow 100_{27} \rightarrow 23_6 \rightarrow 53_3 \rightarrow 0$	52/53	18.45,31.27	
	$54_{10} \rightarrow 0$	10/14	2.92,6.15	

	$228_7 \rightarrow 207_2 \rightarrow 288_{11} \rightarrow 111_5 \rightarrow 0$	25/25	15.68,29.29	
	$172_5 \rightarrow 174_8 \rightarrow 295_6 \rightarrow 0$	19/23	9.77,19.50	
Victoria	$79_7 \rightarrow 276_5 \rightarrow 47_3 \rightarrow 53_7 \rightarrow (\text{cont.})$			92.57,
	$29_{10} \rightarrow 58_4 \rightarrow 73_5 \rightarrow 0$	41/43	24.82,39.87	159.69
	$200_2 \rightarrow 158_{12} \rightarrow 133_5 \rightarrow 134_9 \rightarrow (\text{cont.})$			
	$137_9 \rightarrow 22_6 \rightarrow 23_1 \rightarrow 0$	44/45	23.20,37.50	
	$312_7 \rightarrow 273_{13} \rightarrow 143_{13} \rightarrow 141_8 \rightarrow 26_1 \rightarrow 0$	42/43	19.10,33.53	
	$68_{43} \rightarrow 22_7 \rightarrow 14_1 \rightarrow 0$	51/53	13.38,23.87	
	$62_{23} \rightarrow 244_{21} \rightarrow 12_7 \rightarrow 0$	51/53	16.52,31.12	
Pembroke	$287_{17} \rightarrow 15_{17} \rightarrow 0$	34/36	10.15,19.64	104.70,
	$22_{39} \rightarrow 21_{10} \rightarrow 14_4 \rightarrow 0$	53/53	10.97,19.47	195.96
	$59_{27} \rightarrow 194_{17} \rightarrow 209_6 \rightarrow 207_3 \rightarrow 0$	53/53	20.10,38.48	
	$88_{21} \rightarrow 84_{25} \rightarrow 321_3 \rightarrow 21_2 \rightarrow 0$	51/53	20.48,37.86	
	$44_{33} \rightarrow 17_9 \rightarrow 0$	42/44	13.10,25.51	
	$268_{73} \rightarrow 256_7 \rightarrow 0$	80/80	18.73,24.78	
	$327_{11} \rightarrow 0$	11/11	19.82,30.11	
	$122_{63} \rightarrow 143_{17} \rightarrow 0$	80/80	19.57,26.25	
Canberra <sup>a</sup>	$155_7 \rightarrow 166_{14} \rightarrow 256_4 \rightarrow 206_{28} \rightarrow 0$	53/53	23.35,31.58	189.27,
	$329_{17} \rightarrow 328_{27} \rightarrow 321_3 \rightarrow 330_1 \rightarrow 0$	48/49	29.95,42.46	259.50
	$214_{50} \rightarrow 190_{25} \rightarrow 0$	75/78	18.12,24.17	
	$85_4 \rightarrow 144_{18} \rightarrow 243_9 \rightarrow (\text{cont.})$			
	$232_{13} \rightarrow 230_{14} \rightarrow 42_{22} \rightarrow 0$	80/80	32.85,44.14	
	$35_9 \rightarrow 3_{18} \rightarrow 23_9 \rightarrow 14_5 \rightarrow 143_{31} \rightarrow 0$	72/78	26.88,36.01	
	$66_{11} \rightarrow 58_{10} \rightarrow 69_{13} \rightarrow 71_1 \rightarrow 0$	35/36	13.37,24.05	
	$159_{32} \rightarrow 52_{21} \rightarrow 0$	53/53	11.60,20.81	
Handaq	$223_{13} \rightarrow 290_6 \rightarrow 357_6 \rightarrow (\text{cont.})$			96.37,
	$93_8 \rightarrow 300_{10} \rightarrow 77_{10} \rightarrow 0$	53/53	18.87,31.19	171.47
	$175_{11} \rightarrow 6_9 \rightarrow 338_{14} \rightarrow 332_{17} \rightarrow 0$	51/53	20.90,37.39	
	$87_{34} \rightarrow 377_{13} \rightarrow 375_3 \rightarrow 0$	50/53	17.08,32.44	
	$275_{18} \rightarrow 36_{16} \rightarrow 40_6 \rightarrow 71_3 \rightarrow 0$	43/44	14.55,25.60	
	$11_{38} \rightarrow 0$	38/44	12.10,27.41	
	$42_{13} \rightarrow 31_{24} \rightarrow 80_7 \rightarrow 0$	44/44	15.90,31.61	
	$328_{21} \rightarrow 175_8 \rightarrow 329_{20} \rightarrow (\text{cont.})$			
Valletta	$80_1 \rightarrow 282_1 \rightarrow 0$	51/53	23.60,41.28	108.43,
	$141_{20} \rightarrow 234_{10} \rightarrow 156_{12} \rightarrow 282_6 \rightarrow 0$	48/53	21.85,40.39	204.88
	$213_9 \rightarrow 58_{17} \rightarrow 45_{16} \rightarrow (\text{cont.})$			
	$60_7 \rightarrow 62_1 \rightarrow 259_2 \rightarrow 0$	52/53	21.73,37.38	
	$114_{32} \rightarrow 259_3 \rightarrow 0$	35/36	13.25,26.82	
	$284_{29} \rightarrow 166_{15} \rightarrow 447_1 \rightarrow 0$	45/53	12.25,21.76	
	$152_{32} \rightarrow 157_{10} \rightarrow 137_{11} \rightarrow 0$	53/53	13.02,23.11	
	$466_3 \rightarrow 251_5 \rightarrow 405_{12} \rightarrow (\text{cont.})$			
Birkirkara	$2_{17} \rightarrow 445_{13} \rightarrow 182_3 \rightarrow 0$	53/53	19.30,33.12	93.72,
	$434_{27} \rightarrow 104_{10} \rightarrow 419_9 \rightarrow 182_4 \rightarrow 0$	50/53	18.53,33.26	166.92
	$393_{13} \rightarrow 120_{21} \rightarrow 235_{18} \rightarrow 0$	52/53	13.37,25.05	
	$452_{14} \rightarrow 293_{18} \rightarrow 357_{15} \rightarrow 447_6 \rightarrow 0$	53/53	17.25,30.61	

Hamrun	$95_{24} \rightarrow 239_5 \rightarrow 466_3 \rightarrow 76_{13} \rightarrow 0$	45/53	13.80,23.92	99.12, 178.11
	$346_2 \rightarrow 343_{17} \rightarrow 244_3 \rightarrow 47_5 \rightarrow 482_{19} \rightarrow 0$	46/53	17.50,30.38	
	$123_{25} \rightarrow 210_{19} \rightarrow 215_8 \rightarrow 0$	52/53	12.23,21.09	
	$261_{16} \rightarrow 386_{13} \rightarrow 365_{16} \rightarrow 0$	45/53	14.82,28.88	
	$436_{21} \rightarrow 437_8 \rightarrow 43_7 \rightarrow 47_1 \rightarrow 0$	37/44	12.27,21.95	
	$497_{20} \rightarrow 165_{15} \rightarrow 174_{10} \rightarrow 365_1 \rightarrow 0$	46/53	17.33,31.73	
	$3_{27} \rightarrow 382_{23} \rightarrow 0$	50/53	11.17,20.18	
Cardiff	$292_{27} \rightarrow 0$	27/27	15.55,38.55	69.63, 154.27
	$456_{47} \rightarrow 0$	47/49	14.70,33.27	
	$120_{45} \rightarrow 362_2 \rightarrow 0$	47/49	18.73,38.40	
	$281_{23} \rightarrow 401_{12} \rightarrow 0$	35/37	20.65,44.05	
M. Keynes	$199_{25} \rightarrow 418_{52} \rightarrow 0$	77/78	17.50,33.29	59.05, 113.25
	$182_{40} \rightarrow 253_{27} \rightarrow 0$	67/70	14.95,29.66	
	$165_{36} \rightarrow 205_{14} \rightarrow 0$	50/51	11.40,22.58	
	$293_{43} \rightarrow 373_{37} \rightarrow 0$	80/80	15.20,27.73	
Bridgend <sup>a</sup>	$567_{72} \rightarrow 0$	72/74	24.18,34.09	178.03, 248.62
	$57_{80} \rightarrow 0$	80/80	23.73,33.02	
	$178_{23} \rightarrow 385_{29} \rightarrow 91_7 \rightarrow 0$	59/61	31.38,43.69	
	$363_8 \rightarrow 391_4 \rightarrow 225_1 \rightarrow 74_8 \rightarrow 128_{19} \rightarrow 0$	40/43	33.02,44.55	
	$57_{80} \rightarrow 0$	80/80	23.73,33.02	
	$246_2 \rightarrow 432_{21} \rightarrow 0$	23/23	21.08,30.97	
Edinburgh-2	$507_{18} \rightarrow 225_3 \rightarrow 389_6 \rightarrow 0$	27/27	20.90,29.29	59.13, 110.76
	$142_{80} \rightarrow 0$	80/80	13.40,24.83	
	$142_{80} \rightarrow 0$	80/80	13.40,24.83	
	$509_{42} \rightarrow 575_{38} \rightarrow 0$	80/80	21.98,44.70	
	$575_{80} \rightarrow 0$	80/80	10.35,16.40	
Edinburgh-1	$673_{53} \rightarrow 664_{25} \rightarrow 380_2 \rightarrow 0$	80/80	21.07,37.79	140.55, 263.19
	$665_{61} \rightarrow 664_{19} \rightarrow 0$	80/80	19.23,38.99	
	$143_{26} \rightarrow 88_{38} \rightarrow 45_{15} \rightarrow 0$	79/80	13.43,21.28	
	$92_{69} \rightarrow 0$	69/70	14.33,29.02	
	$270_{69} \rightarrow 0$	69/70	11.22,20.41	
	$40_{56} \rightarrow 52_{18} \rightarrow 0$	74/74	17.15,32.99	
	$92_{80} \rightarrow 0$	80/80	15.25,29.94	
	$380_{57} \rightarrow 288_{23} \rightarrow 0$	80/80	16.07,30.00	
$409_{12} \rightarrow 6_{57} \rightarrow 0$	69/70	12.80,22.76		
Adelaide	$452_{80} \rightarrow 0$	80/80	16.32,32.89	128.32, 230.54
	$1094_{29} \rightarrow 596_{20} \rightarrow 508_{29} \rightarrow 0$	78/78	18.18,33.59	
	$1086_{37} \rightarrow 634_{39} \rightarrow 0$	76/78	16.50,30.95	
	$1123_{28} \rightarrow 746_{16} \rightarrow 452_{16} \rightarrow 683_1 \rightarrow 0$	61/61	23.92,41.28	
	$1072_{60} \rightarrow 0$	60/61	7.22,10.68	
	$146_{62} \rightarrow 1072_{16} \rightarrow 0$	78/78	16.48,31.31	
	$604_{28} \rightarrow 183_{37} \rightarrow 242_{14} \rightarrow 129_1 \rightarrow 0$	80/80	23.15,39.82	
$1072_{52} \rightarrow 0$	52/53	6.55,10.02		

	1747 <sub>67</sub> → 0	67/70	14.25,18.90	
	435 <sub>80</sub> → 0	80/80	19.60,26.60	
	1395 <sub>28</sub> → 1764 <sub>31</sub> → 603 <sub>20</sub> → 812 <sub>1</sub> → 0	80/80	25.97,35.07	
	726 <sub>77</sub> → 0	77/78	19.28,26.25	
Brisbane <sup>a</sup>	276 <sub>28</sub> → 90 <sub>21</sub> → 211 <sub>30</sub> → 0	79/80	19.28,25.46	205.28,
	646 <sub>35</sub> → 1461 <sub>3</sub> → 1243 <sub>22</sub> → 1249 <sub>4</sub> → 0	64/65	23.92,31.40	275.76
	1024 <sub>45</sub> → 1038 <sub>35</sub> → 0	80/80	18.30,24.09	
	812 <sub>52</sub> → 370 <sub>25</sub> → 0	77/78	17.82,23.46	
	1098 <sub>52</sub> → 1505 <sub>27</sub> → 0	79/80	22.17,29.73	
	1373 <sub>74</sub> → 0	74/78	24.70,34.79	

## 6.5 Monte Carlo Simulations of Correlated SLN Travel Times

This section outlines a systematic approach for generating Monte Carlo simulations of correlated SLN travel times. The methodology, presented in Section 6.5.1, supports the analysis of transportation networks under uncertainty. In each sample, a set of travel time realizations – one for each arc featured in a routing solution – is drawn from the joint distribution, yielding a fixed travel time configuration across the network. Repeating this process across multiple samples produces a distribution of network outcomes. This enables assessment of travel time reliability, evaluation of route-level performance metrics, and insights into the likelihood of timely arrivals under varying network conditions. Section 6.5.2 applies this approach to evaluate all the solutions presented in Chapters 4 to 6 across multiple Monte Carlo samples.

### 6.5.1 Monte Carlo Sampling Methodology

Consider a routing solution given by the set of routes  $\mathcal{R} = \{R_1, R_2, \dots, R_k\}$ . Recall from Chapter 3 that  $V'_1$  denotes the set of visited bus stops in  $\mathcal{R}$ , assumed to have cardinality  $n'$ . Furthermore, let  $E'_1 \subseteq E_1$  denote the set of all  $n'(n' + 1)$  arcs  $(u, v)$  where  $u, v \in V'_1 \cup \{v_0\}$  and  $u \neq v$ . For each of the  $I$  Monte Carlo samples, a set of travel time realizations  $\{t_e^i\}_{e \in E'_1}$ , for  $i \in \{1, \dots, I\}$ , is generated by the following steps.

Step 1: The stops in  $V'_1 \cup \{v_0\}$  are ordered in increasing order in the list `1st`, with indices ranging from 0 to  $n'$ .

Step 2: The arcs  $e = (u, v) \in E'_1$  are ordered using the indexing expression

---

<sup>a</sup>For this problem instance, the reliability level  $\kappa = 0.95$ .

$$n' \text{pos}_{\mathbf{1st}}(u) + \text{pos}_{\mathbf{1st}}(v) - \mathbf{1}_{\{\text{pos}_{\mathbf{1st}}(v) > \text{pos}_{\mathbf{1st}}(u)\}},$$

where  $\text{pos}_{\mathbf{1st}}(u)$  denotes the index of stop  $u$  in the list  $\mathbf{1st}$  (and likewise for stop  $v$ ), and  $\mathbf{1}_{\{x > y\}}$  is the indicator function equal to 1 when  $x > y$ , and 0 otherwise. This ordering ensures that each arc has a unique index in the range from 0 to  $n'(n' + 1) - 1$ , which will be used in the vectors and matrices defined in the following steps.

- Step 3: The mean travel times  $\mathbb{E}[T_e], \forall e \in E'_1$ , are arranged in an  $n'(n' + 1) \times 1$  vector  $\boldsymbol{\tau}$ , following the ordering defined in Step 2.
- Step 4: The parameters  $\mu_e, \forall e \in E'_1$  (as defined in Equation (5.6)), are arranged in an  $n'(n' + 1) \times 1$  vector  $\boldsymbol{\mu}$ , following the ordering defined in Step 2.
- Step 5: An  $n'(n' + 1) \times I$  matrix  $\mathbf{Z}$  of independent standard normal values is generated. Each column corresponds to a distinct Monte Carlo sample and will later be transformed into correlated SLN travel times.
- Step 6: The  $n'(n' + 1) \times n'(n' + 1)$  correlation matrix  $\boldsymbol{\rho}$  for the arc travel times is constructed, as described in Section 6.3.
- Step 7: The correlation matrix  $\boldsymbol{\rho}$  for the SLN distributions is transformed into an  $n'(n' + 1) \times n'(n' + 1)$  correlation matrix  $\boldsymbol{\rho}^{\mathcal{N}}$  for the corresponding underlying normal distributions using the formula:

$$\rho_{e_1 e_2}^{\mathcal{N}} = \frac{1}{\sigma_{e_1} \sigma_{e_2}} \ln \left( \rho_{e_1 e_2} \sqrt{\left( \exp[\sigma_{e_1}^2] - 1 \right) \left( \exp[\sigma_{e_2}^2] - 1 \right) + 1} \right),$$

$\forall e_1, e_2 \in E'_1$ . The reader is referred to the work of Žerovnik et al. (2013) for the derivation of this transformation. If necessary,  $\boldsymbol{\rho}^{\mathcal{N}}$  is adjusted to ensure positive semi-definiteness, as described in Section 6.3.3.

- Step 8: The  $n'(n' + 1) \times n'(n' + 1)$  covariance matrix  $\boldsymbol{\Sigma}$  for the underlying normal distributions with entries  $\Sigma_{e_1 e_2} = \sigma_{e_1} \sigma_{e_2} \rho_{e_1 e_2}^{\mathcal{N}}, \forall e_1, e_2 \in E'_1$ , is computed.
- Step 9: The Cholesky factorization  $\boldsymbol{\Sigma} = \mathbf{L}\mathbf{L}^T$  is performed, where  $\mathbf{L}$  is an  $n'(n' + 1) \times n'(n' + 1)$  lower-triangular matrix known as the Cholesky factor of  $\boldsymbol{\Sigma}$ .
- Step 10: The  $n'(n' + 1) \times I$  matrix  $\mathbf{P} = \mathbf{L}\mathbf{Z}$  is computed. Each column of  $\mathbf{P}$  is a multivariate normal vector with mean  $\mathbf{0}$  and covariance matrix  $\boldsymbol{\Sigma}$ .
- Step 11: The  $n'(n' + 1) \times I$  matrix  $\mathbf{O}$  of random and correlated SLN travel times is generated by

$$O_{ei} = (1 - \alpha)\tau_e + \exp[\mu_e + P_{ei}], \quad \forall e \in E'_1, i \in \{1, 2, \dots, I\}.$$

Note that, first,  $\mu_e$  is added to  $P_{ei}$  to shift the mean of arc  $e \in E'_1$  from 0

to  $\mu_e$ . Then, the exponential function is applied to transform the result into a lognormal value. Finally, the term  $(1 - \alpha)\tau_e$  is added to shift the value and thus obtain an SLN travel time.

An example illustrating the steps described above is now provided for clarity. Continuing with the example presented in Section 6.3.3, consider the Mgarr instance with  $V'_1 = \{11, 31, 39\}$  ( $n' = 3$ ),  $\alpha = 0.2$ ,  $\beta = 0.5$ , and  $I = 5$ . Then,  $\text{lst} = [0, 11, 31, 39]$  and the 12 arcs in  $E'_1$  are ordered as shown in Figure 6.3. For example, the arcs  $(0, 39)$  and  $(31, 11)$  have corresponding indices 2 and 7, respectively. The vectors and matrices generated by Steps 3 to 11 are the following:

$$\boldsymbol{\tau} = (192, 116, 152, 204, 76, 229, 127, 76, 144, 149, 212, 159)^T,$$

$$\boldsymbol{\mu} = (2.66, 2.15, 2.42, 2.72, 1.73, 2.83, 2.24, 1.73, 2.37, 2.40, 2.76, 2.47)^T,$$

$$\mathbf{Z} = \begin{pmatrix} 1.12 & 0.30 & 0.07 & 0.07 & -1.42 \\ 1.52 & -0.29 & -0.13 & -0.17 & -1.76 \\ -0.09 & 1.37 & 1.13 & -0.36 & 1.22 \\ -1.34 & 0.43 & -0.12 & 1.41 & -0.12 \\ 2.01 & 0.23 & 0.60 & 1.63 & 1.59 \\ 0.23 & -0.06 & -0.97 & 0.59 & -0.78 \\ -0.44 & -0.35 & -0.88 & -0.44 & -0.54 \\ -1.32 & -0.11 & 0.91 & 0.82 & 0.23 \\ -1.03 & 0.48 & 1.29 & -0.73 & -1.61 \\ 0.99 & 0.11 & -0.38 & 0.11 & 0.35 \\ -1.73 & 1.66 & 2.30 & -0.47 & 1.26 \\ -1.17 & 1.07 & -0.70 & 0.14 & 0.40 \end{pmatrix},$$

$\boldsymbol{\rho}$  is shown in Figure 6.3,

$$\boldsymbol{\rho}^{\mathcal{N}} = \begin{pmatrix} 1.00 & 0.53 & 0.50 & 0.12 & 0.02 & 0.69 & 0.69 & 0.52 & 0.54 & 0.64 & 0.01 & 0.13 \\ 0.53 & 1.00 & 0.59 & 0.18 & 0.50 & 0.43 & 0.09 & -0.02 & 0.65 & 0.22 & -0.00 & 0.11 \\ 0.50 & 0.59 & 1.00 & 0.61 & 0.12 & 0.03 & 0.11 & -0.02 & 0.60 & 0.29 & 0.16 & -0.02 \\ 0.12 & 0.18 & 0.61 & 1.00 & 0.53 & 0.14 & -0.07 & 0.07 & 0.44 & 0.45 & 0.65 & 0.38 \\ 0.02 & 0.50 & 0.12 & 0.53 & 1.00 & 0.48 & 0.01 & 0.24 & 0.27 & 0.11 & 0.71 & 0.66 \\ 0.69 & 0.43 & 0.03 & 0.14 & 0.48 & 1.00 & 0.40 & 0.69 & 0.57 & 0.63 & 0.24 & 0.67 \\ 0.69 & 0.09 & 0.11 & -0.07 & 0.01 & 0.40 & 1.00 & 0.66 & -0.20 & 0.16 & 0.34 & 0.06 \\ 0.52 & -0.02 & -0.02 & 0.07 & 0.24 & 0.69 & 0.66 & 1.00 & 0.08 & 0.17 & 0.52 & 0.72 \\ 0.54 & 0.65 & 0.60 & 0.44 & 0.27 & 0.57 & -0.20 & 0.08 & 1.00 & 0.67 & -0.11 & 0.35 \\ 0.64 & 0.22 & 0.29 & 0.45 & 0.11 & 0.63 & 0.16 & 0.17 & 0.67 & 1.00 & -0.02 & 0.17 \\ 0.01 & -0.00 & 0.16 & 0.65 & 0.71 & 0.24 & 0.34 & 0.52 & -0.11 & -0.02 & 1.00 & 0.64 \\ 0.13 & 0.11 & -0.02 & 0.38 & 0.66 & 0.67 & 0.06 & 0.72 & 0.35 & 0.17 & 0.64 & 1.00 \end{pmatrix},$$

$$\Sigma = \begin{pmatrix} 1.98 & 1.05 & 0.99 & 0.24 & 0.05 & 1.36 & 1.36 & 1.02 & 1.08 & 1.27 & 0.01 & 0.26 \\ 1.05 & 1.98 & 1.17 & 0.36 & 0.99 & 0.85 & 0.19 & -0.03 & 1.29 & 0.44 & -0.00 & 0.23 \\ 0.99 & 1.17 & 1.98 & 1.21 & 0.24 & 0.06 & 0.22 & -0.04 & 1.19 & 0.57 & 0.32 & -0.04 \\ 0.24 & 0.36 & 1.21 & 1.98 & 1.06 & 0.29 & -0.13 & 0.13 & 0.87 & 0.90 & 1.29 & 0.75 \\ 0.05 & 0.99 & 0.24 & 1.06 & 1.98 & 0.95 & 0.01 & 0.47 & 0.53 & 0.22 & 1.41 & 1.30 \\ 1.36 & 0.85 & 0.06 & 0.29 & 0.95 & 1.98 & 0.79 & 1.36 & 1.12 & 1.25 & 0.49 & 1.34 \\ 1.36 & 0.19 & 0.22 & -0.13 & 0.01 & 0.79 & 1.98 & 1.31 & -0.40 & 0.32 & 0.67 & 0.11 \\ 1.02 & -0.03 & -0.04 & 0.13 & 0.47 & 1.36 & 1.31 & 1.98 & 0.16 & 0.33 & 1.04 & 1.43 \\ 1.08 & 1.29 & 1.19 & 0.87 & 0.53 & 1.12 & -0.40 & 0.16 & 1.98 & 1.33 & -0.22 & 0.70 \\ 1.27 & 0.44 & 0.57 & 0.90 & 0.22 & 1.25 & 0.32 & 0.33 & 1.33 & 1.98 & -0.03 & 0.34 \\ 0.01 & -0.00 & 0.32 & 1.29 & 1.41 & 0.49 & 0.67 & 1.04 & -0.22 & -0.03 & 1.98 & 1.26 \\ 0.26 & 0.23 & -0.04 & 0.75 & 1.30 & 1.34 & 0.11 & 1.43 & 0.70 & 0.34 & 1.26 & 1.98 \end{pmatrix},$$

$$L = \begin{pmatrix} 1.41 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.75 & 1.19 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.71 & 0.54 & 1.09 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.17 & 0.19 & 0.91 & 1.05 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.03 & 0.81 & -0.20 & 1.03 & 0.49 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.97 & 0.10 & -0.62 & 0.64 & 0.12 & 0.48 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.96 & -0.45 & -0.20 & -0.03 & 0.68 & -0.59 & \approx 0 & 0 & 0 & 0 & 0 & 0 \\ 0.73 & -0.48 & -0.27 & 0.33 & 0.91 & 0.46 & \approx 0 & \approx 0 & 0 & 0 & 0 & 0 \\ 0.77 & 0.60 & 0.30 & 0.33 & -0.53 & 0.75 & \approx 0 & \approx 0 & \approx 0 & 0 & 0 & 0 \\ 0.90 & -0.19 & 0.04 & 0.72 & -0.78 & 0.13 & \approx 0 & \approx 0 & \approx 0 & \approx 0 & 0 & 0 \\ 0.01 & -0.01 & 0.29 & 0.98 & 0.95 & -0.18 & \approx 0 & 0 \\ 0.18 & 0.08 & -0.19 & 0.84 & 0.67 & 0.86 & \approx 0 \end{pmatrix},$$

$$P = \begin{pmatrix} 1.58 & 0.43 & 0.10 & 0.10 & -2.00 \\ 2.65 & -0.12 & -0.11 & -0.15 & -3.16 \\ 1.51 & 1.55 & 1.21 & -0.43 & -0.62 \\ -0.99 & 1.68 & 0.88 & 1.13 & 0.39 \\ 0.88 & 0.06 & -0.16 & 2.18 & -1.06 \\ 0.80 & -0.32 & -1.11 & 1.66 & -2.58 \\ 1.68 & 0.33 & 0.89 & 0.94 & 0.72 \\ 1.60 & 0.31 & -0.12 & 2.45 & 0.54 \\ 0.41 & 0.44 & -0.77 & -0.10 & -3.25 \\ -1.78 & 0.50 & -0.56 & -0.08 & -2.32 \\ 0.52 & 1.05 & 0.95 & 2.73 & 1.88 \\ 0.76 & 0.23 & -0.74 & 2.87 & -0.33 \end{pmatrix}, \text{ and}$$

$$\mathbf{O} = \begin{pmatrix} 222.86 & 175.44 & 169.36 & 169.41 & 155.53 \\ 215.00 & 100.43 & 100.55 & 100.20 & 93.16 \\ 172.91 & 174.82 & 159.39 & 128.92 & 127.69 \\ 168.81 & 244.58 & 199.59 & 210.26 & 185.60 \\ 74.46 & 66.79 & 65.62 & 110.71 & 62.76 \\ 221.14 & 195.57 & 188.78 & 272.48 & 184.49 \\ 152.21 & 114.74 & 124.50 & 125.71 & 120.97 \\ 88.80 & 68.53 & 65.80 & 126.36 & 70.48 \\ 131.28 & 131.78 & 120.13 & 124.89 & 115.62 \\ 121.07 & 137.46 & 125.55 & 129.40 & 120.29 \\ 196.04 & 214.50 & 210.20 & 410.97 & 272.86 \\ 152.49 & 142.11 & 132.83 & 334.83 & 135.69 \end{pmatrix}.$$

The final matrix  $\mathbf{O}$  presents the travel times associated with each arc across five distinct realizations. Each row corresponds to a specific arc, and each column represents one of the five travel time samples obtained through Monte Carlo simulation. The samples were generated based on the mean travel times specified in the vector  $\boldsymbol{\tau}$ . Thus, the  $i^{\text{th}}$  element of  $\boldsymbol{\tau}$  represents the expected travel time of arc  $i$ , around which the stochastic samples in  $\mathbf{O}$  were generated. An examination of  $\mathbf{O}$  reveals the variability and skewness inherent in the shifted lognormal travel times. Across the five realizations, some arcs (e.g., arcs 0 and 6) exhibit relatively stable travel times close to their expected values, while others (e.g., arcs 10 and 11) show pronounced deviations, reflecting the stochastic nature of the Monte Carlo samples. A positively skewed distribution is evident in arcs with occasional high travel times, which capture rare but significant delays typical of real transportation networks. This variability, combined with the imposed correlation structure, provides a realistic representation of network travel times and enables subsequent analyses to assess route feasibility under uncertain conditions.

### 6.5.2 Performance Evaluation of Routing Solutions

The routing solutions presented in Chapters 4 to 6 were evaluated using the Monte Carlo sampling methodology described above, with  $I = 1000$  samples. The source code used for this evaluation can be downloaded from (Sciortino, 2025d). Table 6.7 reports the results, showing the average number FS of Monte Carlo samples for which the 25 solutions of each scenario (deterministic, independent, and correlated) achieved journey times not exceeding  $m_t = 45$  minutes. For each scenario, the table also provides the average total journey time (TJT, in minutes), calculated only over the samples for which all routes were feasible.

Table 6.7: Monte Carlo results for  $I = 1,000$  samples, averaged across 25 solutions for parameters  $\alpha = 0.2, \beta = 0.5$  and presented as mean  $\pm$  standard deviation. The column FS indicates the number of samples out of 1,000 for which the solution is feasible (all routes with journey times not exceeding 45 minutes). TJT stands for Total Journey Time (minutes), averaged across the feasible samples.

Location	Deterministic Scenario			Correlated Scenario			Independent Scenario		
	$k$	FS	TJT	$k$	FS	TJT	$k$	FS	TJT
Mgarr	4	996.92 $\pm$ 1.63	34.81 $\pm$ 0.16	4	997.00 $\pm$ 1.91	35.06 $\pm$ 0.28	4.00	996.92 $\pm$ 1.63	35.17 $\pm$ 0.27
Mellicha	4	996.28 $\pm$ 1.54	38.34 $\pm$ 0.16	4	997.32 $\pm$ 1.80	38.43 $\pm$ 0.19	4.00	997.20 $\pm$ 1.50	38.40 $\pm$ 0.29
Porthcawl	1	988.36 $\pm$ 3.58	21.00 $\pm$ 0.48	2	996.64 $\pm$ 1.96	17.82 $\pm$ 0.11	2.00	996.88 $\pm$ 1.72	17.84 $\pm$ 0.12
Qrendi	5	995.24 $\pm$ 2.22	50.12 $\pm$ 0.22	5	996.20 $\pm$ 1.83	50.77 $\pm$ 0.90	5.00	995.56 $\pm$ 1.73	51.48 $\pm$ 1.31
Suffolk	3	883.16 $\pm$ 16.60	91.07 $\pm$ 0.92	4	954.32 $\pm$ 6.15	93.02 $\pm$ 1.08	4.00	952.32 $\pm$ 6.24	91.82 $\pm$ 0.87
Senglea	6	991.40 $\pm$ 3.25	48.43 $\pm$ 1.15	6	995.56 $\pm$ 2.29	49.34 $\pm$ 1.39	6.00	994.60 $\pm$ 2.42	48.43 $\pm$ 1.15
Victoria	4	986.40 $\pm$ 3.67	72.29 $\pm$ 0.57	5	991.68 $\pm$ 2.29	71.56 $\pm$ 0.30	4.32	989.92 $\pm$ 3.90	72.67 $\pm$ 1.67
Pembroke	7	988.28 $\pm$ 3.82	70.34 $\pm$ 0.81	7	991.20 $\pm$ 2.45	73.25 $\pm$ 2.00	7.00	991.24 $\pm$ 2.85	72.58 $\pm$ 1.26
Canberra	7	913.16 $\pm$ 9.73	127.62 $\pm$ 0.65	8	932.04 $\pm$ 9.33	138.82 $\pm$ 1.13	8.00	931.12 $\pm$ 7.57	138.50 $\pm$ 1.12
Handaq	6	991.96 $\pm$ 2.72	67.40 $\pm$ 0.78	6	992.64 $\pm$ 3.28	68.63 $\pm$ 1.28	6.00	992.48 $\pm$ 2.89	68.39 $\pm$ 1.07
Valletta	6	978.88 $\pm$ 4.48	76.87 $\pm$ 0.81	6	986.72 $\pm$ 3.98	81.14 $\pm$ 1.99	6.00	986.20 $\pm$ 3.97	81.87 $\pm$ 2.00
Birkirkara	6	989.44 $\pm$ 4.08	66.42 $\pm$ 2.66	6	994.60 $\pm$ 2.90	66.10 $\pm$ 2.13	6.00	993.00 $\pm$ 2.86	65.88 $\pm$ 2.46
Hamrun	7	992.88 $\pm$ 2.76	66.81 $\pm$ 1.43	7	995.32 $\pm$ 2.30	67.78 $\pm$ 1.29	7.00	994.84 $\pm$ 2.51	67.24 $\pm$ 1.31
Cardiff	2	927.20 $\pm$ 7.20	50.20 $\pm$ 0.17	4	983.00 $\pm$ 3.00	53.42 $\pm$ 0.60	4.00	981.36 $\pm$ 4.07	52.31 $\pm$ 0.30
M. Keynes	4	994.68 $\pm$ 2.36	34.89 $\pm$ 0.61	4	995.08 $\pm$ 1.96	34.59 $\pm$ 0.62	4.00	995.72 $\pm$ 1.79	34.31 $\pm$ 0.66
Bridgend	6	863.16 $\pm$ 24.24	129.35 $\pm$ 0.55	7	926.60 $\pm$ 6.51	138.46 $\pm$ 0.83	7.00	923.20 $\pm$ 7.74	138.29 $\pm$ 0.97
Edinburgh-2	4	993.52 $\pm$ 1.66	30.79 $\pm$ 0.26	4	993.52 $\pm$ 1.66	30.79 $\pm$ 0.26	4.00	994.40 $\pm$ 2.04	31.56 $\pm$ 0.21
Edinburgh-1	9	987.64 $\pm$ 3.52	82.14 $\pm$ 1.60	9	989.40 $\pm$ 2.38	84.40 $\pm$ 2.62	9.00	989.92 $\pm$ 3.05	83.79 $\pm$ 2.48
Adelaide	8	984.00 $\pm$ 4.79	76.43 $\pm$ 1.31	8	987.20 $\pm$ 4.26	78.66 $\pm$ 1.94	8.00	986.44 $\pm$ 3.40	77.35 $\pm$ 1.59
Brisbane	10	949.92 $\pm$ 9.84	141.69 $\pm$ 4.13	10	953.92 $\pm$ 7.55	140.78 $\pm$ 3.82	10.00	955.00 $\pm$ 9.31	141.58 $\pm$ 3.55

Across all the locations, the correlated and independent scenarios yield higher FS values compared with the deterministic baseline, demonstrating that the incorporation of stochastic variability in travel times enhances the robustness of solutions. This improvement is particularly pronounced for locations that are more sensitive to travel time variability, such as Suffolk, Bridgend, and Cardiff, where FS increases by tens of samples relative to the deterministic scenario. While the independent scenario also improves FS relative to the deterministic baseline, it slightly underperforms the correlated scenario in these sensitive networks, reflecting its inability to capture interdependencies among route travel times. In addition to higher FS, the correlated scenario exhibits slightly greater stability across Monte Carlo samples, with marginally lower standard deviations in FS compared with the independent scenario. This indicates that explicitly modelling correlation among travel times produces more consistent and reliable solutions. We note that in a few cases, the FS values for the correlated scenario are slightly lower than the expected proportion  $\kappa$  due to Step 7 of the sampling methodology, in which correlations for the SLN distributions are transformed to the corresponding underlying normal distributions, potentially causing minor distortions.

The TJT results in Table 6.7 show consistent patterns across locations and scenarios. For locations with shorter routes, such as M̄garr, Mellieħa, Porthcawl, Milton Keynes, and Edinburgh-2, TJT remains nearly identical under deterministic, correlated, and independent scenarios, with differences typically less than one minute. For locations with medium journey times, such as Qrendi, Victoria, Pembroke, and Valletta, TJT under correlated and independent scenarios is slightly higher than the deterministic baseline, reflecting the effect of stochastic variability, but standard deviations remain small, indicating stable routing performance. For locations with longer or more dispersed routes, such as Canberra and Bridgend, the increases in TJT under stochastic scenarios are more pronounced, with differences of up to 11 minutes relative to the deterministic case. Comparing the correlated and independent scenarios, the TJT values are generally similar, though the correlated scenario occasionally produces slightly higher TJT, consistent with accounting for interdependencies among route travel times.

Overall, these results demonstrate that the incorporation of variability provides slightly more conservative but realistic estimates of journey times, enhancing the reliability of solutions without substantially compromising efficiency. Among the approaches, the correlated scenario seems to offer the best balance between feasibility and realism. By capturing the joint variability among route travel times during solution generation, it reduces the likelihood of extreme combinations of travel times that can render solutions infeasible, producing higher

FS values in sensitive locations. Simultaneously, it yields journey times that better reflect real-world conditions, ensuring that solutions are both feasible and practically reliable.

To further assess the robustness of the routing solutions in Chapters 4 to 6, we evaluated them using Monte Carlo samples based on three alternative distributions identified as representative of travel times in Section 5.3: the shifted gamma, generalized Pareto, and Burr distributions. For each arc, the parameters of these alternative distributions were calibrated so that their mean, variance, and skewness matched those of the corresponding SLN travel time distribution, ensuring statistical consistency across distributional assumptions. The target moments were first computed, with the mean adjusted to account for the shift governed by parameter  $\alpha$ . In the shifted gamma case, closed-form relationships between the distribution's parameters and moments were inverted analytically to recover the required parameters. For the generalized Pareto and Burr distributions, which do not admit closed-form formulae linking parameters to moments, parameters were estimated numerically using the Levenberg–Marquardt algorithm (Levenberg, 1944, Marquardt, 1963). This is an iterative procedure in which the residuals between the theoretical and target moments are evaluated, and the parameters are updated to minimize these residuals. The optimization stops when the changes in the objective function or in the parameters fall below specified tolerances, when the gradient of the objective function is sufficiently small, or when a maximum number of iterations is reached.

In addition, the same correlation structure as used previously was preserved throughout the Monte Carlo procedure using a Gaussian copula. Correlated standard normal samples were first generated according to the correlation matrix  $\rho^N$  in Step 7. These samples were then converted into uniform random variables on  $[0,1]$  by applying the standard normal CDF. Finally, the uniform samples were transformed into arc travel time realizations by applying the inverse CDF of the fitted alternative distribution for each arc. In this way, only the marginal distributions were altered, while the dependence structure between arcs remained unchanged. This approach enables a controlled evaluation of the sensitivity of routing solutions to deviations in marginal distributional form (such as heavier tails), while maintaining the same correlation structure.

The results of this sensitivity analysis are presented in Table 6.8. Overall, the majority of instances achieve feasibility levels well above 75% across all three scenarios, including the deterministic one. This indicates that most routing structures remain stable even when uncertainty is not explicitly incorporated during optimization. Locations such as Mgarr, Porthcawl, and Edinburgh-2 exhibit consistently high feasibility rates across all alternative distributions and modelling

Table 6.8: Monte Carlo results for  $I = 1,000$  samples of the shifted gamma, generalized Pareto, and Burr distributions, averaged across 25 solutions and presented as mean  $\pm$  standard deviation. The values indicate the number of samples out of 1,000 for which the solution is feasible (all routes with journey times not exceeding 45 minutes).

Location	Deterministic Scenario			Correlated Scenario			Independent Scenario		
	Gamma	Gen. Pareto	Burr	Gamma	Gen. Pareto	Burr	Gamma	Gen. Pareto	Burr
Mgarr	913.6 $\pm$ 9.0	948.3 $\pm$ 6.6	910.6 $\pm$ 9.4	923.8 $\pm$ 8.1	955.3 $\pm$ 5.2	922.5 $\pm$ 8.0	925.3 $\pm$ 5.8	955.4 $\pm$ 5.1	922.0 $\pm$ 9.7
Mellieha	903.8 $\pm$ 10.6	939.8 $\pm$ 6.5	893.0 $\pm$ 8.6	907.3 $\pm$ 10.5	947.3 $\pm$ 7.7	902.2 $\pm$ 10.9	909.1 $\pm$ 7.4	945.0 $\pm$ 7.3	901.0 $\pm$ 11.9
Porthcawl	924.6 $\pm$ 9.3	917.0 $\pm$ 11.5	834.0 $\pm$ 12.9	951.0 $\pm$ 5.3	969.0 $\pm$ 2.7	945.0 $\pm$ 6.4	951.0 $\pm$ 5.3	969.1 $\pm$ 3.4	944.4 $\pm$ 6.5
Qrendi	882.1 $\pm$ 11.6	929.1 $\pm$ 9.3	868.9 $\pm$ 13.3	879.7 $\pm$ 10.8	929.0 $\pm$ 10.7	871.8 $\pm$ 9.3	875.0 $\pm$ 13.7	924.8 $\pm$ 11.1	869.2 $\pm$ 10.5
Suffolk	575.7 $\pm$ 19.3	493.2 $\pm$ 17.6	154.2 $\pm$ 16.9	660.6 $\pm$ 14.9	674.6 $\pm$ 17.8	394.9 $\pm$ 22.7	656.2 $\pm$ 12.3	666.8 $\pm$ 12.1	389.1 $\pm$ 13.4
Senglea	890.3 $\pm$ 9.8	923.7 $\pm$ 8.5	870.8 $\pm$ 10.1	894.4 $\pm$ 14.1	934.1 $\pm$ 9.0	884.4 $\pm$ 16.2	894.0 $\pm$ 12.1	930.7 $\pm$ 10.6	879.9 $\pm$ 13.3
Victoria	740.7 $\pm$ 15.1	809.7 $\pm$ 11.2	619.5 $\pm$ 16.5	776.6 $\pm$ 10.3	864.2 $\pm$ 8.7	731.1 $\pm$ 13.4	753.0 $\pm$ 20.0	832.4 $\pm$ 21.6	649.1 $\pm$ 51.8
Pembroke	831.0 $\pm$ 15.3	878.5 $\pm$ 12.2	791.3 $\pm$ 13.6	822.7 $\pm$ 13.7	883.9 $\pm$ 8.7	800.2 $\pm$ 13.2	825.1 $\pm$ 9.9	884.1 $\pm$ 10.2	797.8 $\pm$ 13.7
Camberra	634.2 $\pm$ 20.3	602.7 $\pm$ 19.4	325.2 $\pm$ 13.6	654.7 $\pm$ 19.6	645.8 $\pm$ 20.0	418.7 $\pm$ 16.2	656.0 $\pm$ 22.1	643.8 $\pm$ 22.7	416.3 $\pm$ 17.3
Handaq	820.1 $\pm$ 12.7	883.6 $\pm$ 9.6	783.7 $\pm$ 19.8	819.6 $\pm$ 13.7	887.5 $\pm$ 10.7	797.2 $\pm$ 18.5	820.4 $\pm$ 14.7	884.0 $\pm$ 16.8	791.6 $\pm$ 27.8
Valletta	805.3 $\pm$ 12.7	835.4 $\pm$ 13.8	716.2 $\pm$ 16.1	798.5 $\pm$ 19.6	853.2 $\pm$ 17.6	728.6 $\pm$ 18.8	796.3 $\pm$ 12.3	844.5 $\pm$ 12.4	724.1 $\pm$ 22.5
Birkirkara	824.2 $\pm$ 13.0	880.6 $\pm$ 15.2	784.0 $\pm$ 23.1	835.4 $\pm$ 13.3	900.9 $\pm$ 10.3	819.5 $\pm$ 24.3	831.4 $\pm$ 17.4	892.5 $\pm$ 15.2	809.6 $\pm$ 22.1
Hamrun	837.2 $\pm$ 10.9	895.7 $\pm$ 8.1	824.2 $\pm$ 13.9	847.5 $\pm$ 15.2	911.8 $\pm$ 10.9	842.7 $\pm$ 16.8	845.4 $\pm$ 14.7	907.7 $\pm$ 10.0	845.6 $\pm$ 17.5
Cardiff	821.8 $\pm$ 13.5	733.7 $\pm$ 12.7	502.5 $\pm$ 16.0	856.6 $\pm$ 12.0	871.5 $\pm$ 11.8	776.4 $\pm$ 16.6	862.5 $\pm$ 11.0	869.2 $\pm$ 10.2	781.8 $\pm$ 11.6
M. Keynes	915.0 $\pm$ 9.1	941.6 $\pm$ 7.8	900.2 $\pm$ 9.6	915.0 $\pm$ 10.3	943.2 $\pm$ 9.0	903.4 $\pm$ 11.7	916.7 $\pm$ 10.7	943.4 $\pm$ 9.6	905.6 $\pm$ 10.5
Bridgend	640.3 $\pm$ 23.3	521.4 $\pm$ 41.8	234.7 $\pm$ 50.4	677.5 $\pm$ 13.9	634.4 $\pm$ 15.9	387.9 $\pm$ 22.8	677.0 $\pm$ 17.8	627.8 $\pm$ 16.6	384.4 $\pm$ 20.0
Edinburgh-2	942.2 $\pm$ 7.3	951.2 $\pm$ 7.4	916.1 $\pm$ 10.4	942.2 $\pm$ 7.3	951.2 $\pm$ 7.4	916.1 $\pm$ 10.4	928.6 $\pm$ 6.7	947.6 $\pm$ 4.8	911.7 $\pm$ 7.0
Edinburgh-1	832.4 $\pm$ 14.2	878.6 $\pm$ 13.8	799.4 $\pm$ 21.0	830.8 $\pm$ 15.0	884.4 $\pm$ 14.1	807.0 $\pm$ 20.9	836.6 $\pm$ 17.6	891.1 $\pm$ 13.5	812.6 $\pm$ 20.6
Adelaide	821.3 $\pm$ 12.7	860.4 $\pm$ 10.9	765.2 $\pm$ 18.5	816.3 $\pm$ 9.9	868.1 $\pm$ 9.3	773.9 $\pm$ 23.7	814.9 $\pm$ 17.6	863.9 $\pm$ 11.6	763.9 $\pm$ 16.1
Brisbane	683.6 $\pm$ 17.6	699.9 $\pm$ 24.9	524.3 $\pm$ 36.0	680.1 $\pm$ 18.0	713.2 $\pm$ 20.4	534.8 $\pm$ 29.8	673.0 $\pm$ 17.9	703.5 $\pm$ 20.8	526.0 $\pm$ 33.1

assumptions. Nevertheless, noticeable differences arise in several networks where accounting for uncertainty leads to improved performance. In particular, Suffolk, Canberra, Cardiff, and Bridgend show substantially lower feasibility rates under the deterministic scenario compared to both stochastic scenarios. These instances illustrate the limitations of purely deterministic optimization in networks that are more sensitive to travel time variability.

A comparison of the two stochastic settings reveals generally modest differences, although the correlated scenario frequently performs slightly better. This pattern is observed for instances such as Suffolk (Burr), Victoria (all three distributions), Birkirkara (Generalized Pareto), Edinburgh-2 (Gamma), and Adelaide (Burr), where modelling dependence between travel times results in higher feasibility rates than under the independence assumption. Although the improvements are limited in magnitude, their consistency across multiple instances suggests that incorporating correlation provides an additional layer of robustness.

Overall, the findings indicate that most solutions maintain high feasibility, demonstrating strong robustness to deviations from the SLN assumptions. For a few more challenging networks where feasibility falls below 75%, such as Suffolk, Canberra, and Bridgend, refinement of the solution framework is recommended to enhance robustness under alternative travel time distributions.

## 6.6 Summary

This chapter examined the impact of incorporating spatial travel time correlations into the SBRP framework introduced earlier in the thesis. We began by acknowledging that the topology of real-world road networks inherently induces travel time dependencies due to physical traffic dynamics, such as the conservation of flow. Consequently, the practice of assuming independent travel times, common in many simplified frameworks, introduces bias in estimating route travel time variability. This bias can lead to suboptimal or infeasible routing solutions, particularly under TTR constraints.

To address this, we proposed a systematic method to classify and estimate correlations between pairs of arcs based on spatial proximity, directional flow, and structural overlap. These correlations were integrated into the routing algorithm, with adjustments to ensure positive semi-definiteness of the correlation matrix using the scaled spectral method. Computational results have demonstrated the advantages of this approach, yielding solutions that satisfy chance constraints on on-time arrivals. In contrast, solutions obtained under deterministic or independent assumptions often violate these constraints or underestimate travel time variability. The incorporation of correlations enables more accurate estimations

of route travel time distributions, resulting in a more robust and realistic routing framework under uncertainty, especially under strict service-level constraints.

At the end of this chapter, we also presented a step-by-step Monte Carlo sampling approach for generating SLN travel times with a predefined correlation structure. The implementation of this approach demonstrated that routing solutions must be developed within a framework that explicitly incorporates travel time correlations, as this both increases the likelihood of feasibility across all routes and produces more realistic journey times reflective of real-world conditions. In summary, Chapter 6 underscores that modelling spatial travel time correlations is not merely an incremental enhancement but an important requirement for producing reliable and practically applicable routing solutions. The proposed approach serves as a methodological bridge between empirical studies of traffic variability and the development of optimization algorithms capable of addressing the complexities of real-world transportation networks.



# Chapter 7

## Conclusions and Future Research

This chapter presents a summary of the research conducted, the main conclusions drawn, and potential avenues for future work. Section 7.1 outlines the key considerations in the field of school bus routing. Section 7.2 highlights the significant findings of this thesis, emphasizing its contributions to the field. The work presented in this thesis lays a solid foundation for a range of future research opportunities, discussed in Section 7.3. Section 7.4 then provides an overview of the resources created and used throughout the research, while Section 7.5 offers some closing reflections.

### 7.1 The Problem Investigated

This thesis has investigated the static and single-school variant of the School Bus Routing Problem (S-SBRP), a combinatorial optimization problem within the broader class of Vehicle Routing Problems (VRPs). The research has focused on the first three subproblems identified in the decomposition proposed by Desrosiers et al. (1981): data preparation, bus stop selection, and route generation.

The data preparation subproblem employs a dedicated geographic information system to construct a spatial network incorporating the school, eligible student residences (with associated student counts), candidate bus stop locations, and available bus capacities. Bus stop selection concentrates on determining which candidate stops should be visited and assigning students to them feasibly, subject to maximum walking distance constraints. Route generation then centres on designing bus routes that are safe, punctual, and economically viable, while complying with key logistical constraints such as vehicle capacity and maximum allowable riding time.

In practice, government administrators often take responsibility for designing school bus routes and subsequently issue public tenders, inviting third-party transport operators to bid for the delivery of these predefined services. Once

contracts are awarded, operators are publicly funded and are required to meet service standards concerning punctuality, vehicle maintenance and upkeep, and student well-being. In this context, governments typically evaluate performance using three core criteria: efficiency, effectiveness, and equity.

Efficiency aims to reduce costs relative to service levels, often by minimizing the number of routes and total bus travel time, though this may involve trade-offs such as the need for longer routes with fewer buses. Effectiveness assesses how well the service meets student demand and satisfaction, using indicators such as student riding time, walking distance, and time loss compared to direct travel. Equity addresses fairness in service distribution, seeking balanced bus loads and riding times across students. Together, these three aspects provide a foundation for enhancing overall service quality, which can help boost ridership and encourage a shift from private to public transport. This transition towards more sustainable travel not only alleviates traffic congestion during peak hours but also contributes to reducing greenhouse gas emissions – an increasingly urgent priority in the face of global climate change.

## 7.2 Summary of Findings

This section summarizes the findings and contributions of this thesis in relation to the research aims outlined in Section 1.3.

### 7.2.1 First Research Aim

*Aim: To evaluate a new mixed integer programming formulation for the S-SBRP that accounts for several realistic operational constraints.*

In the existing literature, mathematical formulations that address both the bus stop selection and route generation subproblems simultaneously – while also incorporating multiple realistic features such as maximum walking distance, maximum riding time, heterogeneous fleets, capacity constraints, bus dwell times, and multistops – are scarce. In Chapter 3, we proposed a new mixed integer programming (MIP) model capable of handling all these features and applicable to all types of travel times, regardless of whether they satisfy the triangle inequality.

The proposed model employs a hierarchical objective function that accounts for efficiency, effectiveness, and equity. Efficiency is prioritized by minimizing the number of routes, which reduces capital and driver-related costs and helps mitigate vehicle shortages. It is further enhanced by minimizing route journey times, a criterion that also contributes to effectiveness by reducing student riding

times. Effectiveness is additionally considered by assigning students to their nearest selected bus stops. Equity considerations are incorporated through ride-time balancing, which promotes a more equitable distribution of travel times among students.

A set of twenty real-world problem instances from Malta, the UK, and Australia – half of which were generated specifically for this research using the Bing Maps Routes Application Programming Interface – was used to evaluate the model. These instances vary in size and network structure, with the largest featuring more than 1800 potential bus stops and 750 students. When solving these instances using Gurobi Optimizer in Python, it was immediately evident that the model is unfit for identifying feasible, let alone optimal, solutions within a reasonable time frame when the number of potential bus stops exceeds approximately 200.

This lack of scalability was anticipated, given that we proved the problem to be NP-hard by showing its equivalence to the time-constrained split delivery VRP in Chapter 3. We therefore conclude that the model is suitable only for generating feasible solutions to relatively small or structurally simple instances (for example, a few dozen potential bus stops or clustered student home locations). For larger, more complex instances, heuristic approaches are essential for obtaining high-quality solutions within practical computational time limits.

## 7.2.2 Second Research Aim

*Aim: To evaluate and compare the performance of various bespoke constructive and improvement (local search) heuristics for the S-SBRP.*

### 7.2.2.1 Constructive Heuristics

Chapter 4 described our iterated local search algorithmic framework for the S-SBRP. Within the framework, three different constructive heuristics – Parallel Nearest Neighbour Heuristic (PNNH), Sequential Random Heuristic (SRH), and Regret Insertion Heuristic (RIH) – were compared, with the first and third being original contributions. Recall that PNNH builds routes backward and in parallel using the nearest neighbour heuristic; SRH constructs routes sequentially and randomly; and RIH, assigns stops to routes in non-increasing order of their regret values, extending Pacheco and Marti (2006)’s heuristic to handle multistops.

Computational results show that PNNH performs the worst on average, generating solutions with total journey times significantly longer than those produced by SRH and RIH. This approach relies on greedy, locally optimal decisions by selecting the nearest next stop for the current route without considering the broader

solution structure, often resulting in unnecessarily long routes. Furthermore, the fixed order of route construction limits flexibility, as routes built later must accommodate whichever stops remain. This issue could be alleviated by adopting a more adaptive route selection strategy during construction – such as choosing to extend the route whose nearest next stop is closest – thereby maintaining better global coordination.

The performance of SRH and RIH is similar, with no statistically significant differences in their total journey times. Ultimately, SRH was selected for further use due to its slightly better average performance and because it is fast and easy to implement. This makes SRH practical for generating quick solutions as a starting point for improvement heuristics. In contrast, RIH is more sophisticated and computationally intensive, as it employs a strategic insertion method designed to reduce the likelihood of poor early choices.

### 7.2.2.2 Improvement Heuristics

The framework described in Chapter 4 includes six improvement heuristics aimed at enhancing solution quality. Three of these are intra-route operators, which modify a single route: exchange, which swaps the positions of two stops within the same route; two-opt, which reverses a segment (i.e., a sequence of consecutive stops) of the route; and generalized Or-opt, which relocates a segment to a new position within the same route. The remaining three are inter-route operators, acting on pairs of routes: Or-exchange, which transfers a segment from one route to another; cross-exchange, which swaps segments between two routes; and create multistop, which splits an occurrence of a stop in one route into two separate occurrences across two routes.

Sensitivity analysis was conducted to evaluate the influence of each neighbourhood operator on the overall performance of the algorithm. This involved systematically removing one of the six operators at a time and comparing the algorithm’s performance without that specific operator to its performance with all operators included. Computational results show that the two-opt operator has no observable impact on performance, while the generalized Or-opt has a greater effect than the exchange operator. Specifically, removing the generalized Or-opt operator led to more cases (8 vs. 3 out of 20) in which the total journey time increased on average. All inter-route operators were found to influence the algorithm’s performance to some extent, with cross-exchange being the most impactful and create multistop the least. Notably, removing the cross-exchange operator resulted in the need for an additional route for the Suffolk instance, and caused a statistically significant increase in the median total journey time for half of the instances. On average, the relative increase in the total journey time

reached up to 20.96%.

Although this is not directly linked to the research aim being discussed, it is worth recalling that the algorithm originally followed the steepest descent (SD) strategy, wherein each iteration evaluates all six neighbourhoods and executes the best-improving move among them. While this approach ensures that the most beneficial move is always selected, it can be computationally intensive. To explore a more efficient alternative, the algorithm was modified to use random variable neighbourhood descent (RVND), in which a single neighbourhood is selected at random in each iteration, and the best-improving move within that neighbourhood is applied. This introduces randomness and reduces the search space per iteration, resulting in faster convergence. Computational results show that RVND is both faster and marginally better on average, making it a more efficient alternative in practice.

### 7.2.3 Third Research Aim

*Aim: To investigate the impact of different mechanisms for varying bus stop selection on the quality of the resulting solutions.*

The framework described in Chapter 4 includes a destroy-and-repair operator designed to modify the selection of bus stops and reconstruct the set of routes based on the updated stop choices. Four variants of this operator were developed, denoted PM 1 through PM 4, with all except the second being original contributions. Recall that the initial set of visited bus stops is selected greedily by first including all compulsory stops, followed by the iterative inclusion of non-compulsory stops in non-increasing order of the number of uncovered student addresses they serve.

PM 1 employs a random restart strategy; PM 2 modifies the set of visited stops from the previous iteration; PM 3 modifies the set of visited stops associated with the most recent best feasible solution; and PM 4 randomly selects between PM 2 and PM 3, each with a probability of 0.5. Computational results show that PM 1 requires substantially more computational time as it does not leverage information from previous iterations, leading to longer local search phases due to the need to process entirely new route sets. Moreover, PM 3 and PM 4 were found to produce statistically significantly better results than PM 1 and PM 2 across most instances. Between the two top-performing variants, PM 4 is slightly faster overall as it reuses the previous iteration's stop subset 50% of the time, resulting in fewer changes and faster local search.

### 7.2.4 Fourth Research Aim

*Aim: To explore the effects of incorporating stochastic travel times, as opposed to deterministic ones, within a heuristic algorithmic framework.*

In Chapter 5, we investigated how solutions generated by a framework that treats travel times as deterministic parameters differ from those produced by a framework that models travel times as random variables. To achieve this, travel time along each arc was modelled as a shifted lognormal (SLN) random variable, with the mean equal to the driving time provided in the instance data. The distribution was parametrized by two adjustable factors: the ratio  $\alpha$  of excess travel time to mean travel time, and the coefficient of variation  $\beta$ .

The framework introduced in Chapter 4 was extended to incorporate on-time arrival chance constraints, requiring that each route be completed within a predefined time limit (45 minutes in this study) with a probability at or above a user-defined reliability threshold  $\kappa$ . To our knowledge, this extended chance-constrained formulation is the first of its kind to address VRPs with split deliveries under travel time uncertainty. A notable contribution of the proposed framework is the inclusion of a mechanism to verify whether a specified reliability level  $\kappa$  is attainable for fixed values of the parameters  $\alpha$  and  $\beta$ . This mechanism highlights the interplay between these parameters and reliability, offering insight into the feasibility of meeting reliability targets and emphasizing the potential need to reconsider bus stop placements to improve reliability.

Through testing various combinations of  $\alpha$  and  $\beta$ , we analysed the effects of altering the deviation between mean and minimal travel times and the degree of travel time variability on solutions. Computational results reveal that increases in either  $\alpha$  or  $\beta$  typically lead to greater dispersion in journey times, higher percentile travel times, and, in certain cases, an increased number of routes required to satisfy reliability constraints. However, due to the mathematical properties of the SLN distribution, increasing  $\alpha$  can occasionally reduce percentile travel times when the decrease in variability offsets the slower rise in mean travel time.

In addition, results under the baseline and moderate parameter settings ( $\alpha = 0.2, \beta = 0.5$ ) demonstrate that neglecting travel time variability during the planning phase often leads to infeasible solutions with respect to the on-time arrival constraints across most problem instances. These findings underscore the critical importance of integrating travel time variability and prioritizing reliability-focused routing strategies, rather than relying solely on mean travel times. Consequently, the incorporation of stochastic travel times and reliability considerations should serve as a key driver for future algorithmic developments in VRPs.

### 7.2.5 Fifth Research Aim

*Aim: To assess whether incorporating correlations between travel times leads to better solutions, and to evaluate the additional computational complexity introduced by this consideration.*

The investigation undertaken in Chapter 6 demonstrated that incorporating spatial correlations between arc travel times enhances the realism and reliability of routing solutions. By introducing a classification scheme for arc pairs – encompassing adjacent, nearby, and overlapping categories – the study captured the spatial dependencies that naturally arise in transportation networks. In particular, the treatment of overlapping arcs, an aspect not previously addressed in the literature, provided a novel means of representing shared physical infrastructure and traffic flow interdependencies.

The scaled spectral method was employed to ensure positive semi-definiteness of the resulting correlation matrices, enabling coherent stochastic modelling while preserving numerical stability. Experimental results revealed that, across the benchmark instances, several deterministic and independent solutions became infeasible when correlations were introduced, highlighting that neglecting dependencies tends to underestimate travel time variability. Conversely, solutions generated under correlated assumptions achieved higher feasibility rates and greater stability across Monte Carlo samples, particularly in networks more sensitive to travel time fluctuations.

From a computational standpoint, accounting for correlations introduced additional demands in data processing, memory usage, and runtime. The preprocessing stage for overlap detection, the construction and adjustment of large-scale correlation matrices, and the Cholesky factorization of the corresponding covariance matrices collectively increased computational complexity. Nevertheless, precomputing overlap data allows correlation information to be reused efficiently in future experiments. Multithreading (using eight threads) was applied to the preprocessing and matrix construction stages to accelerate computation. While the additional complexity was certainly non-trivial, it was offset by the resulting gains in realism and solution reliability, confirming that correlation modelling can be operationally viable when supported by appropriate data structures and preprocessing strategies.

Finally, sensitivity analyses using alternative travel time distributions – including shifted gamma, generalized Pareto, and Burr – confirmed the robustness of SLN-based correlated solutions. Feasibility remained high across most locations, even under distributional deviations, indicating that the proposed approach generalizes well beyond the SLN framework. Only a few complex networks exhib-

ited reduced feasibility, highlighting opportunities for refinement through context-specific statistical and correlation modelling. Overall, this research establishes that explicitly representing correlations between travel times yields solutions that are not only more reliable but also better aligned with the stochastic nature of real transportation networks, with only manageable increases in computational complexity.

## 7.3 Future Research

The research presented in this thesis gives rise to several possible extensions. This section outlines suggestions for future work, involving both conceptual improvements and implementation strategies.

### 7.3.1 Cost Function

Future work could consider enriching the cost function with factors beyond solely minimizing the number of routes and total journey times. One important extension would be to incorporate walking distances to bus stops, following Lewis and Smith-Miles (2018), either by minimizing the aggregate walking distance for all students or by penalizing cases where individual walking distances are too long. Furthermore, emissions or fuel consumption could be explicitly included to address environmental sustainability, recognizing that these depend not only on travel time but also on operational factors such as stopping frequency, acceleration, and idling. Another critical consideration is student safety, which could be modelled by assigning penalties based on the safety rating of each road segment. This would guide the routing framework to avoid roads with elevated accident rates, inadequate lighting, heavy traffic, or other hazards, thereby enhancing the overall safety of the student transportation system.

### 7.3.2 Bus Stop Selection Subproblem

More advanced approaches to the bus stop selection subproblem could be explored by moving beyond the current greedy method. Mathematical models could be developed to explicitly optimize objectives such as minimizing the number of stops or total student walking distance, while incorporating constraints such as stop capacities, vehicle capacities, and maximum allowable walking distance. These models may need to balance operational efficiency with student safety, as reducing the number of stops can increase walking distances and raise safety concerns – particularly in areas with limited pedestrian infrastructure. Moreover, tailored strategies may be necessary for different school types, such as general education,

special education, or mixed schools, which present varying constraints and priorities in stop selection. In addition to optimization-based methods, future work could investigate alternative heuristic strategies. For example, heuristics could prioritize stop selection based on student density and safety indicators associated with nearby road segments, enabling faster decision-making while still accounting for accessibility and risk.

### 7.3.3 Constructive and Improvement Heuristics

Future work could explore alternative constructive and improvement heuristics to better diversify the search process and improve solution quality. On the constructive side, a hybrid heuristic could be developed that switches between different insertion strategies depending on the current route load or length – for example, using nearest neighbour insertions when routes are lightly loaded, and switching to regret-based insertions as routes approach capacity. This approach combines the simplicity of building routes based on proximity with the strategic foresight of regret-based choices. Another idea is to adjust regret values dynamically during route construction, such as giving higher priority to stops with many nearby students or those located in unsafe areas to improve both efficiency and safety. Additionally, cluster-based route construction could be explored, where stops are first grouped by geographic criteria, and routes are then built within these clusters before being merged.

On the improvement side, new neighbourhood operators could be exploited. For example, the three-opt move, applied within a single route, removes three edges and reconnects the segments in a different order to create more compact routes. We also suggest exploring compound moves that allow for more flexible redistribution of stops and greater improvements in route structure. For example, multiroute Or-opt or multiroute cross-exchange could be implemented to perform segment relocations or swaps in a linked chain – such as moving a segment from Route A to Route B, then subsequently moving a segment from Route B to Route C. To date, such embedded neighbourhood approaches have not been applied to the SBRP, representing a promising direction for future research. Moreover, operators that temporarily merge multiple routes for reoptimization and then split them to satisfy constraints like maximum riding time could yield more balanced and efficient configurations. This enables the discovery of routing patterns that might be overlooked when routes are considered in isolation.

### 7.3.4 Hybrid Approaches

Another promising avenue for future work involves hybridizing the proposed iterated local search framework with components from other metaheuristics to leverage their complementary strengths. For example, integrating a tabu search mechanism could guide the search away from recently explored solutions, thereby improving diversification during the local search phase. Alternatively, simulated annealing techniques could be applied to the acceptance criterion, allowing occasional uphill moves and improving the ability to escape local optima. Such hybridisations have the potential to enhance solution quality and search effectiveness, particularly on challenging problem instances.

### 7.3.5 Softer Approach to Travel Time Reliability

In the current framework, the hard constraint  $P_{100\kappa}(R) \leq m_t$  treats all lateness violations equally, without distinguishing between routes that slightly exceed the percentile threshold and those that perform substantially worse. A natural extension for future research is the development of a softer formulation that retains the percentile-based objective while providing greater flexibility. One approach is to incorporate penalties for exceeding the percentile threshold directly into the objective function. This allows the optimization to remain feasible even when some routes slightly surpass the threshold, while discouraging large violations. Formally, the objective could be expressed as the sum of route percentiles plus a weighted penalty for any exceedances, ensuring solutions emphasize high-reliability routes without imposing a strict cut-off.

### 7.3.6 Travel Time Distributions

While the SLN distribution has proven effective in modelling travel time data in several empirical studies – particularly due to its ability to capture positive skewness and preserve tractable correlations across links – future research should aim to move beyond the a priori selection of a distribution. Ideally, a data-driven approach would be adopted, whereby travel time data is available at sufficient resolution and scale to empirically evaluate multiple candidate distributions. This would enable the selection of the most appropriate distribution based on goodness-of-fit measures such as the Anderson–Darling test (well-suited for skewed or heavy-tailed data), the Akaike Information Criterion, or the Bayesian Information Criterion.

For example, the Burr distribution has emerged as a compelling candidate, offering flexibility and superior empirical fit in several studies, especially in settings

with heavy-tailed or multimodal travel time characteristics. However, its lack of closed-form expressions – particularly when summing Burr-distributed random variables – limits its use in analytical routing frameworks. Future work could address this by adopting approximation techniques such as surrogate distributions or piecewise approximations. A surrogate distribution mimics key statistical properties (e.g., moments, skewness, and tail behaviour) of a more complex distribution while remaining analytically tractable. A piecewise approximation, by contrast, partitions the domain into intervals and fits simpler functions or distributions within each interval.

Another promising direction involves the use of mixture distributions, which represent travel times as weighted combinations of component distributions, each capturing distinct traffic regimes (e.g., free-flow, congested, etc.). Parameters and weights of such mixtures can be estimated using expectation-maximization algorithms or Bayesian inference applied to high-resolution data. Travel time reliability (TTR) metrics such as means, percentiles, or standard deviations can then be computed as weighted aggregates of the component distributions' statistics. Where closed-form solutions are unavailable, simulation-based techniques (e.g., parametric bootstrapping) may facilitate the use of these more complex distributions through empirical fitting, allowing for a more flexible and realistic representation of traffic dynamics.

It is important to note that the choice of travel time distribution may influence the selection and interpretation of TTR metrics. For instance, metrics based on means and standard deviations assume symmetric distributions and may be misleading when travel times are skewed. In such cases, percentile-based metrics are preferable as they do not rely on symmetry assumptions and better capture the impact of extreme delays. Therefore, any future shift in distributional assumptions should be accompanied by a reassessment of the TTR metric to ensure alignment with the underlying statistical properties.

### 7.3.7 Spatial Correlations

More advanced techniques might also be employed to capture how travel times are correlated across space. When sufficient historical data is available, correlation matrices can be estimated to quantify how travel times covary across the network. However, this approach may not adequately capture complex dependencies such as tail dependence. In the context of travel times, tail dependence refers to the tendency for delays on different road segments to spike simultaneously during rare but severe events, like heavy storms, even if under normal conditions the travel times are only weakly related. Copulas, such as the  $t$ -copula, provide

a more flexible approach to model tail dependence. They work by modelling the marginal distribution of each travel time separately and then linking them through a copula function, which constructs a full joint distribution reflecting both individual travel time characteristics and their complex dependencies.

Beyond copulas, Gaussian processes offer a powerful framework to model spatial correlations by expressing how travel times on different road segments relate based on geographic proximity. The key element of a Gaussian process is the kernel function, which measures similarity between pairs of locations and defines how spatial dependence decreases with distance or other spatial features. Furthermore, machine learning techniques such as graph neural networks and other spatially-aware deep learning models can be trained on large historical travel time datasets to uncover intricate spatial patterns. These models treat the road network as a graph and can learn how traffic conditions propagate across connected segments.

### 7.3.8 Dynamic and Multischool Settings

An interesting direction for future research lies in extending the current S-SBRP framework to accommodate mixed-load strategies, where buses transport students from multiple schools on the same route. This approach is particularly relevant in countries where such mixed loads are permitted and school schedules are sufficiently aligned. In such contexts, mixed-load routing encourages route consolidation, reducing redundant or overlapping routes that would otherwise serve nearby schools individually, which in turn can contribute to decreased operational costs, traffic congestion, and environmental impact.

This study has addressed both the static deterministic and static stochastic versions of the SBRP. In the static stochastic setting, uncertainties such as student absenteeism and travel times are modelled probabilistically, and routes are planned accordingly before execution. However, once these routes are fixed, they remain unchanged during operation, regardless of how actual conditions may differ from the predictions. A more realistic extension is the dynamic SBRP, which can respond to real-time information received shortly before or even during route execution. For example, if student absenteeism information becomes available while a bus is already en route, the route could be updated on the fly by skipping stops for absent students. Within dynamic and stochastic frameworks, a promising strategy is to integrate predictive absenteeism modelling. This involves estimating the likelihood of student absences to proactively design routes that are robust against such uncertainties and then adjust them in response to real-time data. Incorporating these dynamic and predictive capabilities could substantially

enhance the responsiveness and operational efficiency of school transportation systems.

## 7.4 Summary of Available Resources

The following is an overview of the resources created and used throughout the research:

- The ten problem instances, generated specifically for this research with a focus on the Maltese context, along with a detailed description of each instance's format, are available at (Sciortino, 2024d).
- The Python source code for the MIP model in Chapter 3, along with example output files (including console output, MIP solutions, and interactive route visualizations), is available at (Sciortino, 2024c).
- The C++ source code for the deterministic scenario's framework in Chapter 4, along with the SRH-4 results using the RVND implementation and interactive route visualizations, is available at (Sciortino, 2024b).
- The results for the twelve algorithm variants (PNNH-1 to 4, SRH-1 to 4, and RIH-1 to 4), under deterministic travel times and using the SD implementation, are available at (Sciortino, 2024a).
- The C++ source code for the independent scenario's framework in Chapter 5, along with the results for the nine configurations of the parameters  $\alpha$  and  $\beta$  and interactive route visualizations for the configuration  $\alpha = 0.2, \beta = 0.5$ , is available at (Sciortino, 2025c).
- The Python source code for detecting overlapping road segments between distinct arcs in a network is available at (Sciortino, 2025a).
- The C++ source code for the correlated scenario's framework in Chapter 6, along with the results and interactive route visualizations for the configuration  $\alpha = 0.2, \beta = 0.5$ , is available at (Sciortino, 2025b).
- The C++ source code for the Monte Carlo sampling methodology (SLN, shifted gamma, generalized Pareto, and Burr distributions) in Chapter 6 is available at (Sciortino, 2025d).

## 7.5 Closing Remarks

This thesis has undertaken a comprehensive investigation into the design and optimization of school bus routing systems, prioritizing solutions that are both realistic and practically implementable, while rigorously accounting for complex operational constraints and policy considerations. By addressing critical facets such as bus stop selection, route generation, and travel time uncertainty, and

by developing and evaluating a diverse suite of exact and heuristic methodologies, this work contributes both theoretical insights and practical tools to the field. Although numerous opportunities for future research remain, the methods, findings, and resources established here provide a solid foundation for continued exploration and refinement. Ultimately, by striving to improve the efficiency, effectiveness, and equity of school transportation, this line of research can play a meaningful role in fostering safer, greener, and more accessible mobility for future generations worldwide.

# Bibliography

- Aashtiani, H. Z. and Iravani, H. (2002). Application of dwell time functions in transit assignment model. *Transportation Research Record: Journal of the Transportation Research Board*, 1817(1):88–92.
- Abu-Dayya, A. and Beaulieu, N. (1994). Outage probabilities in the presence of correlated lognormal interferers. *IEEE Transactions on Vehicular Technology*, 43(1):164–173.
- Adulyasak, Y. and Jaillet, P. (2016). Models and algorithms for stochastic and robust vehicle routing with deadlines. *Transportation Science*, 50(2):608–626.
- Agra, A., Christiansen, M., Figueiredo, R., Hvattum, L. M., Poss, M., and Requejo, C. (2013). The robust vehicle routing problem with time windows. *Computers & Operations Research*, 40(3):856–866.
- Al-Deek, H. and Emam, E. B. (2006). New methodology for estimating reliability in transportation networks with degraded link capacities. *Journal of Intelligent Transportation Systems*, 10(3):117–129.
- Angel, R. D., Caudle, W. L., Noonan, R., and Whinston, A. (1972). Computer-assisted school bus scheduling. *Management Science*, 18(6):B279–B288.
- Ansari, A., Farrokhvar, L., and Kamali, B. (2021). Integrated student to school assignment and school bus routing problem for special needs students. *Transportation Research Part E: Logistics and Transportation Review*, 152:102416.
- Ansari Esfeh, M., Kattan, L., Lam, W. H., Ansari Esfe, R., and Salari, M. (2020). Compound generalized extreme value distribution for modeling the effects of monthly and seasonal variation on the extreme travel delays for vulnerability analysis of road network. *Transportation Research Part C: Emerging Technologies*, 120:102808.
- Archetti, C., Bouchard, M., and Desaulniers, G. (2011). Enhanced branch and price and cut for vehicle routing with split deliveries and time windows. *Transportation Science*, 45:285–298.
- Archetti, C. and Speranza, M. G. (2012). Vehicle routing problems with split deliveries. *International Transactions in Operational Research*, 19(1-2):3–22.
- Arias-Rojas, J. S., Jiménez, J. F., and Montoya-Torres, J. R. (2012). Solving of school bus routing problem by ant colony optimization. *Revista EIA*, 9(17):193–

208.

- Arroyo, S. and Kornhauser, A. L. (2005). Modeling travel time distributions on a road network. In Deen, T., editor, *Proceedings of the 84th Annual Meeting of the Transportation Research Board*, Washington, DC. Transportation Research Board.
- Asakura, Y. (1999). Reliability measures of an origin and destination pair in a deteriorated road network with variable flows. In Bell, M. G. H., editor, *Transportation Networks: Recent Methodological Advances*, Selected Proceedings of the 4th EURO Transportation Meeting, pages 273–287, Leeds, England. Emerald Publishing.
- Asensio, J. and Matas, A. (2008). Commuters’ valuation of travel time variability. *Transportation Research Part E: Logistics and Transportation Review*, 44(6):1074–1085.
- Asghari, M., Fathollahi-Fard, A. M., Mirzapour Al-e-hashem, S. M. J., and Dulebenets, M. A. (2022). Transformation and linearization techniques in optimization: A state-of-the-art survey. *Mathematics*, 10:283.
- Atkins, A. (2003). VCU study finds cell phones are not the leading cause of distracted driving. *VCU News*. Available at: [https://www.news.vcu.edu/article/vcu\\_study\\_finds\\_cell\\_phones\\_are\\_not\\_the\\_leading\\_cause\\_of\\_distracted](https://www.news.vcu.edu/article/vcu_study_finds_cell_phones_are_not_the_leading_cause_of_distracted) (Accessed: January 23rd, 2025).
- Avdoshin, S. and Beresneva, E. (2019). Constructive heuristics for capacitated vehicle routing problem: a comparative study. *Proceedings of the Institute for System Programming of the RAS*, 31:145–156.
- Babaei, M. and Rajabi-Bahaabadi, M. (2019). School bus routing and scheduling with stochastic time-dependent travel times considering on-time arrival reliability. *Computers & Industrial Engineering*, 138:106125.
- Babin, G., Deneault, S., and Laporte, G. (2007). Improvements to the Or-opt heuristic for the symmetric travelling salesman problem. *Journal of the Operational Research Society*, 58(3):402–407.
- Bakach, I., Campbell, A. M., Ehmke, J. F., and Urban, T. L. (2021). Solving vehicle routing problems with stochastic and correlated travel times and makespan objectives. *EURO Journal on Transportation and Logistics*, 10:100029.
- Banerjee, D. and Smilowitz, K. (2019). Incorporating equity into the school bus scheduling problem. *Transportation Research Part E: Logistics and Transportation Review*, 131:228–246.
- Banik, S., Bachu, A., and Vanajakshi, L. (2021). Travel time reliability. *International Encyclopedia of Transportation*, pages 109–121.
- Barahimi, A. H., Eydi, A., and Aghaie, A. (2022). Urban transportation network reliability calculation considering correlation among the links comprising

- a route. *Scientia Iranica*, 29(3):1742–1754.
- Barakat, R. (1976). Sums of independent lognormally distributed random variables. *Journal of the Optical Society of America*, 66(3):211–216.
- Bates, J., Polak, J., Jones, P., and Cook, A. (2001). The valuation of reliability for personal travel. *Transportation Research Part E: Logistics and Transportation Review*, 37(2):191–229.
- Bauer, D., Tulic, M., and Scherrer, W. (2019). Modelling travel time uncertainty in urban networks based on flating taxi data. *European Transport Research Review*, 11:46.
- Baum, E. B. (1986). Iterated descent: A better algorithm for local search in combinatorial optimization problems. Technical report, Caltech, Pasadena, California.
- Baxter, J. (1981). Local optima avoidance in depot location. *Journal of the Operational Research Society*, 32(9):815–819.
- Beasley, J. (1983). Route first–cluster second methods for vehicle routing. *Omega*, 11(4):403–408.
- Beaulieu, N. C., Abu-Dayya, A. A., and McLane, P. J. (1995). Estimating the distribution of a sum of independent lognormal random variables. *IEEE Transactions on Communications*, 43(12):2869–2873.
- Beaulieu, N. C. and Xie, Q. (2004). An optimal lognormal approximation to lognormal sum distributions. *IEEE Transactions on Vehicular Technology*, 53:479–489.
- Bektaş, T. and Elmastaş, S. (2007). Solving school bus routing problems through integer programming. *Journal of the Operational Research Society*, 58(12):1599–1604.
- Bennett, B. T. and Gazis, D. C. (1972). School bus routing by computer. *Transportation Research*, 6(4):317–325.
- Berggren, F. and Slimane, S. B. (2004). A simple bound on the outage probability with lognormally distributed interferers. *IEEE Communications Letters*, 8(5):271–273.
- Berhan, E., Beshah, B., Kitaw, D., and Abraham, A. (2014). Stochastic vehicle routing problem: A literature survey. *Journal of Information & Knowledge Management*, 13(3):1450022.
- Berry, D. S. and Belmont, D. M. (1951). Distribution of vehicle speeds and travel times. In Neyman, J., editor, *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 589–602, Berkeley, California. University of California Press.
- Bertsimas, D., Delarue, A., and Martin, S. (2019). Optimizing schools’ start time and bus routes. *Proceedings of the National Academy of Sciences*, 116(13):5943–

5948.

- Binart, S., Dejax, P., Gendreau, M., and Semet, F. (2016). A 2-stage method for a field service routing problem with stochastic travel and service times. *Computers & Operations Research*, 65:64–75.
- Bodin, L. D. and Berman, L. (1979). Routing and scheduling of school buses by computer. *Transportation Science*, 13(2):113–129.
- Bögl, M., Doerner, K. F., and Parragh, S. N. (2015). The school bus routing and scheduling problem with transfers. *Networks*, 65(2):180–203.
- Bomboi, F., Buchheim, C., and Prunte, J. (2022). On the stochastic vehicle routing problem with time windows, correlated travel times, and time dependency. *4OR*, 20(2):217–239.
- Bouyahia, Z., Haddad, H., Jabeur, N., and Sidi Moh, A. N. (2018). Optimization of chartered buses routes under uncertainties using probabilistic vehicle routing problem modeling. *Procedia Computer Science*, 130:644–651.
- Bowerman, R. L., Calamai, P. H., and Brent Hall, G. (1994). The spacefilling curve with optimal partitioning heuristic for the vehicle routing problem. *European Journal of Operational Research*, 76(1):128–142.
- Bowerman, R. L., Calamai, P. H., and Brent Hall, G. (1995). A multi-objective optimization approach to urban school bus routing: Formulation and solution method. *Transportation Research Part A: Policy and Practice*, 29(2):107–123.
- Braaten, S., Gjønnnes, O., Hvattum, L. M., and Tirado, G. (2017). Heuristics for the robust vehicle routing problem with time windows. *Expert Systems with Applications*, 77(1):136–147.
- Braca, J., Bramel, J., Posner, B., and Simchi-Levi, D. (1997). A computerized approach to the New York City school bus routing problem. *IIE Transactions*, 29(8):693–702.
- Braekers, K., Ramaekers, K., and Van Nieuwenhuysse, I. (2016). The vehicle routing problem: State of the art classification and review. *Computers & Industrial Engineering*, 99:300–313.
- Bramel, J. and Simchi-Levi, D. (1995). A location based heuristic for general routing problems. *Operations Research*, 43(4):649–660.
- Caceres, H., Batta, R., and He, Q. (2017). School bus routing with stochastic demand and duration constraints. *Transportation Science*, 51(4):1349–1364.
- Caceres, H., Batta, R., and He, Q. (2019). Special need students school bus routing: Consideration for mixed load and heterogeneous fleet. *Socio-Economic Planning Sciences*, 65:10–19.
- Calvete, H. I., Galé, C., and Iranzo, J. A. (2022). Approaching the Pareto front in a biobjective bus route design problem dealing with routing cost and individuals’ walking distance by using a novel evolutionary algorithm. *Mathematics*,

- 10(9):1390.
- Calvete, H. I., Galé, C., Iranzo, J. A., and Toth, P. (2020). A partial allocation local search matheuristic for solving the school bus routing problem with bus stop selection. *Mathematics*, 8(8):1214.
- Calvete, H. I., Galé, C., Iranzo, J. A., and Toth, P. (2023). The school bus routing problem with student choice: a bilevel approach and a simple and effective metaheuristic. *International Transactions in Operational Research*, 30(2):1092–1119.
- Campbell, J. F., North, J. W., and Ellegood, W. A. (2015). Modeling mixed load school bus routing. In Sebastian, H.-J., Kaminsky, P., and Müller, T., editors, *Quantitative Approaches in Logistics and Supply Chain Management*, Proceedings of the 8th Workshop on Logistics and Supply Chain Management, Berkeley, California, October 3rd and 4th, 2013, pages 3–30, Cham, Switzerland. Springer.
- Carabott, S. (2023). Transport operators blame small van ban for early school pick-up times. *Times of Malta*. Available at: <https://timesofmalta.com/article/transport-operators-blame-small-van-ban-early-pickup-times>. 1061877 (Accessed: May 27th, 2024).
- Carrion, C. and Levinson, D. (2012). Value of travel time reliability: A review of current evidence. *Transportation Research Part A: Policy and Practice*, 46(4):720–741.
- Carrion, C. and Levinson, D. (2013). Valuation of travel time reliability from a GPS-based experimental design. *Transportation Research Part C: Emerging Technologies*, 35:305–323.
- Chalkia, E., Salanova Grau, J. M., Bekiaris, E., Ayfandopoulou, G., Ferarini, C., and Mitsakis, E. (2016). Safety bus routing for the transportation of pupils to school. In Yannis, G. and Cohen, S., editors, *Traffic Safety*, Volume 4 of the Research for Innovative Transports Set, chapter 18, pages 283–299. John Wiley & Sons, Ltd, Chichester, West Sussex, United Kingdom.
- Chan, K. S., Lam, W. H., and Tam, M. L. (2009). Real-time estimation of arterial travel times with spatial travel time covariance relationships. *Transportation Research Record: Journal of the Transportation Research Board*, 2121(1):102–109.
- Chapleau, L., Ferland, J.-A., and Rousseau, J.-M. (1985). Clustering for routing in densely populated areas. *European Journal of Operational Research*, 20(1):48–57.
- Chen, A., Kasikitwiwat, P., and Ji, Z. (2003a). Solving the overlapping problem in route choice with paired combinatorial logit model. *Transportation Research*

- Record: Journal of the Transportation Research Board*, 1857(1):65–73.
- Chen, B. Y., Lam, W. H. K., Sumalee, A., and Li, Z.-l. (2012a). Reliable shortest path finding in stochastic networks with spatial correlated link travel times. *International Journal of Geographical Information Science*, 26(2):365–386.
- Chen, B. Y., Shi, C., Zhang, J., Lam, W. H. K., Li, Q., and Xiang, S. (2017a). Most reliable path-finding algorithm for maximizing on-time arrival probability. *Transportmetrica B*, 5(3):248–264.
- Chen, C., Skabardonis, A., and Varaiya, P. (2003b). Travel-time reliability as a measure of service. *Transportation Research Record: Journal of the Transportation Research Board*, 1855:74–79.
- Chen, D., Laval, J., Zheng, Z., and Ahn, S. (2012b). A behavioral car-following model that captures traffic oscillations. *Transportation Research Part B: Methodological*, 46(6):744–761.
- Chen, L., Hà, M. H., Langevin, A., and Gendreau, M. (2014a). Optimizing road network daily maintenance operations with stochastic service and travel times. *Transportation Research Part E: Logistics and Transportation Review*, 64:88–102.
- Chen, M., Yu, G., Chen, P., and Wang, Y. (2017b). A copula-based approach for estimating the travel time reliability of urban arterial. *Transportation Research Part C: Emerging Technologies*, 82:1–23.
- Chen, P., Tong, R., Lu, G., and Wang, Y. (2018a). The  $\alpha$ -reliable path problem in stochastic road networks with link correlations: A moment-matching-based path finding algorithm. *Expert Systems with Applications*, 110:20–32.
- Chen, P., Tong, R., Lu, G., and Wang, Y. (2018b). Exploring travel time distribution and variability patterns using probe vehicle data: Case study in Beijing. *Journal of Advanced Transportation*, 2018(1):3747632.
- Chen, P., Yin, K., and Sun, J. (2014b). Application of finite mixture of regression model with varying mixing probabilities to estimation of urban arterial travel times. *Transportation Research Record*, 2442(1):96–105.
- Chen, P., Zeng, W., Chen, M., Yu, G., and Wang, Y. (2019). Modeling arterial travel time distribution by accounting for link correlations: a copula-based approach. *Journal of Intelligent Transportation Systems*, 23(1):28–40.
- Chen, X., Cheng, Z., and Sun, L. (2024). Bayesian inference for link travel time correlation of a bus route. *Transportmetrica B: Transport Dynamics*, 12(1):2416181.
- Chen, X., Kong, Y., Dang, L., Hou, Y., and Ye, X. (2015). Exact and meta-heuristic approaches for a bi-objective school bus scheduling problem. *PLOS ONE*, 10(7):1–20.
- Chen, X.-P., Kong, Y.-F., Zheng, T.-H., and Zheng, S.-S. (2016). Metaheuristic

- algorithm for split demand school bus routing problem. *Computer Science*, 43(10):234–241.
- Chen, Z. and Fan, W. D. (2020). Analyzing travel time distribution based on different travel time reliability patterns using probe vehicle data. *International Journal of Transportation Science and Technology*, 9(1):64–75.
- Cheng, T., Haworth, J., and Wang, J. (2011). Spatio-temporal autocorrelation of road network data. *Journal of Geographical Systems*, 14(4):389–413.
- Cheok, C. C. T., Khoo, W. C., and Khoo, H. L. (2024). Bus travel time variability modelling using Burr type XII regression: A case study of Klang Valley. *KSCE Journal of Civil Engineering*, 28(9):3998–4009.
- Chepuri, A., Ramakrishnan, J., Arkatkar, S., Joshi, G., and Pulugurtha, S. S. (2018). Examining travel time reliability-based performance indicators for bus routes using GPS-based bus trajectory data in India. *Journal of Transportation Engineering, Part A: Systems*, 144(5):04018012.
- Chu, J., Yan, S., and Huang, H.-J. (2017). A multi-trip split-delivery vehicle routing problem with time windows for inventory replenishment under stochastic travel times. *Networks and Spatial Economics*, 17:41–68.
- Clarke, G. and Wright, J. W. (1964). Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research*, 12(4):568–581.
- Clenshaw, C. W. and Curtis, A. R. (1960). A method for numerical integration on an automatic computer. *Numerische Mathematik*, 2(1):197–205.
- Corberán, A., Fernández, E., Laguna, M., and Martí, R. (2002). Heuristic solutions to the problem of routing school buses with multiple objectives. *The Journal of the Operational Research Society*, 53(4):427–435.
- Cordeau, J.-F., Laporte, G., Savelsbergh, M. W., and Vigo, D. (2007). Vehicle routing. In Barnhart, C. and Laporte, G., editors, *Handbooks in Operations Research and Management Science: Transportation*, volume 14, pages 367–428. North-Holland Publishing Company, Amsterdam, Netherland.
- Crama, Y., Kolen, A. W. J., and Pesch, E. J. (1995). Local search in combinatorial optimization. In Braspenning, P. J., Thuijsman, F., and Weijters, A. J. M. M., editors, *Artificial Neural Networks: An Introduction to ANN Theory and Practice*, pages 157–174. Springer, Berlin, Heidelberg.
- Croes, G. A. (1958). A method for solving traveling-salesman problems. *Operations Research*, 6(6):791–812.
- Dai, Z., Ma, X., and Chen, X. (2018). Bus travel time modelling using GPS probe and smart card data: A probabilistic approach considering link travel time and station dwell time. *Journal of Intelligent Transportation Systems*, 23(2):175–190.
- Dandy, G. C. and McBean, E. A. (1984). Variability of individual travel time

- components. *Journal of Transportation Engineering*, 110(3):340–357.
- Dang, L.-x., Hou, Y.-e., Liu, Q.-s., and Kong, Y.-f. (2019). A hybrid meta-heuristic algorithm for the bi-objective school bus routing problem. *IAENG International Journal of Computer Science*, 46(3):409–416.
- Dantzig, G. B. and Ramser, J. H. (1959). The truck dispatching problem. *Management Science*, 6(1):80–91.
- Desaulniers, G. (2010). Branch-and-price-and-cut for the split-delivery vehicle routing problem with time windows. *Operations Research*, 58(1):179–192.
- Desrosiers, J., Ferland, J. A., Rousseau, J.-M., Lapalme, G., and Chapleau, L. (1981). An overview of a school busing system. In Jaiswal, N., editor, *Scientific Management of Transport Systems*, Revised and edited version of selected papers presented at the International Conference on Transportation held at New Delhi, India, November 26-28, 1980, pages 235–243. North-Holland Publishing Company, Amsterdam, Netherland.
- Díaz-Ramírez, J., Leal-Garza, C. M., and Gómez-Acosta, C. (2022). A smart school routing and scheduling problem for the new normalcy. *Computers & Industrial Engineering*, 168:108101.
- Díaz-Ramírez, J., Rosas, A., García-Colis, B., and Valencia-Torres, P. (2024). The school bus routing problem: A systematic literature review. *Revista Facultad de Ingeniería Universidad de Antioquia*, 114:79–94. Available at: <https://doi.org/10.17533/udea.redin.20240518> (Accessed: April 18th, 2025).
- Dong, J. and Mahmassani, H. S. (2009). Flow breakdown and travel time reliability. *Transportation Research Record*, 2124(1):203–212.
- Dong, W., Li, M., Vo, B., and Vu, H. (2012). Reliability in stochastic time-dependent traffic networks with correlated link travel times. In *15th International IEEE Conference on Intelligent Transportation Systems*, pages 1626–1631, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Dror, M. and Trudeau, P. (1989). Savings by split delivery routing. *Transportation Science*, 23(2):141–145.
- Dror, M. and Trudeau, P. (1990). Split delivery routing. *Naval Research Logistics*, 37(3):383–402.
- Dulac, G., Ferland, J. A., and Forgues, P. A. (1980). School bus routes generator in urban surroundings. *Computers & Operations Research*, 7(3):199–213.
- Dunlop, D. (2009). Central area bus operational review. Technical report, Opus International Consultants Limited, Wellington, New Zealand.
- Durán-Hormazábal, E. and Tirachini, A. (2016). Estimation of travel time variability for cars, buses, metro and door-to-door public transport trips in Santiago, Chile. *Research in Transportation Economics*, 59:26–39.

- Ehmke, J. F., Campbell, A. M., and Urban, T. L. (2015). Ensuring service levels in routing problems with time windows and stochastic travel times. *European Journal of Operational Research*, 240(2):539–550.
- Eksioglu, B., Vural, A. V., and Reisman, A. (2009). The vehicle routing problem: A taxonomic review. *Computers & Industrial Engineering*, 57(4):1472–1483.
- El Esawey, M. and Sayed, T. (2011). Travel time estimation in urban networks using limited probes data. *Canadian Journal of Civil Engineering*, 38(3):305–318.
- El Esawey, M. and Sayed, T. (2012). A framework for neighbour links travel time estimation in an urban network. *Transportation Planning and Technology*, 35(3):281–301.
- Eldrandaly, K. A. and Abdallah, A. F. (2012). A novel GIS-based decision-making framework for the school bus routing problem. *Geo-spatial Information Science*, 15(1):51–59.
- Elefteriadou, L. (2014). Travel time reliability implementation for the freeway SIS. Technical report, University of Florida Transportation Institute, Florida, USA.
- Eliasson, J. (2007). The relationship between travel time variability and road congestion. In Deakin, E., editor, *Proceedings of the 11th World Conference of Transport Research*, California, USA. University of California, Berkeley.
- Ellegood, W. A., Solomon, S., North, J., and Campbell, J. F. (2020). School bus routing problem: Contemporary trends and research directions. *Omega*, 95:102056.
- Elshaer, R. and Awad, H. (2020). A taxonomic review of metaheuristic algorithms for solving the vehicle routing problem and its variants. *Computers & Industrial Engineering*, 140:106242.
- Emam, E. B. and Al-Deek, H. (2006). Using real-life dual-loop detector data to develop new methodology for estimating freeway travel time reliability. *Transportation Research Record*, 1959(1):140–150.
- Ergun, Ö., Orlin, J., and Steele-Feldman, A. (2006). Creating very large scale neighborhoods out of smaller ones by compounding moves. *Journal of Heuristics*, 12(1):115–140.
- Ermagun, A., Chatterjee, S., and Levinson, D. (2017). Using temporal detrending to observe the spatial correlation of traffic. *PLOS ONE*, 12(5):e0176853.
- Euchi, J. and Mraihi, R. (2012). The urban bus routing problem in the Tunisian case by the hybrid artificial ant colony algorithm. *Swarm and Evolutionary Computation*, 2:15–24.
- Eufinger, L., Kurtz, J., Buchheim, C., and Clausen, U. (2020). A robust approach to the capacitated vehicle routing problem with uncertain costs. *INFORMS*

- Journal on Optimization*, 2(2):79–95.
- Fan, Y., Kalaba, R., and Moore, J. (2005). Shortest paths in stochastic networks with correlated link costs. *Computers & Mathematics with Applications*, 49(9):1549–1564.
- Faraj, M. F., Sarubbi, J. F., Silva, C. M., Porto, M. F., and Nunes, N. T. R. (2014). A real geographical application for the school bus routing problem. In *17th International IEEE Conference on Intelligent Transportation Systems*, volume 4, pages 2762–2767, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Feillet, D., Dejax, P., Gendreau, M., and Gueguen, C. (2006). Vehicle routing with time windows and split deliveries. Technical report, Laboratoire d’Informatique d’Avignon, Avignon, France.
- Fenton, L. F. (1960). The sum of log-normal probability distributions in scatter transmission systems. *IRE Transactions on Communications Systems*, 8(1):57–67.
- Feo, T. A. and Resende, M. G. C. (1989). A probabilistic heuristic for a computationally difficult set covering problem. *Operations Research Letters*, 8(2):67–71.
- Feo, T. A. and Resende, M. G. C. (1995). Greedy randomized adaptive search procedures. *Journal of Global Optimization*, 6:109–133.
- Figliozzi, M. A., Kingdon, L., and Wilkitzki, A. (2007). Analysis of freight tours in a congested urban area using disaggregated data: characteristics and data collection challenges. In *Proceedings 2nd Annual National Urban Freight Conference, Long Beach, California. December, 2007*.
- Fischetti, M., Salazar González, J. J., and Toth, P. (1997). A branch-and-cut algorithm for the symmetric generalized traveling salesman problem. *Operations Research*, 45:378–394.
- Fügenschuh, A. (2009). Solving a school bus scheduling problem with integer programming. *European Journal of Operational Research*, 193(3):867–884.
- Fulin, S. and Yueguang, L. (2012). An improved quantum-behaved particle swarm algorithm and its application in school bus problem. In *2012 Third International Conference on Digital Manufacturing & Automation*, pages 198–201, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Gajewski, B. and Rilett, L. (2005). Estimating link travel time correlation: An application of Bayesian smoothing splines. *Journal of Transportation and Statistics*, 7(2):53–70.
- Galdi, M. and Thebpanya, P. (2016). Optimizing school bus stop placement in Howard County, Maryland: A GIS-based heuristic approach. *International Journal of Applied Geospatial Research (IJAGR)*, 7(1):30–44.

- Garey, M. R. and Johnson, D. S. (1975). Complexity results for multiprocessor scheduling under resource constraints. *SIAM Journal on Computing*, 4(4):397–411.
- Garey, M. R. and Johnson, D. S. (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, USA.
- Gavish, B. and Shlifer, E. (1979). An approach for solving a class of transportation scheduling problems. *European Journal of Operational Research*, 3(2):122–134.
- Gendreau, M., Hertz, A., and Laporte, G. (1992). New insertion and postoptimization procedures for the traveling salesman problem. *Operations Research*, 40:1086–1094.
- Gendreau, M., Laporte, G., and Séguin, R. (1996). Stochastic vehicle routing. *European Journal of Operational Research*, 88(1):3–12.
- Gillett, B. E. and Miller, L. R. (1974). A heuristic algorithm for the vehicle-dispatch problem. *Operations Research*, 22(2):340–349.
- Glover, F. (1996). Ejection chains, reference structures and alternating path methods for traveling salesman problems. *Discrete Applied Mathematics*, 65(1):223–253.
- Gómez, A., Mariño, R., Akhavan-Tabatabaei, R., Medaglia, A. L., and Mendoza, J. E. (2016). On modeling stochastic travel and service times in vehicle routing. *Transportation Science*, 50(2):627–641.
- Graziosi, F., Fuciarelli, L., and Santucci, E. (2001). Second order statistics of the sir for cellular mobile networks in the presence of correlated co-channel interferers. In *IEEE VTS 53rd Vehicular Technology Conference, Spring 2001*, volume 4, pages 2499–2503.
- Guan, L. and Wang, W. (2024). Collaborative optimization of company commuter bus stop selection and route programming considering stochastic travel time and demand. *Transportation Research Record: Journal of the Transportation Research Board*, 2678(1):61–76.
- Guenther, R. P. and Sinha, K. C. (1983). Modeling bus delays due to passenger boardings and alightings. *Transportation Research Record: Journal of the Transportation Research Board*, 915:7–13.
- Guessous, Y., Aron, M., Bhourri, N., and Cohen, S. (2014). Estimating travel time distribution under different traffic conditions. *Transportation Research Procedia*, 3:339–348.
- Guo, F., Gu, X., Guo, Z., Dong, Y., and Wallace, S. W. (2020). Understanding the marginal distributions and correlations of link travel speeds in road networks. 10(1):11821.
- Guo, F., Rakha, H., and Park, S. (2010). Multistate model for travel time reliability. *Transportation Research Record*, 2188(1):46–54.

- Guo, S., Zhou, D., Fan, J., Tong, Q., Zhu, T., Lv, W., Daqing, L., and Havlin, S. (2019a). Identifying the most influential roads based on traffic correlation networks. *EPJ Data Science*, 8:28.
- Guo, X., Liu, Y., and Samaranayake, S. (2018). Solving the school bus routing problem at scale via a compressed shareability network. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 1900–1907. Institute of Electrical and Electronics Engineers (IEEE).
- Guo, X. and Samaranayake, S. (2022). Shareability network based decomposition approach for solving large-scale single school routing problems. *Transportation Research Part C: Emerging Technologies*, 140:103691.
- Guo, Z., Wallace, S. W., and Kaut, M. (2019b). Vehicle routing with space- and time-correlated stochastic travel times: Evaluating the objective function. *INFORMS Journal on Computing*, 31(4):654–670.
- Gutierrez, A., Dieulle, L., Labadie, N., and Velasco, N. (2016). A multi population memetic algorithm for the vehicle routing problem with time windows and stochastic travel and service times. *IFAC-PapersOnLine*, 49(12):1204–1209.
- Han, D., Xiang, H., and Xie, Y. (2022). Solving multi depot school bus routing problem based on simulated annealing algorithm. In Sun, N. and Cen, F., editors, *International Conference on Algorithms, Microchips and Network Applications*, volume 12176, page 1217608, Bellingham, Washington, USA. SPIE.
- Han, J., Lee, C., and Park, S. (2014). A robust scenario approach for the vehicle routing problem with uncertain travel times. *Transportation Science*, 48(3):373–390.
- Hao, D., Mao, X., Wei, Y., Sun, L., and Yang, Y. (2024). Review of research on vehicle routing problems. In Mikusova, M., editor, *International Conference on Smart Transportation and City Engineering (STCE 2023)*, volume 13018, page 130180Y, Bellingham, Washington, USA. SPIE.
- Harsha, M. and Mulangi, R. H. (2022). Probability distributions analysis of travel time variability for the public transit system. *International Journal of Transportation Science and Technology*, 11(4):790–803.
- Hcine, M. and Bouallegue, R. (2015). On the approximation of the sum of log-normals by a log skew normal distribution. *International Journal of Computer Networks and Communications*, 7:135–151.
- He, R. R., Liu, H. X., Kornhauser, A. L., and Ran, B. (2006). A study on temporal and spatial variability of travel time. UC Irvine: Center for Traffic Simulation Studies. Available at: <https://escholarship.org/uc/item/6c5635f2> (Accessed: February 4th, 2025).
- Herman, R. and Lam, T. N. (1974). Trip time characteristics of journeys to and from work. In Buckley, D. J., editor, *Transportation and Traffic Theory*,

- Proceedings of the 6th International Symposium on Transportation and Traffic Theory held at the University of New South Wales, Sydney, Australia from 26-28 August 1974, pages 57–86, New York, USA. Elsevier.
- Hollander, Y. (2006). Direct versus indirect models for the effects of unreliability. *Transportation Research Part A: Policy and Practice*, 40:699–711.
- Hou, L. and Zhou, H. (2010). Stochastic vehicle routing problem with uncertain demand and travel time and simultaneous pickups and deliveries. In *2010 Third International Joint Conference on Computational Science and Optimization*, pages 32–35, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Hou, Y., Liu, B., Dang, L., He, W., and Gu, W. (2022a). A local search-based metaheuristic algorithm framework for the school bus routing problem. *Engineering Letters*, 30(1).
- Hou, Y., Zhao, N., Dang, L., , and Liu, B. (2021). A hybrid metaheuristic algorithm for the school bus routing problem with multi-school planning scenarios. *Engineering Letters*, 29(4):1397–1406.
- Hou, Y., Zhao, N., and Dang, L. (2022b). Solving multi-objective school bus routing problem using an improved NSGA-II algorithm. *Engineering Letters*, 30(2):397–405.
- Hu, C., Lu, J., Liu, X., and Zhang, G. (2018). Robust vehicle routing problem with hard time windows under demand and travel time uncertainty. *Computers & Operations Research*, 94:139–153.
- Hülagü, S. and Celikoglu, H. B. (2022). Environment-friendly school bus routing problem with heterogeneous fleet: A large-scale real case. *IEEE Transactions on Intelligent Transportation Systems*, 23(4):3461–3471.
- Huo, L., Yan, G., Fan, B., Wang, H., and Gao, W. (2014). School bus routing problem based on ant colony optimization algorithm. In *2014 IEEE Conference and Expo Transportation Electrification Asia-Pacific (ITEC Asia-Pacific)*, pages 1–5, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Iklassov, Z., Sobirov, I., Solozabal, R., and Takáč, M. (2024). Reinforcement learning for solving stochastic vehicle routing problem. In Yamikoğlu, B. and Buntine, W., editors, *Proceedings of the 15th Asian Conference on Machine Learning (ACML 2023) held on November 11-14, 2023, Istanbul, Turkey*, volume 222 of *Proceedings of Machine Learning Research*, pages 502–517. PMLR.
- Iman, R. L. and Conover, W. J. (1982). A distribution-free approach to inducing rank correlation among input variables. *Communications in Statistics - Simulation and Computation*, 11(3):311–334.
- Isukapati, I. K., List, G. F., Williams, B. M., and Karr, A. F. (2013). Syn-

- thesizing route travel time distributions from segment travel time distributions. *Transportation Research Record: Journal of the Transportation Research Board*, 2396:71–81.
- Jaillet, P., Qi, J., and Sim, M. (2016). Routing optimization under uncertainty. *Operations Research*, 64(1):186–200.
- Ji, Z., Kim, Y. S., and Chen, A. (2011). Multi-objective  $\alpha$ -reliable path finding in stochastic networks with correlated link costs: A simulation-based multi-objective genetic algorithm approach (SMOGA). *Expert Systems with Applications*, 38(3):1515–1528.
- Jiang, L. and Mahmassani, H. S. (2014). City logistics: Freight distribution management with time-dependent travel times and disruptive events. *Transportation Research Record: Journal of the Transportation Research Board*, 2410(1):85–95.
- Johnson, D. and McGeoch, L. (1997). The traveling salesman problem: A case study in local optimization. In Aarts, E. H. L. and Lenstra, J. K., editors, *Local Search in Combinatorial Optimization*, pages 215–310. John Wiley & Sons, Inc, Hoboken, New Jersey, USA.
- Kang, M. H., Kim, S.-K., Felan, J., Choi, H. R., and Cho, M. (2015). Development of a genetic algorithm for the school bus routing problem. *International Journal of Software Engineering and its Applications*, 9(5):107–126.
- Kaparias, I., Bell, M. G. H., and Belzner, H. (2008). A new measure of travel time reliability for in-vehicle navigation systems. *Journal of Intelligent Transportation Systems: Technology, Planning, and Operations*, 12(4):202–211.
- Kaut, M. and Wallace, S. W. (2007). Evaluation of scenario-generation methods for stochastic programming. *Pacific Journal of Optimization*, 3(2):257–271.
- Kazagli, E. and Koutsopoulos, H. N. (2013). Estimation of arterial travel time from automatic number plate recognition data. *Transportation Research Record*, 2391(1):22–31.
- Kek, A., Cheu, R., and Meng, Q. (2008). Distance-constrained capacitated vehicle routing problems with flexible assignment of start and end depots. *Mathematical and Computer Modelling*, 47(1-2):140–152.
- Kenyon, A. S. and Morton, D. P. (2003). Stochastic vehicle routing with random travel times. *Transportation Science*, 37(1):69–82.
- Khoo, W. C., Lim, K. S., and Cheong, H. T. (2021). Travel time reliability modelling with Burr distribution. *Malaysian Journal of Mathematical Sciences*, 15(2):313–322.
- Kieu, L. M., Bhaskar, A., and Chung, E. (2015). Public transport travel time variability definitions and monitoring. *Journal of Transportation Engineering*, 141(1):04014068.

- Kim, B.-I., Kim, S., and Park, J. (2012). A school bus scheduling problem. *European Journal of Operational Research*, 218(2):577–585.
- Kim, J. and Mahmassani, H. S. (2015). Compound gamma representation for modeling travel time variability in a traffic network. *Transportation Research Part B: Methodological*, 80:40–63.
- Kim, T. and Park, B.-J. (2013). Model and algorithm for solving school bus problem. *Journal of Emerging Trends in Computing and Information Sciences*, 4(8):596–600.
- Kinable, J., Spieksma, F., and Vanden Berghe, G. (2014). School bus routing—a column generation approach. *International Transactions in Operational Research*, 21(3):453–478.
- Knoop, V. L., van Zuylen, H. J., and Hoogendoorn, S. P. (2008). Capacity reduction at incidents: Empirical data collected from a helicopter. *Transportation Research Record*, 2071:19–25.
- Knoop, V. L., van Zuylen, H. J., and Hoogendoorn, S. P. (2009). Microscopic traffic behaviour near incidents. In Lam, W. H. K., Wong, S. C., and Lo, H. K., editors, *Transportation and Traffic Theory 2009: Golden Jubilee*, pages 75–97, New York, USA. Springer.
- Köksal Ahmed, E., Li, Z., Veeravalli, B., and Ren, S. (2021). Reinforcement learning-enabled genetic algorithm for school bus scheduling. *Journal of Intelligent Transportation Systems*, 26(3):269–283.
- Kolesar, P. and Walker, W. (1972). *An Algorithm for the Dynamic Relocation of Fire Companies*. RAND Corporation, Santa Monica, CA.
- Komijan, A. R., Ghasemi, P., Khalili-Damghani, K., and Yazdi, F. H. (2021). A new school bus routing problem considering gender separation, special students and mix loading: A genetic algorithm approach. *Journal of Optimization in Industrial Engineering*, 14(31).
- Konstantakopoulos, G. D., Gayialis, S. P., and Kechagias, E. P. (2022). Vehicle routing problem and related algorithms for logistics distribution: a literature review and classification. *Operational Research*, 22(3):2033–2062.
- Kotoula, K., Morfoulaki, M., Ayfantopoulou, G., and Tzenos, P. (2017). Calculating optimal school bus routing and its impact on safety and the environment. *Transportation Research Record: Journal of the Transportation Research Board*, 2647:142–150.
- Kumar, Y. and Jain, S. (2015). School bus routing based on branch and bound approach. In *2015 International Conference on Computer, Communication and Control (IC4)*, pages 1–4, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Labadie, N., Prins, C., and Prodhon, C. (2016). *Metaheuristics for Vehicle Rout-*

- ing Problems*. Wiley-IEEE Press, 1st edition.
- Lam, T. C. and Small, K. A. (2001). The value of time and reliability: measurement from a value pricing experiment. *Transportation Research Part E: Logistics and Transportation Review*, 37(2):231–251.
- Laporte, G. (2009). Fifty years of vehicle routing. *Transportation Science*, 43(4):408–416.
- Laporte, G., Louveaux, F., and Mercure, H. (1992). The vehicle routing problem with stochastic travel times. *Transportation Science*, 26(3):161–170.
- Laporte, G., Nobert, Y., and Desrochers, M. (1985). Optimal routing under capacity and distance restrictions. *Operations Research*, 33(5):1050–1073.
- Laporte, G., Nobert, Y., and Taillefer, S. (1988). Solving a family of multi-depot vehicle routing and location-routing problems. *Transportation Science*, 22(3):161–172.
- Laporte, G. and Semet, F. (2002). Classical heuristics for the capacitated VRP. In Toth, P. and Vigo, D., editors, *The Vehicle Routing Problem*, pages 109–128, Philadelphia, USA. Society for Industrial and Applied Mathematics (SIAM).
- Larsen, J. (1999). *Parallelization of the Vehicle Routing Problem with Time Windows*. PhD thesis, Technical University of Denmark.
- Lecluyse, C., Van Woensel, T., and Peremans, H. (2009). Vehicle routing with stochastic time-dependent travel times. *4OR*, 7(4):363–377.
- Lee, C., Lee, K., and Park, S. (2012). Robust vehicle routing problem with deadlines and travel time/demand uncertainty. *Journal of the Operational Research Society*, 63(9):1294–1306.
- Lei, F., Wang, Y., Lu, G., and Sun, J. (2014). A travel time reliability model of urban expressways with varying levels of service. *Transportation Research Part C: Emerging Technologies*, 48:453–467.
- Letchford, A. N., Lysgaard, J., and Eglese, R. W. (2007). A branch-and-cut algorithm for the capacitated open vehicle routing problem. *Journal of the Operational Research Society*, 58(12):1642–1651.
- Levenberg, K. (1944). A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168.
- Lewis, R. (2017). Available at: <https://rhydLewis.eu/sbrp/>.
- Lewis, R. and Smith-Miles, K. (2018). A heuristic algorithm for finding cost-effective solutions to real-world school bus routing problems. *Journal of Discrete Algorithms*, 52-53:2–17.
- Lewis, R., Smith-Miles, K., and Phillips, K. (2018). The school bus routing problem: An analysis and algorithm. In Brankovic, L., Ryan, J., and Smyth, W. F., editors, *Combinatorial Algorithms: 28th International Workshop, IWOCA 2017, Newcastle, NSW, Australia, July 17-21, 2017, Revised*

- Selected Papers*, volume 10765 of *Lecture Notes in Computer Science (LNCS)*, pages 287–298, Cham, Switzerland. Springer.
- Li, K., Bai, H., Yan, X., Zhu, M., and Wang, X. (2023). School bus routing considering operation and travel costs. *Journal of Information and Computing Science*, 18(2):141–160.
- Li, L. Y. O. and Fu, Z. (2002). The school bus routing problem: A case study. *Journal of the Operational Research Society*, 53(5):552–558.
- Li, M. and Chow, J. (2021). School bus routing problem with a mixed ride, mixed load, and heterogeneous fleet. *Transportation Research Record: Journal of the Transportation Research Board*, 2675:467–479.
- Li, M., Zhou, X., and Roupail, N. M. (2017). Quantifying travel time variability at a single bottleneck based on stochastic capacity and demand distributions. *Journal of Intelligent Transportation Systems*, 21(2):79–93.
- Li, X., Tian, P., and Leung, S. C. (2010a). Vehicle routing problems with time windows and stochastic travel and service times: Models and algorithm. *International Journal of Production Economics*, 125(1):137–145.
- Li, Y. and Chung, S. H. (2019). Disaster relief routing under uncertainty: A robust optimization approach. *IIE Transactions*, 51(8):869–886.
- Li, Z., Hensher, D. A., and Rose, J. M. (2010b). Willingness to pay for travel time reliability in passenger transport: A review and some new empirical evidence. *Transportation Research Part E: Logistics and Transportation Review*, 46(3):384–403.
- Lin, S. (1965). Computer solutions of the traveling salesman problem. *The Bell System Technical Journal*, 44(10):2245–2269.
- Liu, F., Lu, C., Gui, L., Zhang, Q., Tong, X., and Yuan, M. (2023). Heuristics for vehicle routing problem: A survey and recent advances. Available at: <https://arxiv.org/pdf/2303.04147> (Accessed: August 19th, 2024).
- Lomax, T., Schrank, D., Turner, S., and Margiotta, R. (2003). Selecting travel reliability measures. Technical report, Texas Transportation Institute, Texas, USA.
- López Santana, E. R. and Romero Carvajal, J. (2015). A hybrid column generation and clustering approach to the school bus routing problem with time windows. *Ingeniería*, 20(1):101–117.
- Lourenço, H. R., Martin, O. C., and Stützle, T. (2003). Iterated local search. In Glover, F. and Kochenberger, G. A., editors, *Handbook of Metaheuristics*, volume 57, pages 320–353. Springer, New York, USA.
- Low, V. J. M., Khoo, H. L., and Khoo, W. C. (2021). Statistical modelling of bus travel time with Burr distribution. *ITM Web of Conferences*, 36:01011.
- Lyman, K. and Bertini, R. L. (2008). Using travel time reliability measures

- to improve regional transportation planning and operations. *Transportation Research Record: Journal of the Transportation Research Board*, 2046(1):1–10.
- Ma, Z., Ferreira, L., Mesbah, M., and Zhu, S. (2016). Modeling distributions of travel time variability for bus operations. *Journal of Advanced Transportation*, 50(1):6–24.
- Ma, Z., Koutsopoulos, H. N., Ferreira, L., and Mesbah, M. (2017). Estimation of trip travel time distribution using a generalized Markov chain approach. *Transportation Research Part C: Emerging Technologies*, 74:1–21.
- Mandujano, P., Giesen, R., and Ferrer, J.-C. (2012). Model for optimization of locations of schools and student transportation in rural areas. *Transportation Research Record: Journal of the Transportation Research Board*, 2283(1):74–80.
- Marée, S. (2012). *Correcting Non Positive Definite Correlation Matrices*. Bachelor’s thesis, Electrical Engineering, Mathematics and Computer Science, TU Delft.
- Marquardt, D. W. (1963). An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441.
- Martello, S. and Toth, P. (1981). An algorithm for the generalized assignment problem. In Brans, J., editor, *Operational Research '81*, pages 589–603, Amsterdam, Netherland. North-Holland Publishing Company.
- Martínez, L. M. and Viegas, J. M. (2011). Design and deployment of an innovative school bus service in lisbon. *Procedia - Social and Behavioral Sciences*, 20:120–130.
- Masinick, J., Teng, H., and Orochena, N. (2014). The impact of rubbernecking on urban freeway traffic. *Journal of Transportation Technologies*, 4(1):116–125.
- Mehta, N. B., Molisch, A. F., Wu, J., and Zhang, J. (2006). Approximating the sum of correlated lognormal or lognormal-rice random variables. In *2006 IEEE International Conference on Communications*, volume 4, pages 1605–1610.
- Mehta, N. B., Wu, J., Molisch, A. F., and Zhang, J. (2007). Approximating a sum of random variables with a lognormal. *IEEE Transactions on Wireless Communications*, 6(7):2690–2699.
- Messaoud, E. (2022). Solving a stochastic programming with recourse model for the stochastic electric capacitated vehicle routing problem using a hybrid genetic algorithm. *European Journal of Industrial Engineering*, 16(1):71–90.
- Miller, C. E., Tucker, A. W., and Zemlin, R. A. (1960). Integer programming formulation of traveling salesman problems. *Journal of the ACM*, 7(4):326–329.
- Miller-Hooks, E. and Mahmassani, H. S. (1998). Optimal routing of hazardous materials in stochastic, time-varying transportation networks. *Transportation Research Record: Journal of the Transportation Research Board*, 1645(1):143–

151.

- Minocha, B. and Tripathi, S. (2014). Solving school bus routing problem using hybrid genetic algorithm: A case study. In Babu, B. V., Nagar, A., Deep, K., Pant, M., Bansal, J. C., Ray, K., and Gupta, U., editors, *Proceedings of the Second International Conference on Soft Computing for Problem Solving (SocProS 2012)*, pages 93–103, New Delhi, India. Springer India.
- Miranda, D. M. and Conceição, S. V. (2016). The vehicle routing problem with hard time windows and stochastic travel and service time. *Expert Systems with Applications*, 64:104–116.
- Miranda, D. M., de Camargo, R. S., Conceição, S. V., Porto, M. F., and Nunes, N. T. (2018). A multi-loading school bus routing problem. *Expert Systems with Applications*, 101:228–242.
- Miranda, D. M., de Camargo, R. S., Conceição, S. V., Porto, M. F., and Nunes, N. T. (2021). A metaheuristic for the rural school bus routing problem with bell adjustment. *Expert Systems with Applications*, 180:115086.
- Mladenović, N. and Hansen, P. (1997). Variable neighborhood search. *Computers & Operations Research*, 24(11):1097–1100.
- Mohd Azmi, N. I. L., Md Mahdzir, A. H., Xuan, H. Y., Yu, N. X., Tiean, T. M., and Kun, T. Y. (2024). The review of recent trend for school bus routing problem. *Journal of Engineering Technology and Applied Physics*, 6(1):67–75.
- Mokhtari, N.-a. and Ghezavati, V. (2018). Integration of efficient multi-objective ant-colony and a heuristic method to solve a novel multi-objective mixed load school bus routing model. *Applied Soft Computing*, 68:92–109.
- Mulvey, J. M. and Beck, M. P. (1984). Solving capacitated clustering problems. *European Journal of Operational Research*, 18(3):339–348.
- Munakata, K. (2007). Examination of correlation between travel times on links. Master of Engineering Research Report, Department of Civil and Natural Resources Engineering, University of Canterbury, Christchurch, New Zealand.
- Munari, P. and Savelsbergh, M. (2022). Compact formulations for split delivery routing problems. *Transportation Science*, 56(4):1022–1043.
- MVA Consultancy (1996). The benefits of reduced travel time variability. Technical Report 02/C/ 2923, Department for Transport, London, United Kingdom.
- Na, B., Jun, Y., and Kim, B.-I. (2011). Some extensions to the sweep algorithm. *The International Journal of Advanced Manufacturing Technology*, 56(9):1057–1067.
- Newton, R. M. and Thomas, W. H. (1969). Design of school bus routes by computer. *Socio-Economic Planning Sciences*, 3(1):75–85.
- Newton, R. M. and Thomas, W. H. (1974). Bus routing in a multi-school system. *Computers & Operations Research*, 1(2):213–222.

- Ngonyani, B., Mujuni, E., and Mushi, A. (2015). Optimizing schedules for school bus routing problem: The case of Dar Es Salaam schools. *International Journal of Advanced Research in Computer Science*, 6(1):132–136.
- Nicholson, A. (2015). Travel time reliability benefits: Allowing for correlation. *Research in Transportation Economics*, 49:14–21.
- Nie, Y. M. and Wu, X. (2009a). Reliable a priori shortest path problem with limited spatial and temporal dependencies. In Lam, W. H. K., Wong, S. C., and Lo, H. K., editors, *Transportation and Traffic Theory 2009: Golden Jubilee*, Springer Books, pages 169–195. Springer, Boston, Massachusetts.
- Nie, Y. M. and Wu, X. (2009b). Shortest path problem considering on-time arrival probability. *Transportation Research Part B: Methodological*, 43(6):597–613.
- Nikolova, E. (2009). *Strategic algorithms*. PhD thesis, Massachusetts Institute of Technology.
- Ochoa-Zezzatti, A., Carbajal, U., Castillo, O., Mejía, J., Rivera, G., and Gonzalez, S. (2020). Development of a Java library to solve the school bus routing problem. In Banat, M. M. and Paiva, S., editors, *Smart Technologies for Smart Cities*, pages 175–196, Cham, Switzerland. Springer.
- Oluwadare, S., Oguntuyi, I. P., and Nwaiwu, J. C. (2018). Solving school bus routing problem using genetic algorithm-based model. *International Journal of Intelligent Systems and Applications*, 10(3):50–58.
- Or, I. (1976). *Traveling salesman-type combinatorial problems and their relation to the logistics of regional blood banking*. PhD thesis, Northwestern University.
- Ordóñez, F. (2010). Robust vehicle routing. *INFORMS TutORials in Operations Research*, pages 153–178.
- Orloff, C. S. (1976). Route constrained fleet scheduling. *Transportation Science*, 10(2):149–168.
- Osman, I. H. (1993). Metastrategy simulated annealing and tabu search algorithms for the vehicle routing problem. *Annals of Operations Research*, 41:421–451.
- Oyola, J. (2019). The capacitated vehicle routing problem with soft time windows and stochastic travel times. *Revista Facultad de Ingeniería*, 28(50):19–33.
- Oyola, J., Arntzen, H., and Woodruff, D. L. (2017). The stochastic vehicle routing problem, a literature review, part II: solution methods. *EURO Journal on Transportation and Logistics*, 6(4):349–388.
- Oyola, J., Arntzen, H., and Woodruff, D. L. (2018). The stochastic vehicle routing problem, a literature review, part I: models. *EURO Journal on Transportation and Logistics*, 7(3):193–221.
- Ozmen, M. and Sahin, H. (2021). Real-time optimization of school bus routing problem in smart cities using genetic algorithm. In *Proceedings of the Sixth In-*

- ternational Conference on Inventive Computation Technologies (ICICT 2021), pages 1152–1158, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Pacheco, J., Caballero, R., Laguna, M., and Molina, J. (2013). Bi-objective bus routing: An application to school buses in rural areas. *Transportation Science*, 47:397–411.
- Pacheco, J. and Marti, R. (2006). Tabu search for a multi-objective routing problem. *Journal of the Operational Research Society*, 57:29–37.
- Park, D. and Rilett, L. R. (1999). Forecasting freeway link travel times with a multilayer feedforward neural network. *Computer-Aided Civil and Infrastructure Engineering*, 14(5):357–367.
- Park, J. and Kim, B.-I. (2010). The school bus routing problem: A review. *European Journal of Operational Research*, 202(2):311–319.
- Park, J., Tae, H., and Kim, B.-I. (2012). A post-improvement procedure for the mixed load school bus routing problem. *European Journal of Operational Research*, 217(1):204–213.
- Parvasi, S. P., Mahmoodjanloo, M., and Mostafa, S. (2017). A bi-level school bus routing problem with bus stops selection and possibility of demand outsourcing. *Applied Soft Computing*, 61:222–238.
- Paul, O., Handl, J., and López-Ibáñez, M. (2023). Capacitated school bus routing problem with time windows, heterogeneous fleets and travel assistants. In Silva, S. and Paquete, L., editors, *Proceedings of the 2023 Genetic and Evolutionary Computation Conference, GECCO '23 Companion*, page 2374–2377, New York, USA. Association for Computing Machinery.
- Pekel Özmen, E. and Küçükdeniz, T. (2024). Two-stage clustering and routing problem by using fcm and k-means with genetic algorithm. *Sigma Journal of Engineering and Natural Sciences*, 42(4):1030–1038.
- Pérez-Rodríguez, R. and Hernández-Aguirre, A. (2016). Probability model to solve the school bus routing problem with stops selection. *International Journal of Combinatorial Optimization Problems and Informatics*, 7(1):30–39.
- Polus, A. (1979). A study of travel time and reliability on arterial routes. *Transportation*, 8(2):141–151.
- Potvin, J.-Y. and Rousseau, J.-M. (1995). An exchange heuristic for routeing problems with time windows. *The Journal of the Operational Research Society*, 46(12):1433–1446.
- Prakash, A. A. and Srinivasan, K. K. (2018). Pruning algorithms to determine reliable paths on networks with random and correlated link travel times. *Transportation Science*, 52(1):80–101.
- Pu, W. (2011). Analytic relationships between travel time reliability mea-

- sures. *Transportation Research Record: Journal of the Transportation Research Board*, 2254(1):122–130.
- Qin, W., Ji, X., and Liang, F. (2020). Estimation of urban arterial travel time distribution considering link correlations. *Transportmetrica A: Transport Science*, 16(3):1429–1458.
- Rachtan, P., Huang, H., and Gao, S. (2013). Spatiotemporal link speed correlations: Empirical study. *Transportation Research Record: Journal of the Transportation Research Board*, 2390(1):34–43.
- Rajabi-Bahaabadi, M., Shariat-Mohaymany, A., Babaei, M., and Vigo, D. (2021). Reliable vehicle routing problem in stochastic networks with correlated travel times. *Operational Research*, 21:299–330.
- Rakha, H., El-Shawarby, I., and Arafah, M. (2010). Trip travel-time reliability: Issues and proposed solutions. *Journal of Intelligent Transportation Systems*, 14(4):232–250.
- Ramezani, M. and Geroliminis, N. (2012). On the estimation of arterial route travel time distribution with Markov chains. *Transportation Research Part B: Methodological*, 46(10):1576–1590.
- Rashidi, S., Ataeian, S., and Ranjitkar, P. (2023). Estimating bus dwell time: A review of the literature. *Transport Reviews*, 43(1):32–61.
- Rego, C. (2001). Node-ejection chains for the vehicle routing problem: Sequential and parallel algorithms. *Parallel Computing*, 27(3):201–222.
- Reina, P. (2021). Understanding the formation of rubbernecking queues on urban freeways. *Transportation Research Interdisciplinary Perspectives*, 9:100266.
- Ren, J., Jin, W., and Wu, W. (2019). A two-stage algorithm for school bus stop location and routing problem with walking accessibility and mixed load. *IEEE Access*, 7:119519–119540.
- Richardson, A. J. and Taylor, M. (1978). Travel time variability on commuter journeys. *High Speed Ground Transportation Journal*, 12(1):77–99.
- Riera-Ledesma, J. and Salazar-González, J.-J. (2012). Solving school bus routing using the multiple vehicle traveling purchaser problem: A branch-and-cut approach. *Computers & Operations Research*, 39(2):391–404.
- Riera-Ledesma, J. and Salazar-González, J. J. (2013). A column generation approach for a school bus routing problem with resource constraints. *Computers & Operations Research*, 40(2):566–583.
- Ripplinger, D. (2005). Rural school vehicle routing problem. *Transportation Research Record: Journal of the Transportation Research Board*, 1922(1):105–110.
- Ritzinger, U., Puchinger, J., and Hartl, R. F. (2016). A survey on dynamic and stochastic vehicle routing problems. *International Journal of Production*

- Research*, 54(1):215–231.
- Rosenkrantz, D. J., Stearns, R. E., and Lewis, P. M. (1974). Approximate algorithms for the traveling salesperson problem. In *15th Annual Symposium on Switching and Automata Theory*, pages 33–42, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Rostami, B., Desaulniers, G., Errico, F., and Lodi, A. (2017). The vehicle routing problem with stochastic and correlated travel times. Available at: [https://cerc-datascience.polymtl.ca/wp-content/uploads/2018/01/Technical-Report\\_DS4DM-2017-016.pdf](https://cerc-datascience.polymtl.ca/wp-content/uploads/2018/01/Technical-Report_DS4DM-2017-016.pdf) (Accessed: February 7th, 2025).
- Rostami, B., Desaulniers, G., Errico, F., and Lodi, A. (2021). Branch-price-and-cut algorithms for the vehicle routing problem with stochastic and correlated travel times. *Operations Research*, 69(2):436–455.
- Rubin, P. (2011). Perils of ‘Big M’, OR in an OB World. Available at: <https://orinanobworld.blogspot.com/2011/07/perils-of-big-m.html> (Accessed: July 5th, 2024).
- Russell, R. A., Morrel, R. B., and Haddox, B. (1986). Routing special-education school buses. *Interfaces*, 16(5):56–64.
- Russell, R. A. and Urban, T. L. (2008). Vehicle routing with soft time windows and Erlang travel times. *The Journal of the Operational Research Society*, 59(9):1220–1228.
- Safak, A. (1993). Statistical analysis of the power sum of multiple correlated log-normal components. *IEEE Transactions on Vehicular Technology*, 42(1):58–61.
- Sales, L., Melo, C., Bonates, T., and Prata, B. (2018). Memetic algorithm for the heterogeneous fleet school bus routing problem. *Journal of Urban Planning and Development*, 144(2):04018018.
- Samadi-Dana, S., Paydar, M. M., and Jouzdani, J. (2017). A simulated annealing solution method for robust school bus routing. *International Journal of Operational Research*, 28(3):307–326.
- Sansone, K. (2018). No free school transport for 1,000 students in church and independent schools. *MaltaToday*. Available at: [https://www.maltatoday.com.mt/news/national/89608/no\\_free\\_school\\_transport\\_for\\_1000\\_students\\_in\\_church\\_and\\_independent\\_schools](https://www.maltatoday.com.mt/news/national/89608/no_free_school_transport_for_1000_students_in_church_and_independent_schools) (Accessed: May 27th, 2024).
- Sarmiento, A., Ruiz, C. F., Akhavan-Tabatabaei, R., Medaglia, A., and Mendoza, J. E. (2015). Evaluating route reliability for the stochastic vehicle routing problem (SVRP) with correlated parameters. Available at: <https://repositorio.uniandes.edu.co/server/api/core/bitstreams/bb0d3166-e88d-450a-b1bc-01a1950ad166/content> (Accessed: February 12th, 2025).

- Sarubbi, J. F. M., Mesquita, C. M. R., Wanner, E. F., Santos, V. F., and Silva, C. M. (2016). A strategy for clustering students minimizing the number of bus stops for solving the school bus routing problem. In Oktug, S., Ulema, M., Cavdar, C., Granville, L. Z., and dos Santos, C. R. P., editors, *NOMS 2016 - 2016 IEEE/IFIP Network Operations and Management Symposium*, pages 1175–1180, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Savas, E. S. (1978). On equity in providing public services. *Management Science*, 24(8):800–808.
- Schittekat, P., Kinable, J., Sørensen, K., Sevaux, M., Spieksma, F., and Springael, J. (2013). A metaheuristic for the school bus routing problem with bus stop selection. *European Journal of Operational Research*, 229(2):518–528.
- Schittekat, P., Sevaux, M., and Sorensen, K. (2006). A mathematical formulation for a school bus routing problem. In *2006 International Conference on Service Systems and Service Management (ICSSSM'06)*, volume 2, pages 1552–1557, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Schleher, D. C. (1977). Generalized gram-charlier series with application to the sum of log-normal variates. *IEEE Transactions on Information Theory*, 23:275–280.
- Schwartz, S. C. and Yeh, Y. S. (1982). On the distribution function and moments of power sums with log-normal components. *The Bell System Technical Journal*, 61(7):1441–1462.
- Sciortino, M. (2024a). Iterated local search algorithm for the single-school bus routing problem with deterministic travel times: Source code and outputs (Part I). Available at: <https://doi.org/10.5281/zenodo.14379476>.
- Sciortino, M. (2024b). Iterated local search algorithm for the single-school bus routing problem with deterministic travel times: Source code and outputs (Part II). Available at: <https://doi.org/10.5281/zenodo.14380351>.
- Sciortino, M. (2024c). Mixed integer program for the single-school bus routing problem: Source code and outputs. Available at: <https://doi.org/10.5281/zenodo.13121107>.
- Sciortino, M. (2024d). Ten real-world problem instances (Malta) of the single-school bus routing problem. Available at: <https://doi.org/10.5281/zenodo.13120919>.
- Sciortino, M. (2025a). Detection of overlapping road segments between distinct arcs in a network. Available at: <https://doi.org/10.5281/zenodo.17348746>.
- Sciortino, M. (2025b). Iterated local search algorithm for the single-school

- bus routing problem with shifted log normal and correlated travel times: Source code and outputs. Available at: <https://doi.org/10.5281/zenodo.17311521>.
- Sciortino, M. (2025c). Iterated local search algorithm for the single-school bus routing problem with shifted log normal and independent travel times: Source code and outputs. Available at: <https://doi.org/10.5281/zenodo.16874845>.
- Sciortino, M. (2025d). Monte Carlo sampling methodology for the single-school bus routing problem with stochastic and correlated travel times. Available at: <https://doi.org/10.5281/zenodo.17311931>.
- Semba, S. and Mujuni, E. (2019). An empirical performance comparison of meta-heuristic algorithms for school bus routing problem. *Tanzania Journal of Science*, 45(1):81–92.
- Sen, A., Thakuriah, P., Zhu, X., and Karr, A. (1997). Frequency of probe reports and variance of travel time estimates. *Journal of Transportation Engineering*, 123(4):290–297.
- Seshadri, R. and Srinivasan, K. K. (2010). Algorithm for determining most reliable travel time path on network with normally distributed and correlated link travel times. *Transportation Research Record: Journal of the Transportation Research Board*, 2196(1):83–92.
- Sghaier, S. B., Guedria, N. B., and Rafea, M. (2013). Solving school bus routing problem with genetic algorithm. In *2013 International Conference on Advanced Logistics and Transport (ICALT 2013)*, pages 7–12, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Shafahi, A., Wang, Z., and Haghani, A. (2017). Solving the school bus routing problem by maximizing trip compatibility. *Transportation Research Record*, 2667(1):17–27.
- Shafahi, A., Wang, Z., and Haghani, A. (2018a). A matching-based heuristic algorithm for school bus routing problems. Available at: <https://arxiv.org/pdf/1807.05311> (Accessed: August 9th, 2024).
- Shafahi, A., Wang, Z., and Haghani, A. (2018b). Speedroute: Fast, efficient solutions for school bus routing problems. *Transportation Research Part B: Methodological*, 117:473–493.
- Shahabi, M., Unnikrishnan, A., and Boyles, S. D. (2013). An outer approximation algorithm for the robust shortest path problem. *Transportation Research Part E: Logistics and Transportation Review*, 58:52–66.
- Sharma, A. K. (2024). Optimizing school bus routes with ant colony optimization: A dynamic approach to sustainable transportation logistics. *Advanced International Journal of Multidisciplinary Research*, 2(2):1–12.

- Silva, C. M., Sarubbi, J. F. M., Silva, D. F., Porto, M. F., and Nunes, N. T. R. (2015). A mixed load solution for the rural school bus routing problem. In *2015 IEEE 18th International Conference on Intelligent Transportation Systems*, pages 1940–1945, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Siqueira, V., Silva, F., Silva, E., Silva, R., and Lisboa, M. (2016). Implementation of the metaheuristic GRASP applied to the school bus routing problem. *International Journal of e-Education, e-Business, e-Management and e-Learning*, 6(2):137–145.
- Slimane, S. B. (2001). Bounds on the distribution of a sum of independent lognormal random variables. *IEEE Transactions on Communications*, 49(6):975–978.
- Solano-Charris, E., Prins, C., and Santos, A. C. (2015). Local search based metaheuristics for the robust vehicle routing problem with discrete scenarios. *Applied Soft Computing*, 32:518–531.
- Song, G. S. and Kim, M. C. (2022). Theoretical approach for uncertainty quantification in probabilistic safety assessment using sum of lognormal random variables. *Nuclear Engineering and Technology*, 54(6):2084–2093.
- Souza Lima, F. M., Pereira, D. S., Conceição, S. V., and Camargo, R. S. (2017). A multi-objective capacitated rural school bus routing problem with heterogeneous fleet and mixed loads. *4OR*, 15(4):359–386.
- Souza Lima, F. M., Pereira, D. S., Conceição, S. V., and Ramos Nunes, N. T. (2016). A mixed load capacitated rural school bus routing problem with heterogeneous fleet: Algorithms for the Brazilian context. *Expert Systems with Applications*, 56:320–334.
- Spada, M., Bierlaire, M., and Lieblich, T. M. (2005). Decision-aiding methodology for the school bus routing and scheduling problem. *Transportation Science*, 39(4):443–556.
- Srinivasan, K. K., Prakash, A., and Seshadri, R. (2014). Finding most reliable paths on networks with correlated and shifted log-normal travel times. *Transportation Research Part B: Methodological*, 66(8):110–128.
- Sterman, B. P. and Schofer, J. L. (1976). Factors affecting reliability of urban bus services. *Transportation Engineering Journal of ASCE*, 102(1):147–159.
- Sun, L. and Wang, B. (2015). Robust optimisation approach for vehicle routing problems with uncertainty. *Mathematical Problems in Engineering*, 2015(1):901583.
- Sun, Q., Chien, S., Hu, D., and Chen, X. (2021). Optimizing customized transit service considering stochastic bus arrival time. *Journal of Advanced Transportation*, 2021(1):3207025.
- Sun, S., Duan, Z., and Xu, Q. (2018). School bus routing problem in the stochastic

- and time-dependent transportation network. *PLOS ONE*, 13(8):e0202618.
- Susilawati, S., Taylor, M. A. P., and Somenahalli, S. V. C. (2013). Distributions of travel time variability on urban roads. *Journal of Advanced Transportation*, 47(8):720–736.
- Taillard, É., Badeau, P., Gendreau, M., Guertin, F., and Potvin, J.-Y. (1997). A tabu search heuristic for the vehicle routing problem with soft time windows. *Transportation Science*, 31:170–186.
- Tan, S.-Y. and Yeh, W.-C. (2021). The vehicle routing problem: State-of-the-art classification and review. *Applied Sciences*, 11(21):10295.
- Taniguchi, E., Thompson, R. G., Yamada, T., and van Duin, R. (2001). *City Logistics: Network Modelling and Intelligent Transport Systems*. Emerald Group Publishing Limited, Leeds, England.
- Taş, D., Dellaert, N., van Woensel, T., and de Kok, T. (2013). Vehicle routing problem with stochastic travel times including soft time windows and service costs. *Computers & Operations Research*, 40(1):214–224.
- Tavakkoli-Moghaddam, R., Alinaghian, M., Salamat-Bakhsh, A., and Norouzi, N. (2012). A hybrid meta-heuristic algorithm for the vehicle routing problem with stochastic travel times considering the driver’s satisfaction. *Journal of Industrial Engineering International*, 8:4.
- Taylor, M. A. (2017a). Fosgerau’s travel time reliability ratio and the Burr distribution. *Transportation Research Part B: Methodological*, 97:50–63.
- Taylor, M. A. P. and Susilawati, S. (2012). Modelling travel time reliability with the Burr distribution. *Procedia - Social and Behavioral Sciences*, 54:75–83.
- Taylor, R. (2017b). Travel time reliability: Making it there on time, all the time. Technical report, U.S. Department of Transportation, Federal Highway Administration (FHWA) Office of Operations. Available at: [https://ops.fhwa.dot.gov/publications/tt\\_reliability/ttr\\_report.htm](https://ops.fhwa.dot.gov/publications/tt_reliability/ttr_report.htm) (Accessed: October 31st, 2024).
- Thangiah, S. and Nygard, K. (1992). School bus routing using genetic algorithms. In Biswas, G., editor, *Applications of Artificial Intelligence, X: Knowledge-based Systems, 22-24 April 1992, Orlando, Florida*, volume 1707 of *Proceedings of SPIE*, pages 387–398, Bellingham, Washington, USA. SPIE.
- Thompson, R. G., Taniguchi, E., and Yamada, T. (2011). Estimating the benefits of considering travel time variability in urban distribution. *Transportation Research Record: Journal of the Transportation Research Board*, 2238(1):86–96.
- Tirachini, A. (2013). Bus dwell time: The effect of different fare collection systems, bus floor level and age of passengers. *Transportmetrica A: Transport Science*, 9(1):28–49.

- Transportation Research Board and National Academies of Sciences, Engineering, and Medicine (2013). *Transit Capacity and Quality of Service Manual, Third Edition*. The National Academies Press, Washington, DC.
- Transportation Research Board and National Academies of Sciences, Engineering, and Medicine (2022). *Highway Capacity Manual 7th Edition: A Guide for Multimodal Mobility Analysis*. The National Academies Press, Washington, DC.
- Turner, J. K. and Wardrop, J. G. (1951). The variation of journey time in central London. Road Research Laboratory Note No. RN/1511/JKT.JGW.
- Ümit, Ü. and Kiliç, F. (2019). A school bus routing problem using genetic algorithm by reducing the number of buses. *2019 Innovations in Intelligent Systems and Applications Conference (ASYU)*, pages 1–6.
- van Lint, J., van Zuylen, H. J., and Tu, H. (2008). Travel time unreliability on freeways: Why measures based on variance tell only half the story. *Transportation Research Part A: Policy and Practice*, 42(1):258–277.
- van Lint, J. W. C. and van Zuylen, H. J. (2005). Monitoring and predicting freeway travel time reliability: using width and skew of day-to-day travel time distribution. *Transportation Research Record: Journal of the Transportation Research Board*, 1917:54–62.
- Vareias, A. D., Repoussis, P. P., and Tarantilis, C. D. (2019). Assessing customer service reliability in route planning with self-imposed time windows and stochastic travel times. *Transportation Science*, 53(1):256–281.
- Vassallo, A. (2024). Education minister states that school transport will remain free. *TVMnews*. Available at: <https://tvmnews.mt/en/news/education-minister-states-that-school-transport-will-remain-free/> (Accessed: May 27th, 2024).
- Vidal, T., Laporte, G., and Matl, P. (2020). A concise guide to existing and emerging vehicle routing problem variants. *European Journal of Operational Research*, 286(2):401–416.
- Wang, Z. and Haghani, A. (2020). Column generation-based stochastic school bell time and bus scheduling optimization. *European Journal of Operational Research*, 286(3):1087–1102.
- Wang, Z. and Lin, W.-H. (2017). Incorporating travel time uncertainty into the design of service regions for delivery/pickup problems with time windows. *Expert Systems with Applications*, 72:207–220.
- Wardrop, J. G. (1952). Some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers*, 1(3):325–362.
- Wei, T. H., Muda, N., Yunus, A. N. M., Othman, A. R., Aïssa, S., and Ahad, N. A. (2024). Approximation of the lognormal distribution as a solution to

- the sum of lognormal variates. *Communications in Statistics - Simulation and Computation*, pages 1–23.
- White, G. P. (1982). An improvement in the gavish-shlifer algorithm for a class of transportation scheduling problems. *European Journal of Operational Research*, 9(2):190–193.
- Wren, A. and Holliday, A. (1972). Computer scheduling of vehicles from one or more depots to a number of delivery points. *Operational Research Quarterly (1970-1977)*, 23(3):333–344.
- Wu, J., Mehta, N., and Zhang, J. (2005). Flexible lognormal sum approximation method. In *GLOBECOM '05, IEEE Global Telecommunications Conference 2005*, volume 6, pages 3413–3417.
- Wu, N. and Geistefeldt, J. (2014). Standard deviation of travel time in a freeway network—a mathematical quantifying tool for reliability analysis. In Ma, J., Yin, Y., Huang, H., and Pan, D., editors, *CICTP 2014: Safe, Smart, and Sustainable Multimodal Transportation Systems*, Proceedings of the 14th COTA International Conference of Transportation Professionals (CICTP 2014) held in Changsha, China, July 4-7 2014, pages 3292–3303, Reston, Virginia, USA. American Society of Civil Engineers.
- Wu, N. and Geistefeldt, J. (2015). Modeling travel time for reliability analysis in a freeway network. Available at: [https://homepage.rub.de/ning.wu/pdf/TRB2016\\_ModelingReliability\\_WuGeistefeldt\\_resub\\_final.pdf](https://homepage.rub.de/ning.wu/pdf/TRB2016_ModelingReliability_WuGeistefeldt_resub_final.pdf) (Accessed: February 7th, 2025).
- Xing, T. and Zhou, X. (2011). Finding the most reliable path with and without link travel time correlation: A Lagrangian substitution based approach. *Transportation Research Part B: Methodological*, 45(10):1660–1679.
- Xue, Y., Jin, J., Lai, J., Ran, B., and Yang, D. (2011). Empirical characteristics of transit travel time distribution for commuting routes. In *Proceedings of the 90th Annual Meeting of the Transportation Research Board*, number 11-2827, Washington, DC. Transportation Research Board.
- Xue, Z., Deng, X., Chen, B., and Khamis, A. (2023). School bus routing using metaheuristics algorithms. In *2023 IEEE International Conference on Smart Mobility (SM)*, pages 33–38, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Yan, S., Hsiao, F.-Y., and Chen, Y.-C. (2015). Inter-school bus scheduling under stochastic travel times. *Networks and Spatial Economics*, 15(4):1049–1074.
- Yan, S., Wang, S.-S., and Chang, Y.-H. (2014). Cash transportation vehicle routing and scheduling under stochastic travel times. *Engineering Optimization*, 46(3):289–307.
- Yang, L. and Zhou, X. (2017). Optimizing on-time arrival probability and per-

- centile travel time for elementary path finding in time-dependent transportation networks: Linear mixed integer programming reformulations. *Transportation Research Part B: Methodological*, 96:68–91.
- Yang, S. and Wu, Y.-J. (2016). Mixture models for fitting freeway travel time distributions and measuring travel time reliability. *Transportation Research Record*, 2594(1):95–106.
- Yao, B., Cao, Q., Wang, Z., Hu, P., Zhang, M., and Yu, B. (2016). A two-stage heuristic algorithm for the school bus routing problem with mixed load plan. *Transportation Letters*, 8(4):205–219.
- Yellow, P. C. (1970). A computational modification to the savings method of vehicle scheduling. *Operational Research Quarterly (1970-1977)*, 21(2):281–283.
- Yin, Y., Yang, Y., Yu, Y., Wang, D., and Cheng, T. (2023). Robust vehicle routing with drones under uncertain demands and truck travel times in humanitarian logistics. *Transportation Research Part B: Methodological*, 174:102781.
- Yuan, M., Jiang, C., Li, S., Shen, W., Pavlidis, Y., and Li, J. (2014). Message passing algorithm for the generalized assignment problem. In Hsu, C.-H., Shi, X., and Salapura, V., editors, *Network and Parallel Computing*, volume 8707, pages 423–434, Berlin, Heidelberg. Springer.
- Zang, Z., Xu, X., Qu, K., Chen, R., and Chen, A. (2022). Travel time reliability in transportation networks: A review of methodological developments. *Transportation Research Part C: Emerging Technologies*, 143:103866.
- Žerovnik, G., Trkov, A., Smith, D. L., and Capote, R. (2013). Transformation of correlation coefficients between normal and lognormal distribution and implications for nuclear applications. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 727:33–39.
- Zhang, D. and Guo, Z. (2019). On the necessity and effects of considering correlated stochastic speeds in shortest path problems under sustainable environments. *Sustainability*, 12(1):1–14.
- Zhang, D., Wallace, S. W., Guo, Z., Dong, Y., and Kaut, M. (2021). On scenario construction for stochastic shortest path problems in real road networks. *Transportation Research Part E: Logistics and Transportation Review*, 152:102410.
- Zhang, H., Ge, H., Yang, J., and Tong, Y. (2022). Review of vehicle routing problems: Models, classification and solving algorithms. *Archives of Computational Methods in Engineering*, 29:195–221.
- Zhang, J., Lam, W., and Chen, B. Y. (2013). A stochastic vehicle routing problem with travel time uncertainty: Trade-off between cost and customer service. *Networks and Spatial Economics*, 13(4):471–496.

- Zhang, T., Chaovaitwongse, W., and Zhang, Y. (2012). Scatter search for the stochastic travel-time vehicle routing problem with simultaneous pick-ups and deliveries. *Computers & Operations Research*, 39(10):2277–2290.
- Zhang, X. (2022). *A survey of the sum of lognormal distribution and some recent results*. Master’s thesis, University of British Columbia.
- Zhang, Y. and Khani, A. (2019). An algorithm for reliable shortest path problem with travel time correlations. *Transportation Research Part B: Methodological*, 121:92–113.
- Zhang, Y., Max Shen, Z.-J., and Song, S. (2017). Lagrangian relaxation for the reliable shortest path problem with correlated link travel times. *Transportation Research Part B: Methodological*, 104:501–521.
- Zhang, Z., He, Q., Gou, J., and Li, X. (2019). Analyzing travel time reliability and its influential factors of emergency vehicles with generalized extreme value theory. *Journal of Intelligent Transportation Systems*, 23(1):1–11.
- Zhao, T., Nie, Y., Wu, X., and Zhang, Y. (2014). Empirical analysis of the dependence structure in traffic data using copula function. In *Proceedings of 2014 IEEE International Conference on Service Operations and Logistics, and Informatics*, pages 38–42, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Zockaie, A., Mahmassani, H. S., and Nie, Y. M. (2016). Path finding in stochastic time varying networks with spatial and temporal correlations for heterogeneous travelers. *Transportation Research Record: Journal of the Transportation Research Board*, 2567:105–113.
- Zockaie, A., Nie, Y., Wu, X., and Mahmassani, H. (2013). Impacts of correlations on reliable shortest path finding: A simulation-based study. *Transportation Research Record: Journal of the Transportation Research Board*, 2334:1–9.
- Zou, H., Yue, Y., Li, Q., and Shi, Y. (2010). A spatial analysis approach for describing spatial pattern of urban traffic state. In *13th International IEEE Conference on Intelligent Transportation Systems*, pages 557–562, Piscataway, New Jersey, USA. Institute of Electrical and Electronics Engineers (IEEE).
- Zou, Y., Zhu, X., Zhang, Y., and Zeng, X. (2014). A space–time diurnal method for short-term freeway travel time prediction. *Transportation Research Part C: Emerging Technologies*, 43:33–49.



# Appendix A

## Visualizations of MIP Solutions

This appendix visualizes the MIP solutions listed in Table 3.4. In Figures A.1 to A.8, the lime dot represents the school, the red dots represent the visited bus stops, and the yellow dots represent the student addresses. Moreover, each route is shown in a different colour. Note that certain subroutes are not visible since they overlap with other subroutes. The following visualizations and further information on arrival times, dwell times, distances covered, and walking/riding times and distances for each student can be viewed from the interactive visualization files at (Sciortino, 2024c).

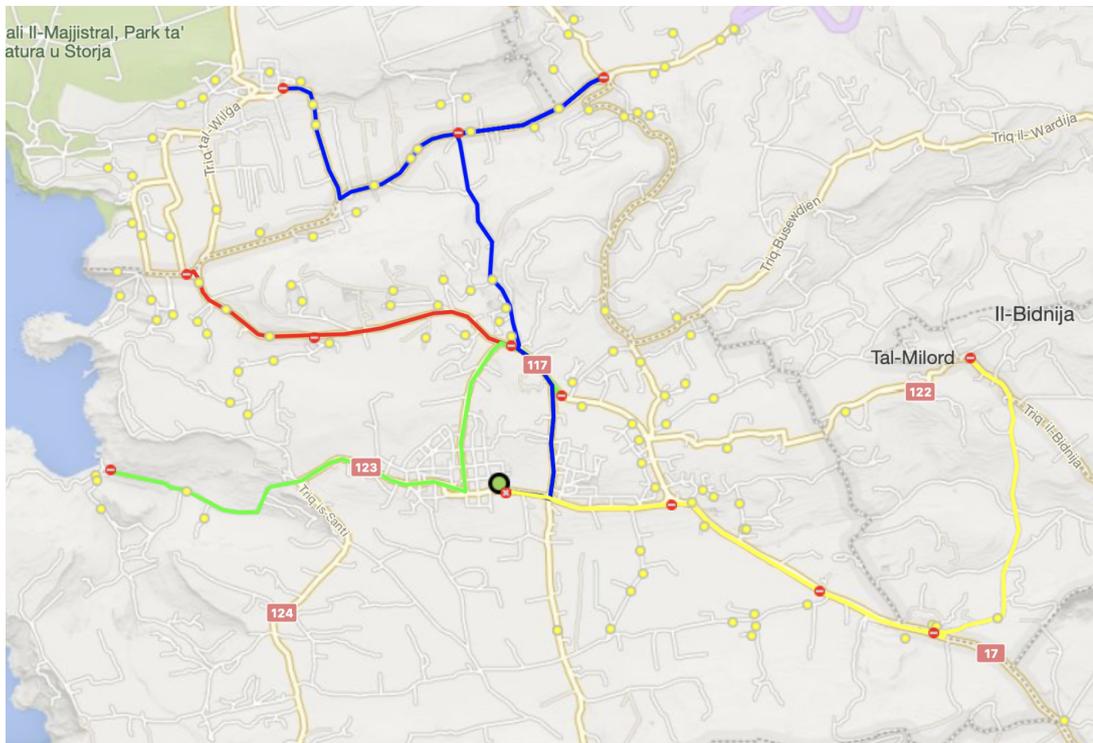


Figure A.1: Mġarr MIP solution with 4 routes (red, lime, blue, yellow), 190 students, average walk 6.06 minutes, and average journey time 13.50 minutes.

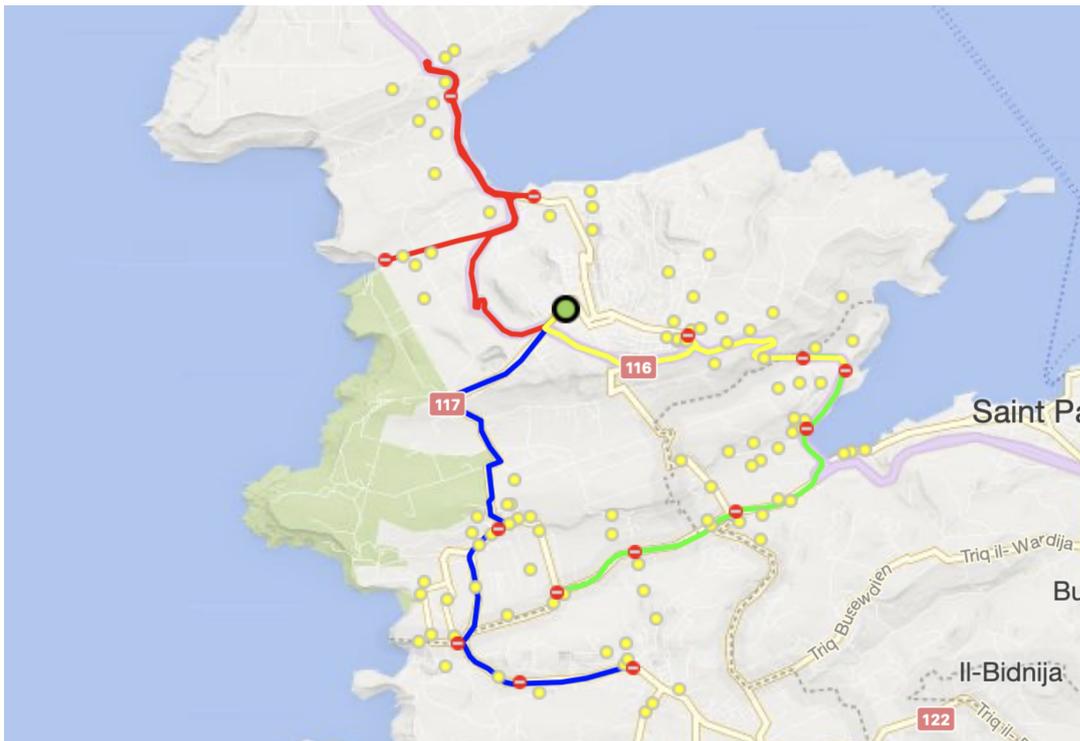


Figure A.2: Mellieħa MIP solution with 4 routes (red, lime, blue, yellow), 171 students, average walk 5.04 minutes, and average journey time 14.07 minutes.

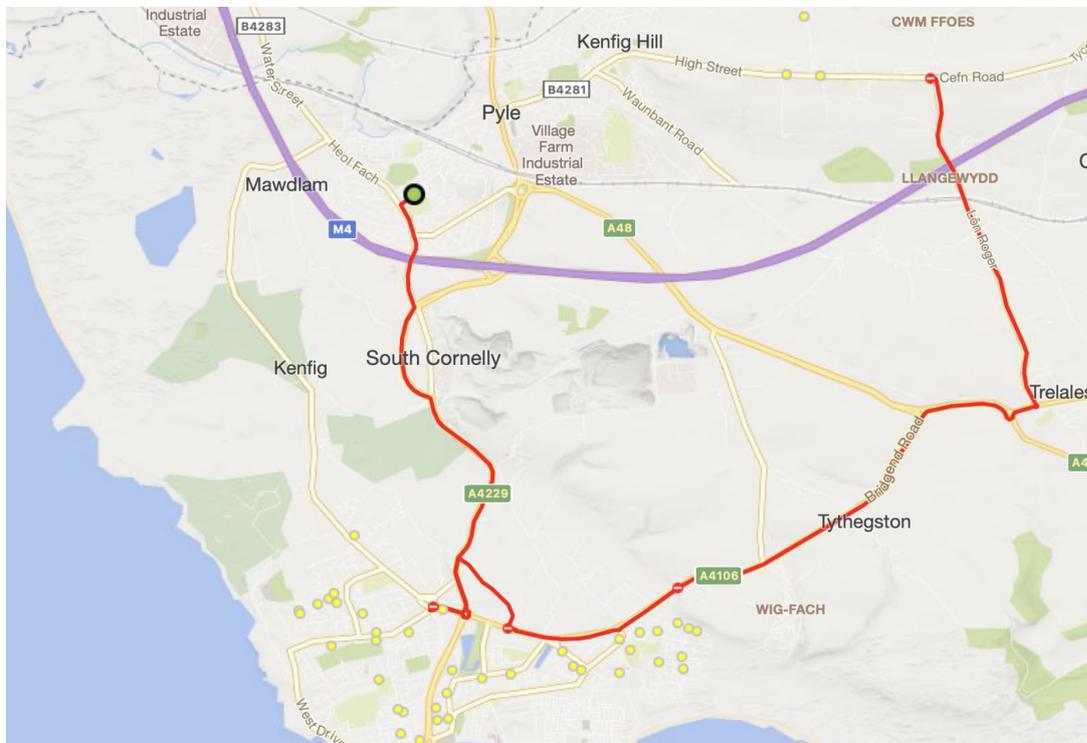


Figure A.3: Porthcawl MIP solution with 1 route (red), 66 students, average walk 14.26 minutes, and average journey time 26.87 minutes.

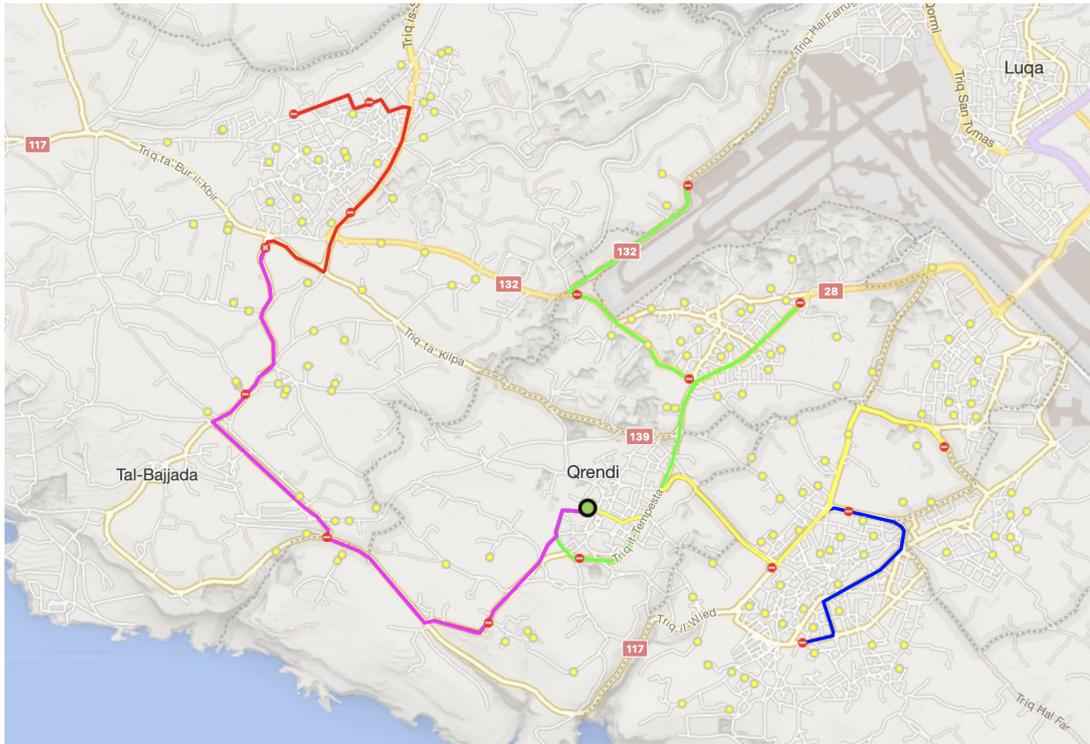


Figure A.4: Qrendi MIP solution with 5 routes (red, lime, blue, yellow, magenta), 255 students, average walk 6.31 minutes, and average journey time 16.03 minutes.

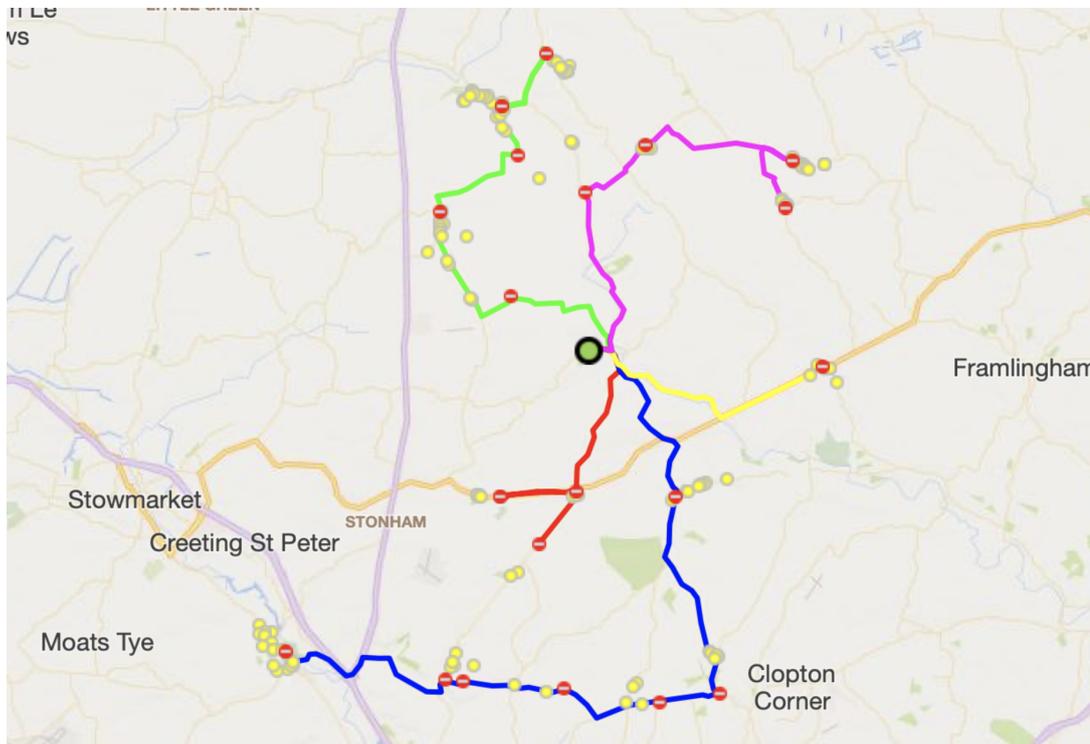


Figure A.5: Suffolk MIP solution with 5 routes (red, lime, blue, yellow, magenta), 209 students, average walk 7.68 minutes, and average journey time 22.85 minutes.

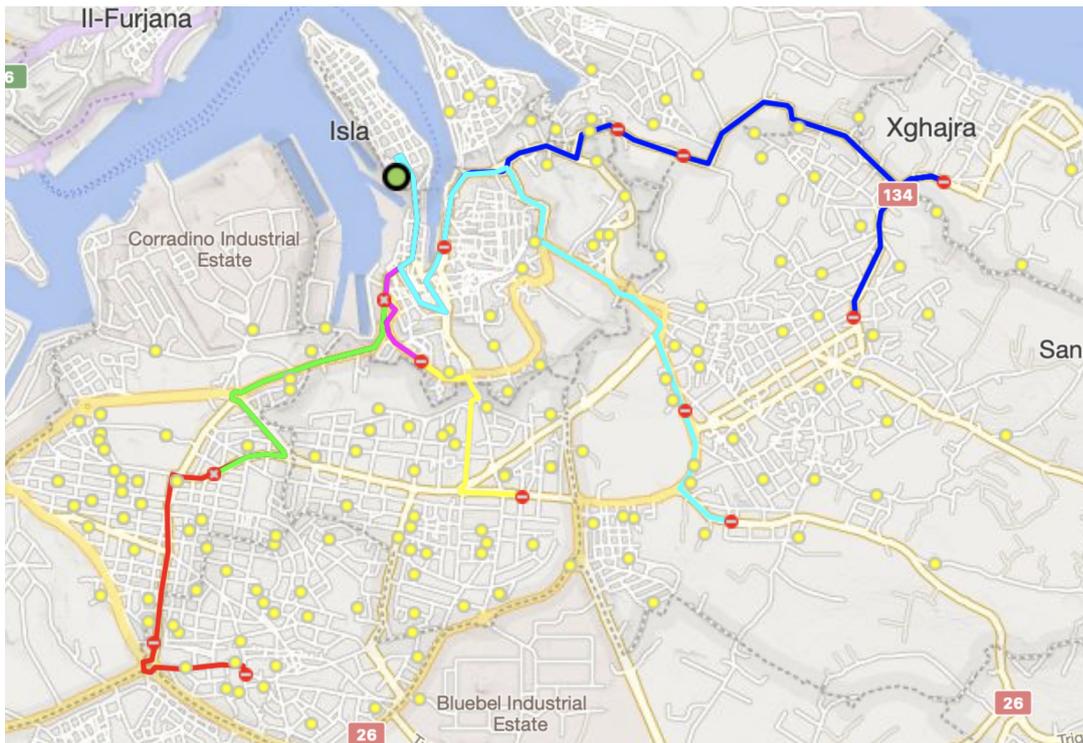


Figure A.6: Senglea MIP solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 266 students, average walk 6.19 minutes, and average journey time 12.88 minutes.

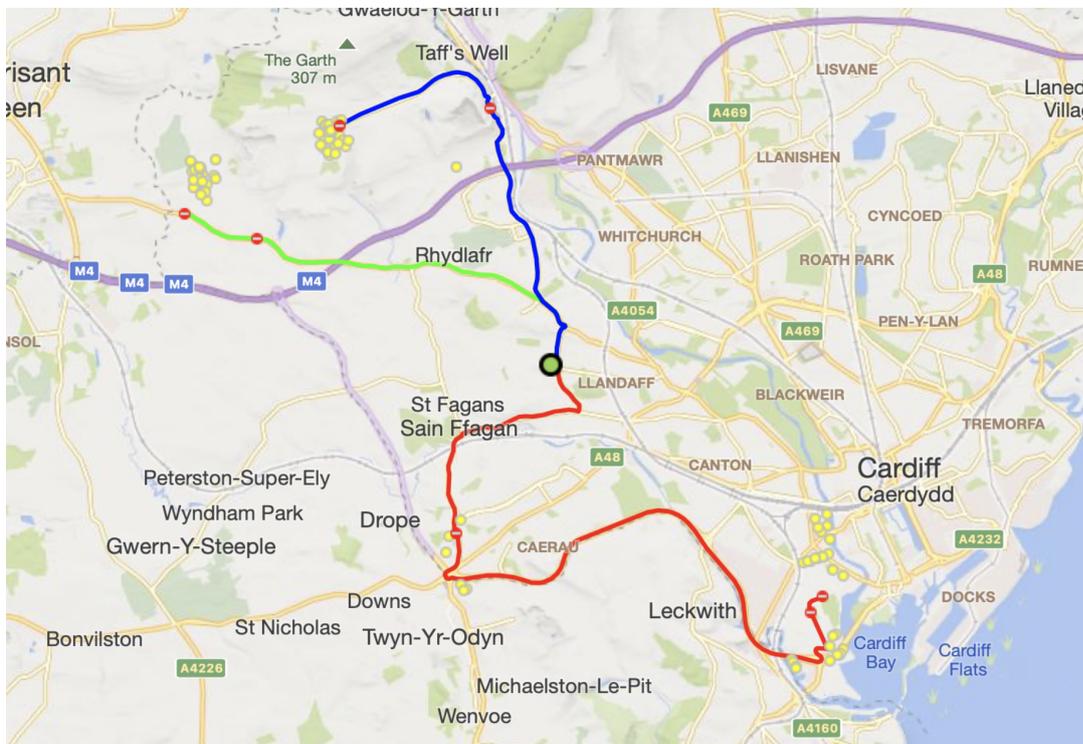


Figure A.7: Cardiff MIP solution with 3 routes (red, lime, blue), 156 students, average walk 11.04 minutes, and average journey time 19.08 minutes.

APPENDIX A. VISUALIZATIONS OF MIP SOLUTIONS

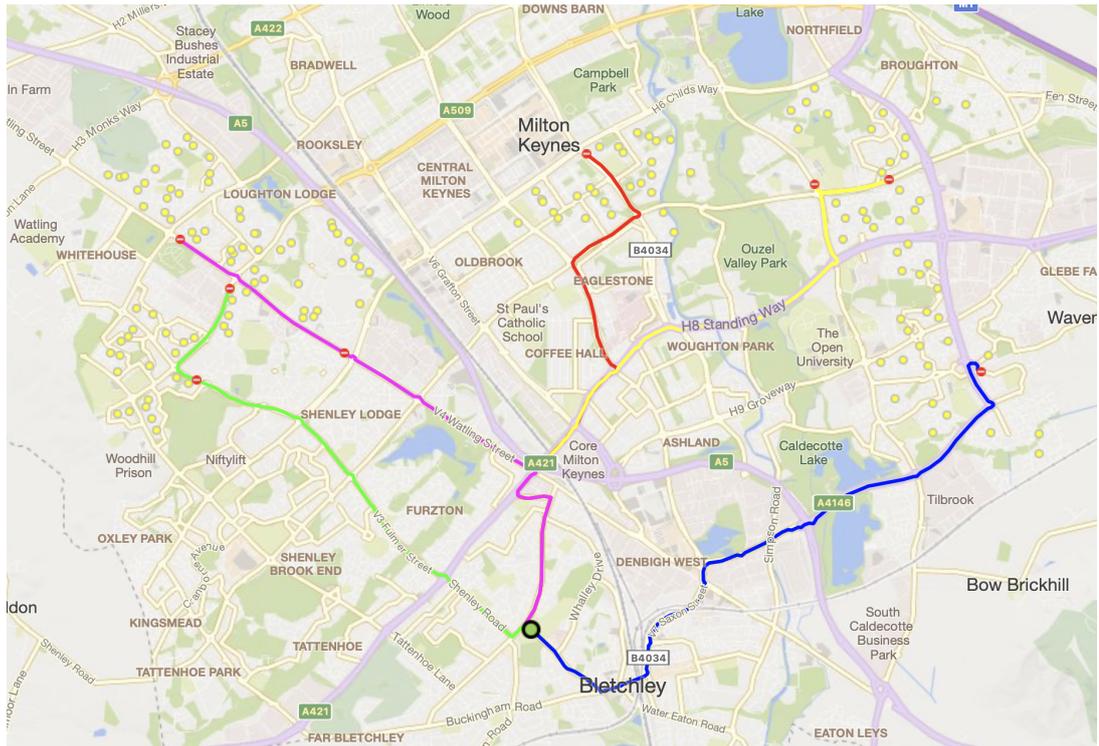


Figure A.8: Milton Keynes MIP solution with 5 routes (red, lime, blue, yellow, magenta), 274 students, average walk 11.40 minutes, and average journey time 13.36 minutes.



# Appendix B

## Heuristic Algorithm

## Deterministic Scenario Results

### B.1 Total Journey Time Boxplots

Boxplots displaying the 25 total journey times (in minutes) reached by each algorithm variant for each instance are shown in Figures B.1 to B.20. In each boxplot, the endpoints of the whiskers represent the minimum and maximum, the endpoints of the box represent the 25th and 75th percentiles, the line inside the box represents the median, the dots represent potential outliers, while the asterisks represent extreme outliers.

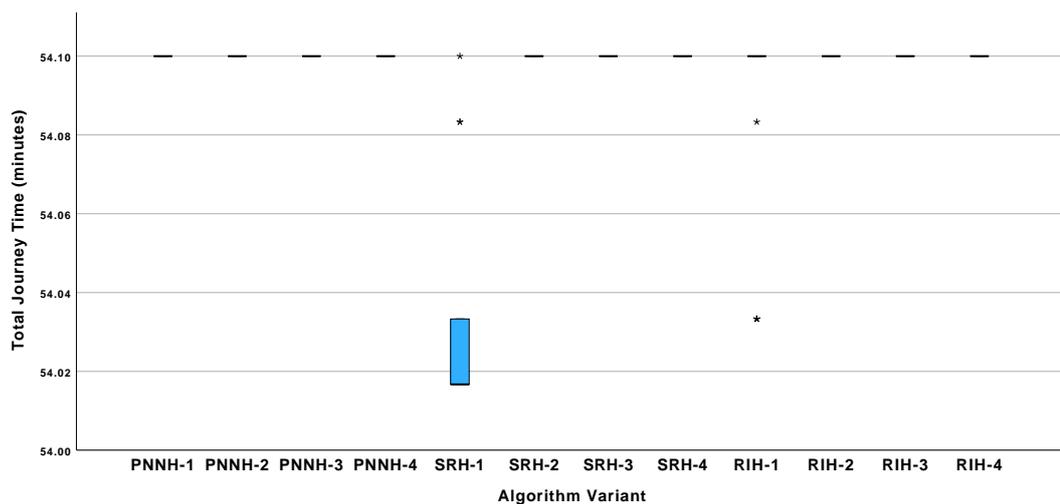


Figure B.1: Total journey time values for the Mğarr instance.

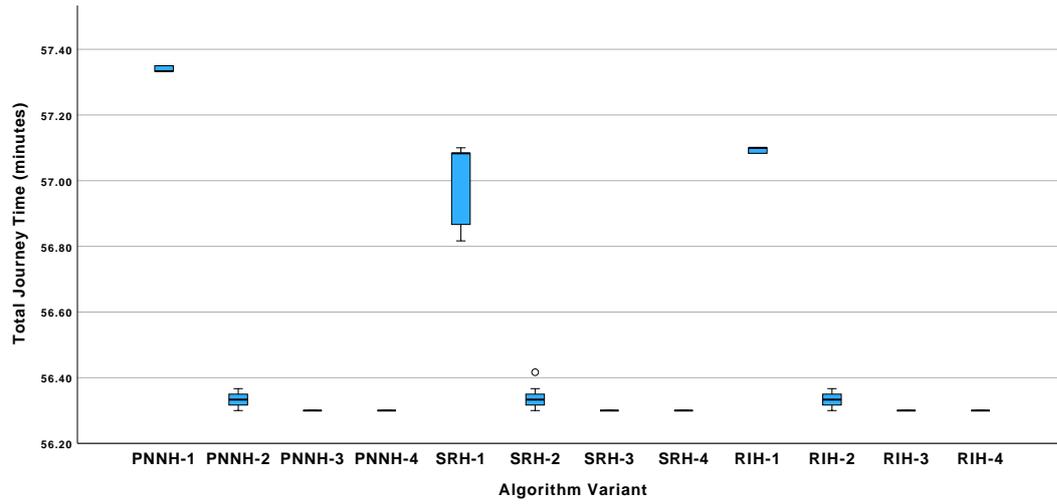


Figure B.2: Total journey time values for the Mellieña instance.

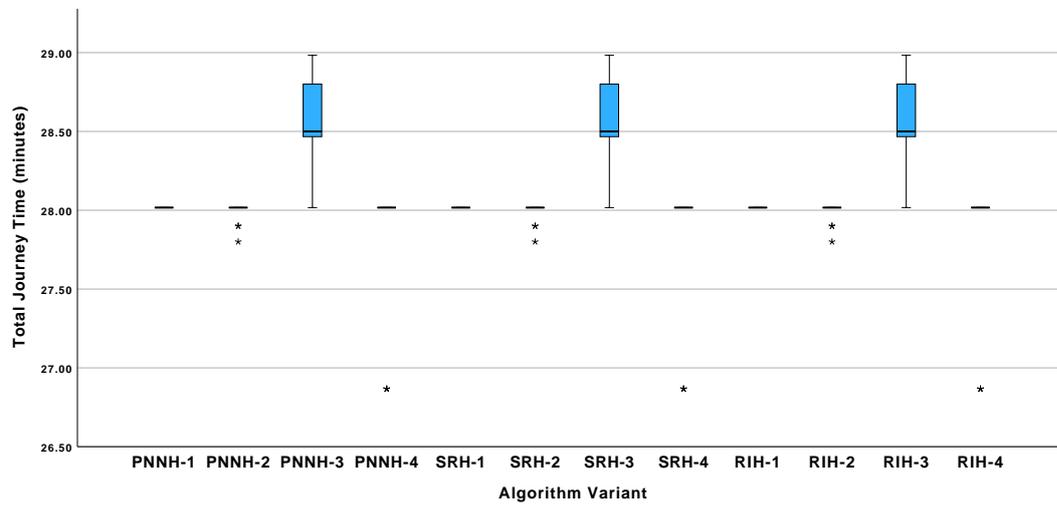


Figure B.3: Total journey time values for the Porthcawl instance.

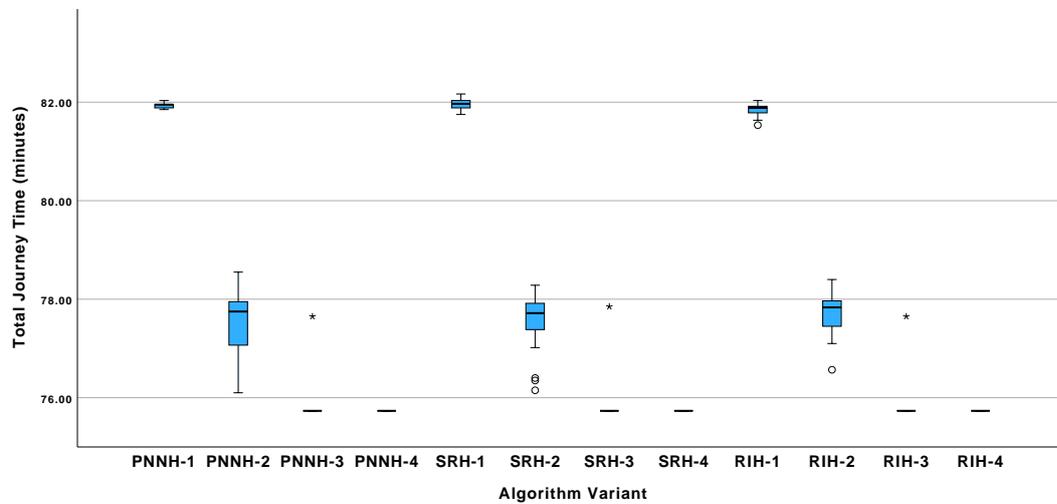


Figure B.4: Total journey time values for the Qrendi instance.

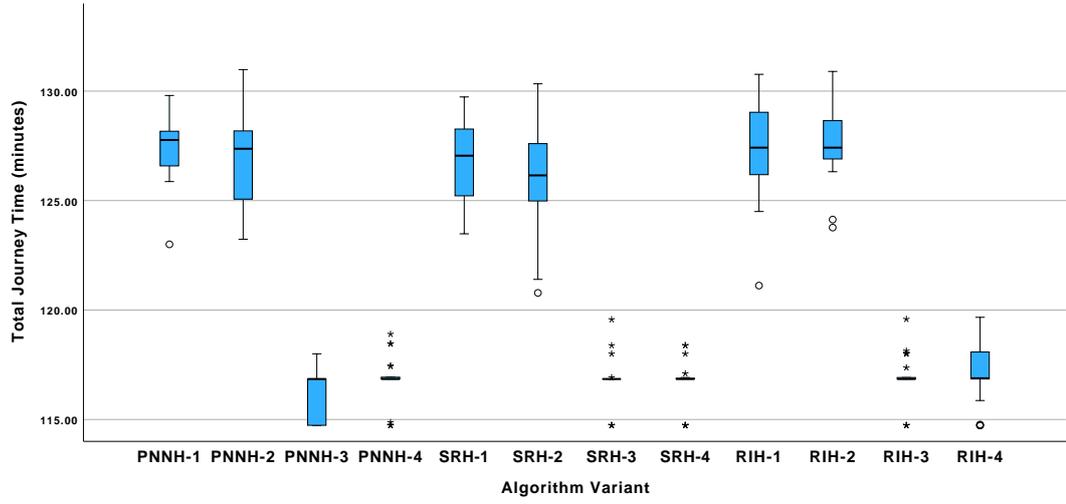


Figure B.5: Total journey time values for the Suffolk instance.

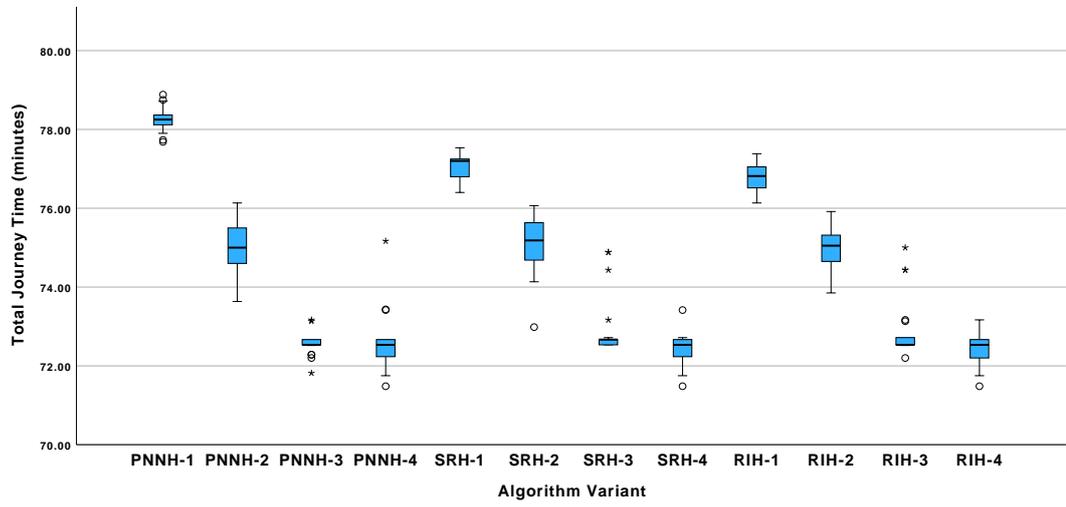


Figure B.6: Total journey time values for the Senglea instance.

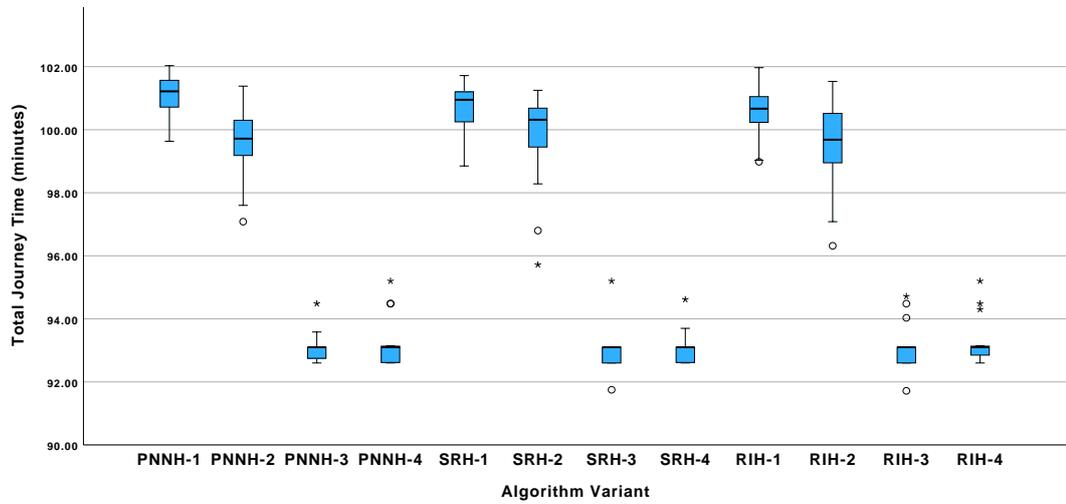


Figure B.7: Total journey time values for the Victoria instance.

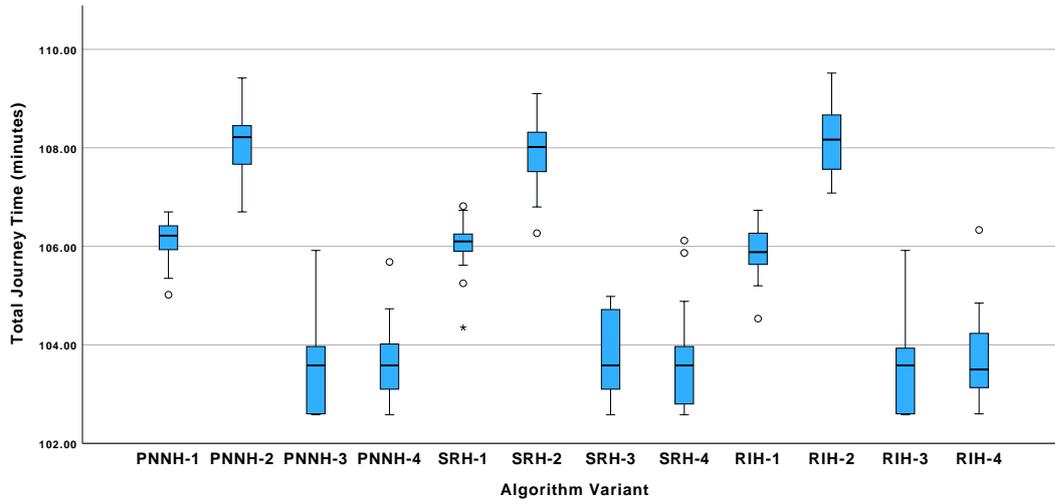


Figure B.8: Total journey time values for the Pembroke instance.

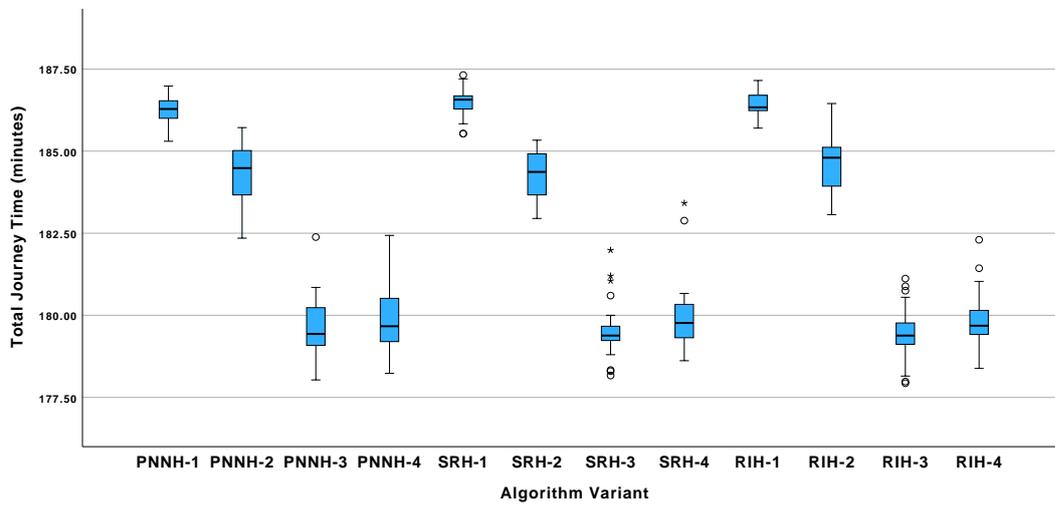


Figure B.9: Total journey time values for the Canberra instance.

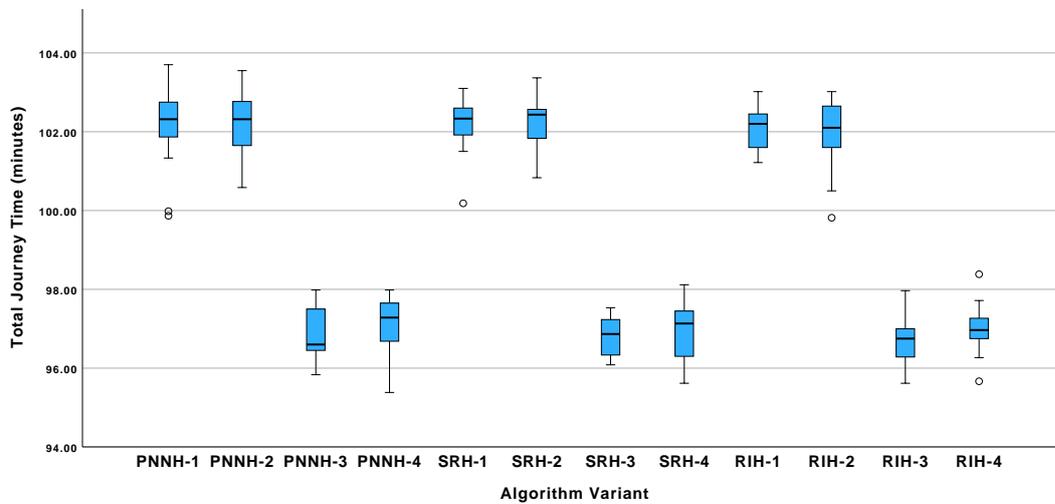


Figure B.10: Total journey time values for the Handaq instance.

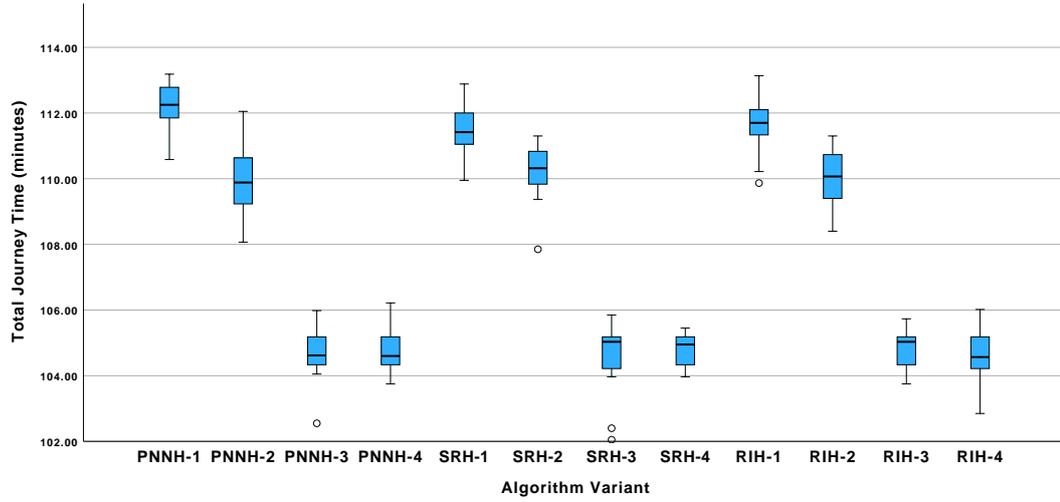


Figure B.11: Total journey time values for the Valletta instance.

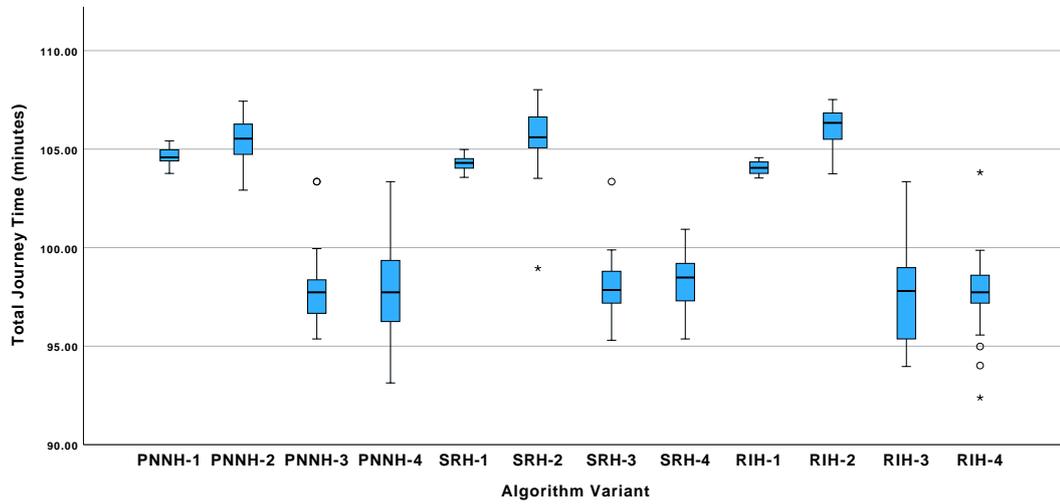


Figure B.12: Total journey time values for the Birkirkara instance.

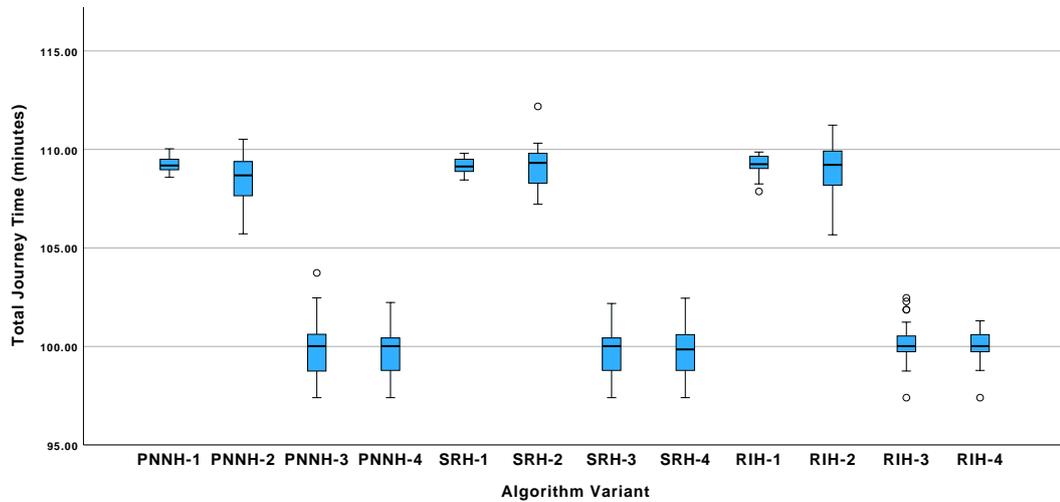


Figure B.13: Total journey time values for the Hamrun instance.

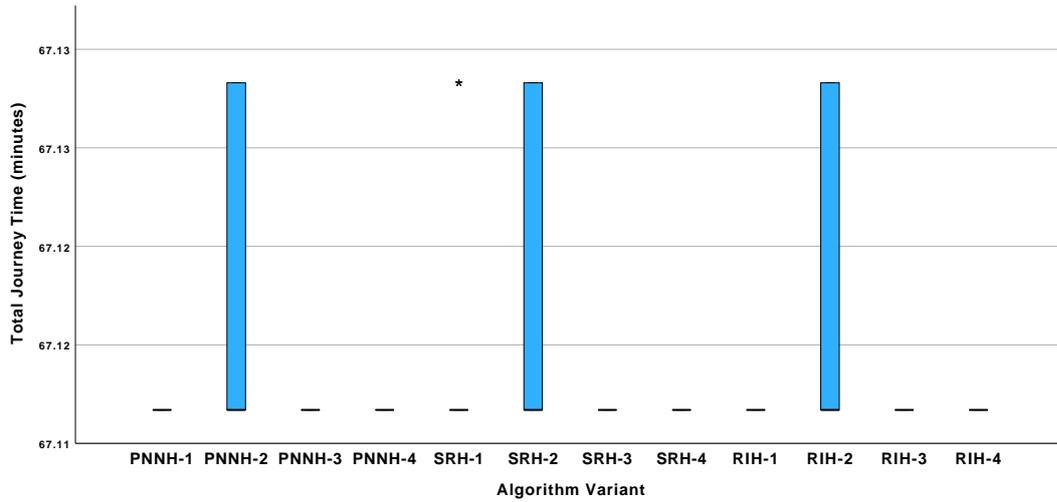


Figure B.14: Total journey time values for the Cardiff instance.

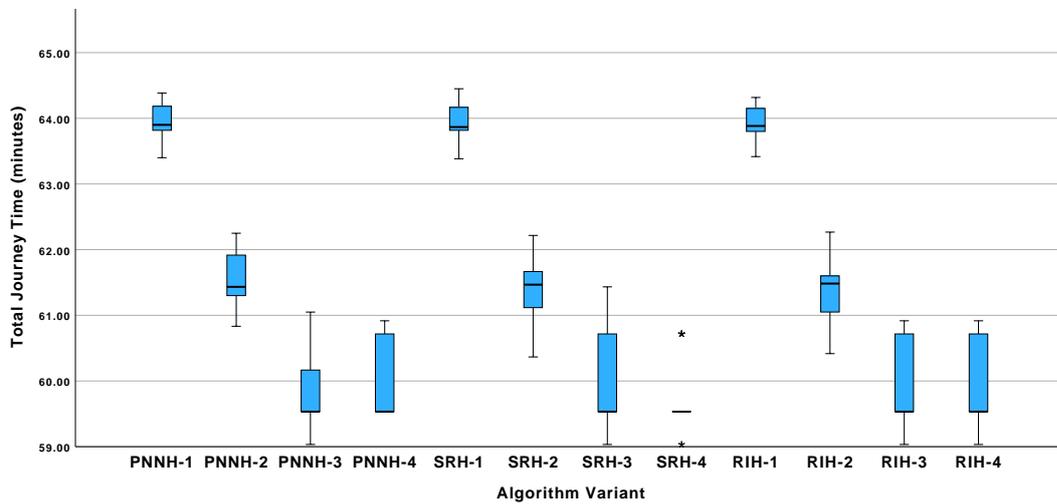


Figure B.15: Total journey time values for the Milton Keynes instance.

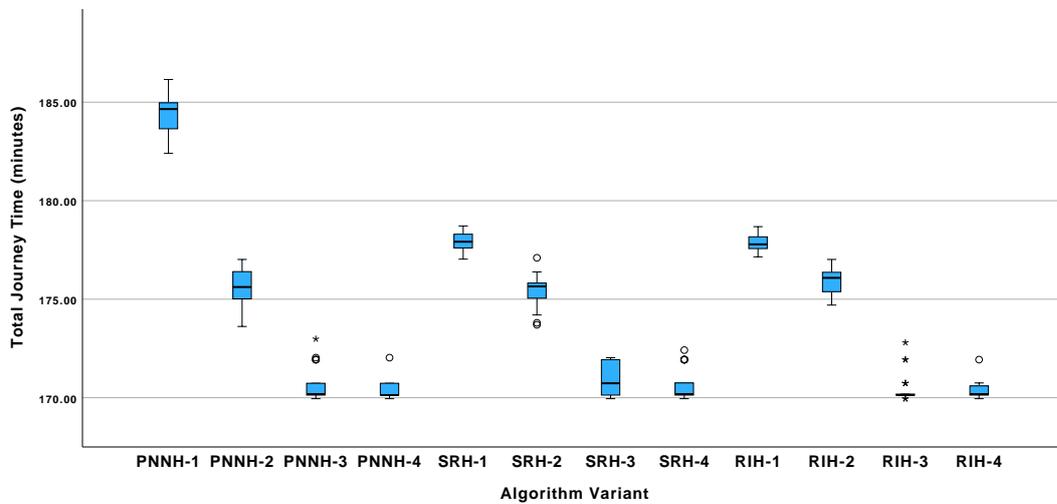


Figure B.16: Total journey time values for the Bridgend instance.

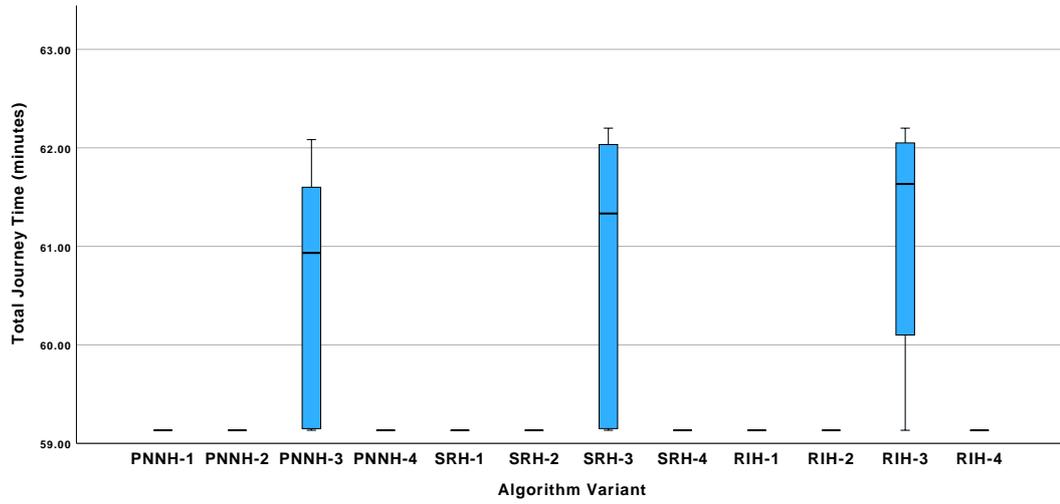


Figure B.17: Total journey time values for the Edinburgh-2 instance.

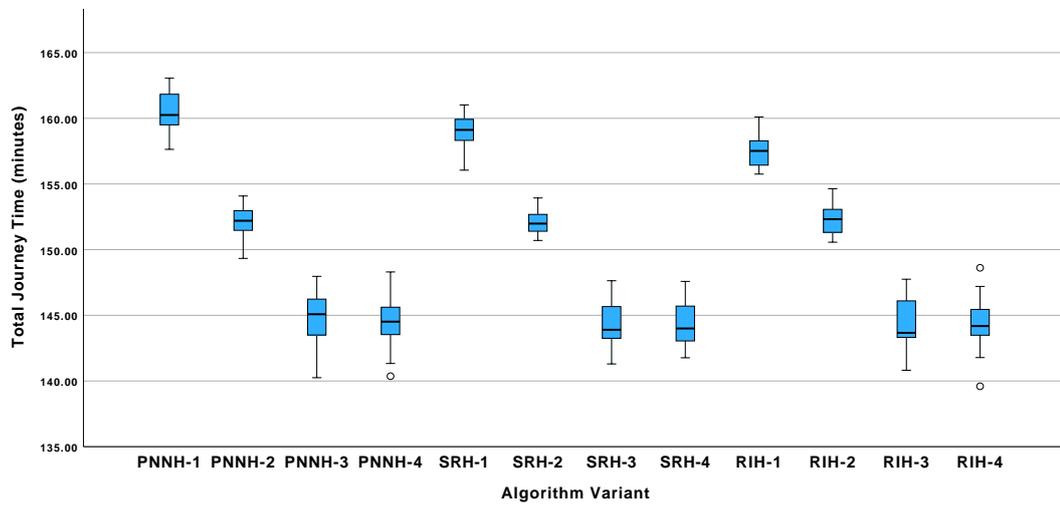


Figure B.18: Total journey time values for the Edinburgh-1 instance.

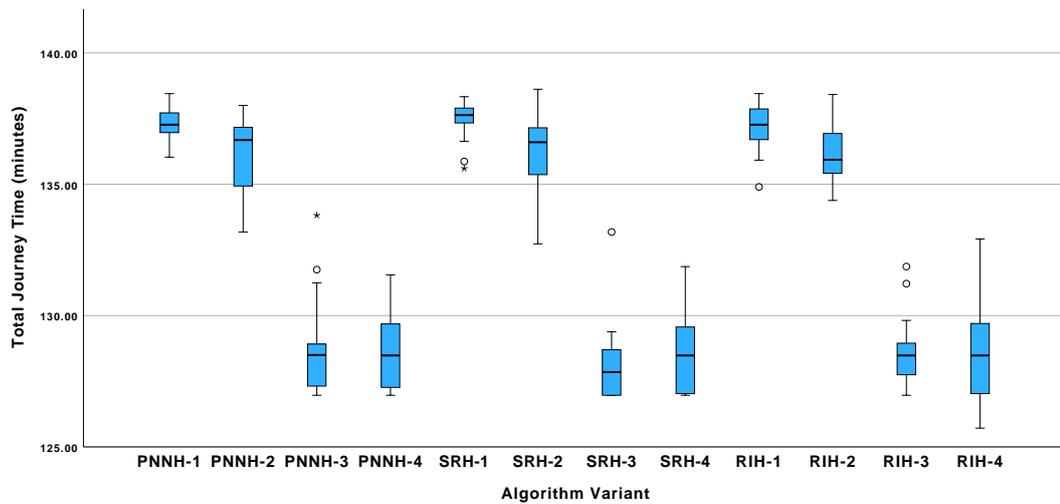


Figure B.19: Total journey time values for the Adelaide instance.

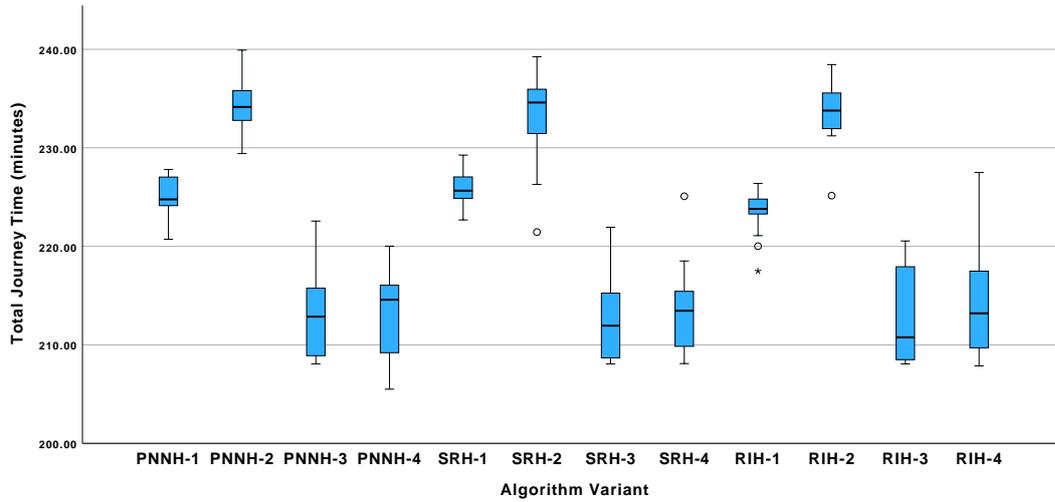


Figure B.20: Total journey time values for the Brisbane instance.

## B.2 Independent-Samples Kruskal-Wallis Tests

The results of the independent-samples Kruskal-Wallis tests for comparing the different algorithm variants are shown in Table B.1.

Table B.1: Kruskal-Wallis test results for comparing the algorithm variants.

Location	Test Statistic	$p$ -value
Mgarr	242.129	< 0.001
Mellieħa	286.480	< 0.001
Porthcawl	214.719	< 0.001
Qrendi	283.537	< 0.001
Suffolk	233.343	< 0.001
Senglea	255.983	< 0.001
Victoria	241.527	< 0.001
Pembroke	245.421	< 0.001
Canberra	253.046	< 0.001
Handaq	226.872	< 0.001
Valletta	248.176	< 0.001
Birkirkara	242.176	< 0.001
Hamrun	226.256	< 0.001
Cardiff	59.508	< 0.001
Milton Keynes	249.561	< 0.001
Bridgend	259.293	< 0.001
Edinburgh-2	255.067	< 0.001
Edinburgh-1	255.006	< 0.001
Adelaide	233.434	< 0.001
Brisbane	246.467	< 0.001

### B.3 Average Results for the Other Algorithm Variants

All results presented in this section are averaged across 25 runs and presented as mean  $\pm$  standard deviation. Time gives the computational time (seconds), FR stands for Feasibility Rate (percentage of iterations yielding a solution with all route journey times being at most 45 minutes), and TJT stands for Total Journey Time (minutes). Bold TJT values represent the best averages across all perturbation mechanisms of the respective constructive heuristic.

Table B.2: PM 3 results for SRH.

Location	$T$	$\underline{L}$	$k$	SRH-3		
				Time (s)	FR (%)	TJT (mins.)
Mgarr	85,450	4	4	175.86 $\pm$ 2.94	100.00 $\pm$ 0.00	54.10 $\pm$ 0.00
Mellieħa	85,450	4	4	157.72 $\pm$ 30.57	100.00 $\pm$ 0.00	<b>56.30</b> $\pm$ 0.00
Porthcawl	284,260	1	1	687.35 $\pm$ 18.69	100.00 $\pm$ 0.00	28.56 $\pm$ 0.29
Qrendi	72,300	5	5	192.04 $\pm$ 5.25	100.00 $\pm$ 0.00	75.82 $\pm$ 0.42
Suffolk	9,380	3	3	77.45 $\pm$ 11.77	31.00 $\pm$ 5.91	116.81 $\pm$ 1.01
Senglea	70,750	6	6	241.47 $\pm$ 11.33	100.00 $\pm$ 0.00	72.88 $\pm$ 0.72
Victoria	5,980	4	4	176.38 $\pm$ 8.19	100.00 $\pm$ 0.00	<b>92.93</b> $\pm$ 0.58
Pembroke	32,610	7	7	226.94 $\pm$ 20.28	100.00 $\pm$ 0.00	<b>103.71</b> $\pm$ 0.83
Canberra	6,610	7	7	110.78 $\pm$ 8.43	100.00 $\pm$ 0.00	<b>179.55</b> $\pm$ 0.88
Handaq	26,020	6	6	243.46 $\pm$ 18.32	100.00 $\pm$ 0.00	<b>96.84</b> $\pm$ 0.47
Valletta	32,090	6	6	325.98 $\pm$ 24.13	100.00 $\pm$ 0.00	<b>104.65</b> $\pm$ 0.87
Birkirkara	43,580	6	6	342.95 $\pm$ 39.98	100.00 $\pm$ 0.00	<b>97.94</b> $\pm$ 1.69
Hamrun	21,620	7	7	259.38 $\pm$ 73.14	100.00 $\pm$ 0.00	99.76 $\pm$ 1.21
Cardiff	82,320	2	2	799.25 $\pm$ 94.06	74.72 $\pm$ 0.37	<b>67.12</b> $\pm$ 0.00
Milton Keynes	63,740	4	4	549.51 $\pm$ 34.30	100.00 $\pm$ 0.00	59.98 $\pm$ 0.62
Bridgend	21,260	5	6	221.82 $\pm$ 3.29	99.89 $\pm$ 0.05	170.86 $\pm$ 0.84
Edinburgh-2	19,430	4	4	543.50 $\pm$ 102.35	100.00 $\pm$ 0.00	60.85 $\pm$ 1.24
Edinburgh-1	31,390	9	9	298.82 $\pm$ 59.51	100.00 $\pm$ 0.00	<b>144.26</b> $\pm$ 1.82
Adelaide	25,070	8	8	328.39 $\pm$ 96.69	100.00 $\pm$ 0.00	<b>128.06</b> $\pm$ 1.35
Brisbane	15,120	10	10	327.15 $\pm$ 53.54	100.00 $\pm$ 0.00	<b>212.43</b> $\pm$ 4.08

Table B.3: PM 1 and PM 2 results for SRH.

Location	$T$	$k$	$k$	SRH-1				SRH-2			
				Time (s)	FR (%)	TJT (mins.)	Time (s)	FR (%)	TJT (mins.)		
Mgarr	85,450	4	4	524.63 ± 6.65	100.00 ± 0.00	<b>54.03</b> ± 0.02	152.00 ± 1.05	100.00 ± 0.00	54.10 ± 0.00		
Mellieha	85,450	4	4	397.51 ± 130.17	100.00 ± 0.00	56.99 ± 0.11	136.44 ± 27.66	100.00 ± 0.00	56.34 ± 0.02		
Porthcawl	284,260	1	1	637.98 ± 5.69	100.00 ± 0.00	28.02 ± 0.00	606.95 ± 5.86	99.96 ± 0.00	28.00 ± 0.05		
Qrendi	72,300	5	5	897.35 ± 8.01	100.00 ± 0.00	81.96 ± 0.10	180.38 ± 2.05	100.00 ± 0.00	77.57 ± 0.57		
Suffolk	9,380	3	3	190.30 ± 5.79	0.09 ± 0.02	126.80 ± 1.72	50.96 ± 3.21	0.12 ± 0.05	125.99 ± 2.47		
Senglea	70,750	6	6	713.66 ± 18.32	100.00 ± 0.00	77.05 ± 0.33	214.88 ± 5.04	100.00 ± 0.00	75.07 ± 0.71		
Victoria	5,980	4	4	353.00 ± 25.31	100.00 ± 0.00	100.70 ± 0.84	151.94 ± 3.89	100.00 ± 0.00	99.87 ± 1.33		
Pembroke	32,610	7	7	892.37 ± 91.89	100.00 ± 0.00	106.08 ± 0.52	184.27 ± 16.02	100.00 ± 0.00	107.88 ± 0.75		
Canberra	6,610	7	7	238.45 ± 14.59	100.00 ± 0.00	186.48 ± 0.45	91.82 ± 6.36	100.00 ± 0.00	184.29 ± 0.75		
Handaq	26,020	6	6	1,073.82 ± 70.76	100.00 ± 0.00	102.25 ± 0.61	206.73 ± 12.63	100.00 ± 0.00	102.30 ± 0.62		
Valetta	32,090	6	6	1,211.02 ± 16.53	100.00 ± 0.00	111.50 ± 0.70	300.93 ± 3.35	100.00 ± 0.00	110.25 ± 0.76		
Birkirkara	43,580	6	6	1,477.08 ± 369.43	100.00 ± 0.00	104.31 ± 0.35	323.28 ± 18.54	100.00 ± 0.00	105.63 ± 1.78		
Hamrun	21,620	7	7	672.96 ± 86.65	100.00 ± 0.00	109.20 ± 0.37	253.32 ± 72.41	100.00 ± 0.00	109.08 ± 1.15		
Cardiff	82,320	2	2	738.53 ± 187.79	41.88 ± 0.14	<b>67.12</b> ± 0.01	584.87 ± 143.15	54.27 ± 0.20	<b>67.12</b> ± 0.01		
Milton Keynes	63,740	4	4	1,256.53 ± 67.34	100.00 ± 0.00	63.95 ± 0.28	503.58 ± 22.90	100.00 ± 0.00	61.37 ± 0.48		
Bridgend	21,260	5	6	949.93 ± 48.95	88.48 ± 0.28	177.90 ± 0.49	203.63 ± 3.44	94.03 ± 0.33	175.39 ± 0.83		
Edinburgh-2	19,430	4	4	553.59 ± 91.42	100.00 ± 0.00	<b>59.13</b> ± 0.00	550.05 ± 104.00	100.00 ± 0.00	<b>59.13</b> ± 0.00		
Edinburgh-1	31,390	9	9	1,165.58 ± 64.74	100.00 ± 0.00	158.91 ± 1.36	282.87 ± 43.25	100.00 ± 0.00	152.05 ± 0.90		
Adelaide	25,070	8	8	1,156.96 ± 37.42	100.00 ± 0.00	137.45 ± 0.65	301.77 ± 87.11	100.00 ± 0.00	136.20 ± 1.46		
Brisbane	15,120	10	10	1,367.12 ± 35.86	100.00 ± 0.00	225.89 ± 1.78	314.41 ± 48.79	100.00 ± 0.00	233.36 ± 4.20		

Table B.4: PM 1 and PM 2 results for PNNH.

Location	$T$	$k$	$k$	PNNH-1				PNNH-2			
				Time (s)	FR (%)	TJT (mins.)	Time (s)	FR (%)	TJT (mins.)		
Mgarr	85,450	4	4	389.71 ± 92.43	100.00 ± 0.00	<b>54.10</b> ± 0.00	140.87 ± 25.58	100.00 ± 0.00	<b>54.10</b> ± 0.00	54.10 ± 0.00	
Mellieha	43,960	4	4	401.23 ± 22.14	100.00 ± 0.00	57.34 ± 0.01	145.26 ± 6.32	100.00 ± 0.00	56.34 ± 0.02	56.34 ± 0.02	
Porthcawl	284,260	1	1	509.52 ± 5.50	100.00 ± 0.00	28.02 ± 0.00	487.68 ± 5.22	99.96 ± 0.00	28.00 ± 0.05	28.00 ± 0.05	
Qrendi	72,300	5	5	776.26 ± 201.66	100.00 ± 0.00	81.93 ± 0.05	188.73 ± 20.64	100.00 ± 0.00	77.47 ± 0.70	77.47 ± 0.70	
Suffolk	9,380	3	3	133.01 ± 1.61	0.08 ± 0.03	127.48 ± 1.40	51.57 ± 2.80	0.13 ± 0.05	126.85 ± 2.05	126.85 ± 2.05	
Senglea	70,750	6	6	644.37 ± 19.89	100.00 ± 0.00	78.26 ± 0.30	209.32 ± 4.32	100.00 ± 0.00	75.01 ± 0.60	75.01 ± 0.60	
Victoria	5,980	4	4	206.41 ± 14.73	100.00 ± 0.00	101.15 ± 0.57	119.14 ± 20.33	100.00 ± 0.00	99.57 ± 1.07	99.57 ± 1.07	
Pembroke	32,610	7	7	579.61 ± 183.27	100.00 ± 0.00	106.13 ± 0.42	232.82 ± 21.42	100.00 ± 0.00	108.10 ± 0.69	108.10 ± 0.69	
Canberra	6,610	7	7	194.19 ± 20.10	100.00 ± 0.00	186.22 ± 0.43	75.14 ± 7.13	100.00 ± 0.00	184.38 ± 0.90	184.38 ± 0.90	
Handaq	26,020	6	6	823.35 ± 206.89	100.00 ± 0.00	102.21 ± 0.89	241.09 ± 7.87	100.00 ± 0.00	102.20 ± 0.86	102.20 ± 0.86	
Valletta	32,090	6	6	925.04 ± 167.37	100.00 ± 0.00	112.13 ± 0.72	264.56 ± 12.11	100.00 ± 0.00	109.96 ± 0.97	109.96 ± 0.97	
Birkirkara	43,580	6	6	1,471.31 ± 374.81	100.00 ± 0.00	104.62 ± 0.42	385.44 ± 10.14	100.00 ± 0.00	105.59 ± 1.14	105.59 ± 1.14	
Hamrun	21,620	7	7	634.53 ± 25.49	100.00 ± 0.00	109.24 ± 0.35	221.10 ± 55.05	100.00 ± 0.00	108.46 ± 1.28	108.46 ± 1.28	
Cardiff	82,320	2	2	988.58 ± 135.14	87.36 ± 0.11	<b>67.12</b> ± 0.00	549.94 ± 129.57	54.27 ± 0.20	<b>67.12</b> ± 0.01	<b>67.12</b> ± 0.01	
Milton Keynes	63,740	4	4	1,241.08 ± 22.29	100.00 ± 0.00	63.94 ± 0.28	505.31 ± 16.54	100.00 ± 0.00	61.58 ± 0.40	61.58 ± 0.40	
Bridgend	21,260	5	6	859.30 ± 76.68	86.28 ± 0.25	184.34 ± 1.02	208.01 ± 16.83	94.00 ± 0.24	175.61 ± 0.92	175.61 ± 0.92	
Edinburgh-2	19,430	4	4	612.15 ± 26.60	100.00 ± 0.00	<b>59.13</b> ± 0.00	613.75 ± 24.71	100.00 ± 0.00	<b>59.13</b> ± 0.00	<b>59.13</b> ± 0.00	
Edinburgh-1	31,390	9	9	1,600.48 ± 550.43	100.00 ± 0.00	160.57 ± 1.52	400.42 ± 115.61	100.00 ± 0.00	152.11 ± 1.19	152.11 ± 1.19	
Adelaide	25,070	8	8	1,351.36 ± 160.57	100.00 ± 0.00	137.31 ± 0.62	469.46 ± 32.27	100.00 ± 0.00	136.17 ± 1.44	136.17 ± 1.44	
Brisbane	15,120	10	10	1,555.30 ± 84.00	100.00 ± 0.00	225.23 ± 1.92	442.11 ± 13.20	100.00 ± 0.00	234.52 ± 2.58	234.52 ± 2.58	

Table B.5: PM 3 and PM 4 results for PNNH.

Location	$T$	$k$	$k$	PNNH-3			PNNH-4		
				Time (s)	FR (%)	TJT (mins.)	Time (s)	FR (%)	TJT (mins.)
Mgarr	85,450	4	4	168.74 ± 11.20	100.00 ± 0.00	54.10 ± 0.00	154.45 ± 25.55	100.00 ± 0.00	54.10 ± 0.00
Mellieħa	43,960	4	4	164.19 ± 8.27	100.00 ± 0.00	56.30 ± 0.00	157.80 ± 8.46	100.00 ± 0.00	56.30 ± 0.00
Porthcawl	284,260	1	1	551.17 ± 16.36	100.00 ± 0.00	28.56 ± 0.29	512.07 ± 11.71	99.98 ± 0.03	27.88 ± 0.38
Qrendi	72,300	5	5	200.03 ± 22.19	100.00 ± 0.00	75.81 ± 0.38	192.37 ± 20.52	100.00 ± 0.00	75.73 ± 0.00
Suffolk	9,380	3	3	77.58 ± 9.46	28.40 ± 8.11	116.14 ± 1.10	71.64 ± 6.64	18.65 ± 5.63	116.79 ± 1.07
Senglea	70,750	6	6	237.92 ± 11.09	100.00 ± 0.00	72.55 ± 0.26	228.82 ± 6.47	100.00 ± 0.00	72.58 ± 0.69
Victoria	5,980	4	4	136.88 ± 25.08	100.00 ± 0.00	93.04 ± 0.39	128.97 ± 22.28	100.00 ± 0.00	93.16 ± 0.64
Pembroke	32,610	7	7	272.74 ± 34.70	100.00 ± 0.00	103.53 ± 0.91	253.54 ± 39.30	100.00 ± 0.00	103.63 ± 0.76
Canberra	6,610	7	7	92.86 ± 5.09	100.00 ± 0.00	179.63 ± 0.96	88.74 ± 3.93	100.00 ± 0.00	179.89 ± 0.90
Handaq	26,020	6	6	291.10 ± 23.01	100.00 ± 0.00	96.91 ± 0.68	268.73 ± 14.81	100.00 ± 0.00	97.13 ± 0.65
Valletta	32,090	6	6	284.97 ± 18.32	100.00 ± 0.00	104.69 ± 0.66	277.89 ± 15.53	100.00 ± 0.00	104.83 ± 0.68
Birkirkara	43,580	6	6	427.89 ± 20.61	100.00 ± 0.00	97.89 ± 2.02	418.10 ± 15.92	100.00 ± 0.00	97.85 ± 2.49
Hamrun	21,620	7	7	239.36 ± 61.28	100.00 ± 0.00	99.85 ± 1.67	234.67 ± 57.19	100.00 ± 0.00	99.67 ± 1.21
Cardiff	82,320	2	2	637.26 ± 121.39	74.73 ± 0.34	67.12 ± 0.00	599.36 ± 136.41	66.82 ± 0.33	67.12 ± 0.00
Milton Keynes	63,740	4	4	549.40 ± 28.69	100.00 ± 0.00	59.86 ± 0.57	543.58 ± 21.75	100.00 ± 0.00	59.95 ± 0.57
Bridgend	21,260	5	6	227.47 ± 17.84	99.88 ± 0.06	170.65 ± 0.83	220.99 ± 19.54	99.38 ± 0.16	170.43 ± 0.45
Edinburgh-2	19,430	4	4	607.94 ± 27.68	100.00 ± 0.00	60.70 ± 1.18	607.63 ± 25.95	100.00 ± 0.00	59.13 ± 0.00
Edinburgh-1	31,390	9	9	407.29 ± 122.47	100.00 ± 0.00	144.64 ± 2.12	394.33 ± 108.81	100.00 ± 0.00	144.49 ± 1.81
Adelaide	25,070	8	8	514.95 ± 39.93	100.00 ± 0.00	128.73 ± 1.65	509.19 ± 30.01	100.00 ± 0.00	128.57 ± 1.38
Brisbane	15,120	10	10	472.59 ± 14.09	100.00 ± 0.00	213.25 ± 4.46	466.00 ± 10.22	100.00 ± 0.00	213.33 ± 4.11

Table B.6: PM 1 and PM 2 results for RIH.

Location	$T$	$k$	$k$	RIH-1				RIH-2			
				Time (s)	FR (%)	TJT (mins.)	Time (s)	FR (%)	TJT (mins.)		
Mgarr	85,450	4	4	134.99 ± 3.11	100.00 ± 0.00	<b>54.09</b> ± 0.03	69.01 ± 2.21	100.00 ± 0.00	54.10 ± 0.00		
Mellieħa	43,960	4	4	134.84 ± 2.41	100.00 ± 0.00	57.09 ± 0.01	67.85 ± 0.49	100.00 ± 0.00	56.33 ± 0.02		
Porthcawl	284,260	1	1	242.96 ± 2.24	100.00 ± 0.00	28.02 ± 0.00	233.51 ± 2.81	99.96 ± 0.00	28.00 ± 0.05		
Qrendi	72,300	5	5	386.76 ± 1.84	100.00 ± 0.00	81.85 ± 0.12	83.63 ± 0.45	100.00 ± 0.00	77.70 ± 0.45		
Suffolk	9,380	3	3	95.84 ± 1.06	0.09 ± 0.03	127.43 ± 2.11	48.10 ± 0.94	0.11 ± 0.05	127.63 ± 1.64		
Senglea	70,750	6	6	310.95 ± 7.36	100.00 ± 0.00	76.77 ± 0.35	96.01 ± 2.21	100.00 ± 0.00	74.97 ± 0.52		
Victoria	5,980	4	4	117.78 ± 2.49	100.00 ± 0.00	100.62 ± 0.76	72.05 ± 6.44	100.00 ± 0.00	99.64 ± 1.31		
Pembroke	32,610	7	7	365.36 ± 15.31	100.00 ± 0.00	105.88 ± 0.50	92.03 ± 4.28	100.00 ± 0.00	108.15 ± 0.68		
Canberra	6,610	7	7	122.44 ± 0.87	100.00 ± 0.00	186.42 ± 0.36	43.22 ± 0.58	100.00 ± 0.00	184.61 ± 0.78		
Handaq	26,020	6	6	383.90 ± 7.36	100.00 ± 0.00	102.09 ± 0.57	94.62 ± 1.00	100.00 ± 0.00	101.99 ± 0.81		
Valletta	32,090	6	6	458.06 ± 8.39	100.00 ± 0.00	111.66 ± 0.77	142.37 ± 1.27	100.00 ± 0.00	109.99 ± 0.81		
Birkirkara	43,580	6	6	809.01 ± 22.96	100.00 ± 0.00	104.08 ± 0.32	148.13 ± 3.76	100.00 ± 0.00	106.18 ± 0.92		
Hamrun	21,620	7	7	431.67 ± 11.45	100.00 ± 0.00	109.21 ± 0.48	164.84 ± 25.35	100.00 ± 0.00	108.87 ± 1.44		
Cardiff	82,320	2	2	492.02 ± 34.58	86.13 ± 0.13	<b>67.12</b> ± 0.00	357.55 ± 32.18	54.27 ± 0.20	<b>67.12</b> ± 0.01		
Milton Keynes	63,740	4	4	559.16 ± 15.87	100.00 ± 0.00	63.92 ± 0.28	226.47 ± 6.42	100.00 ± 0.00	61.38 ± 0.49		
Bridgend	21,260	5	6	823.35 ± 57.81	91.33 ± 0.24	177.83 ± 0.42	203.54 ± 1.45	93.98 ± 0.34	175.92 ± 0.64		
Edinburgh-2	19,430	4	4	303.62 ± 3.16	100.00 ± 0.00	<b>59.13</b> ± 0.00	304.82 ± 7.24	100.00 ± 0.00	<b>59.13</b> ± 0.00		
Edinburgh-1	31,390	9	9	1,114.85 ± 14.25	100.00 ± 0.00	157.42 ± 1.19	246.45 ± 1.94	100.00 ± 0.00	152.30 ± 1.14		
Adelaide	25,070	8	8	1,138.61 ± 31.24	100.00 ± 0.00	137.22 ± 0.87	234.98 ± 8.99	100.00 ± 0.00	136.18 ± 1.14		
Brisbane	15,120	10	10	1,363.16 ± 50.87	100.00 ± 0.00	223.78 ± 1.97	227.36 ± 9.14	100.00 ± 0.00	233.59 ± 2.79		

Table B.7: PM 3 and PM 4 results for RIH.

Location	$T$	$k$	$k$	RIH-3			RIH-4		
				Time (s)	FR (%)	TJT (mins.)	Time (s)	FR (%)	TJT (mins.)
Mgarr	85,450	4	4	78.87 ± 2.31	100.00 ± 0.00	54.10 ± 0.00	75.50 ± 2.28	100.00 ± 0.00	54.10 ± 0.00
Mellieha	43,960	4	4	76.28 ± 0.43	100.00 ± 0.00	<b>56.30</b> ± 0.00	73.14 ± 0.47	100.00 ± 0.00	<b>56.30</b> ± 0.00
Porthcawl	284,260	1	1	260.63 ± 7.38	100.00 ± 0.00	28.56 ± 0.29	244.35 ± 5.42	99.98 ± 0.03	<b>27.88</b> ± 0.38
Qrendi	72,300	5	5	87.49 ± 2.23	100.00 ± 0.00	75.81 ± 0.38	84.24 ± 0.77	100.00 ± 0.00	<b>75.73</b> ± 0.00
Suffolk	9,380	3	3	71.64 ± 3.68	31.55 ± 5.55	<b>116.92</b> ± 1.06	66.21 ± 3.27	18.82 ± 4.25	117.19 ± 1.42
Senglea	70,750	6	6	108.07 ± 5.96	100.00 ± 0.00	72.86 ± 0.70	104.44 ± 4.06	100.00 ± 0.00	<b>72.43</b> ± 0.39
Victoria	5,980	4	4	81.79 ± 3.48	100.00 ± 0.00	<b>93.03</b> ± 0.62	78.29 ± 3.33	100.00 ± 0.00	93.17 ± 0.61
Pembroke	32,610	7	7	118.03 ± 9.07	100.00 ± 0.00	<b>103.55</b> ± 0.90	111.33 ± 6.59	100.00 ± 0.00	103.74 ± 0.84
Canberra	6,610	7	7	52.67 ± 3.14	100.00 ± 0.00	<b>179.46</b> ± 0.84	50.52 ± 2.24	100.00 ± 0.00	179.87 ± 0.82
Handaq	26,020	6	6	109.25 ± 6.65	100.00 ± 0.00	<b>96.72</b> ± 0.64	103.76 ± 6.12	100.00 ± 0.00	96.98 ± 0.53
Valetta	32,090	6	6	159.44 ± 12.49	100.00 ± 0.00	104.83 ± 0.56	150.41 ± 7.92	100.00 ± 0.00	<b>104.65</b> ± 0.73
Birkirkara	43,580	6	6	162.48 ± 7.33	100.00 ± 0.00	97.78 ± 2.46	156.71 ± 5.98	100.00 ± 0.00	<b>97.56</b> ± 2.16
Hamrun	21,620	7	7	143.04 ± 11.70	100.00 ± 0.00	100.12 ± 1.20	154.00 ± 19.54	100.00 ± 0.00	<b>99.92</b> ± 0.91
Cardiff	82,320	2	2	361.74 ± 39.86	74.73 ± 0.35	<b>67.12</b> ± 0.00	349.72 ± 12.61	66.85 ± 0.33	<b>67.12</b> ± 0.00
Milton Keynes	63,740	4	4	250.59 ± 12.03	100.00 ± 0.00	<b>59.88</b> ± 0.66	245.56 ± 11.44	100.00 ± 0.00	59.96 ± 0.63
Bridgend	21,260	5	6	220.80 ± 1.51	99.88 ± 0.07	170.45 ± 0.71	215.62 ± 1.28	99.41 ± 0.13	<b>170.34</b> ± 0.42
Edinburgh-2	19,430	4	4	300.80 ± 2.25	100.00 ± 0.00	61.10 ± 1.14	302.63 ± 2.36	100.00 ± 0.00	<b>59.13</b> ± 0.00
Edinburgh-1	31,390	9	9	258.28 ± 8.16	100.00 ± 0.00	<b>144.22</b> ± 1.90	256.17 ± 7.76	100.00 ± 0.00	144.53 ± 1.90
Adelaide	25,070	8	8	254.76 ± 8.91	100.00 ± 0.00	<b>128.52</b> ± 1.23	247.39 ± 9.53	100.00 ± 0.00	128.70 ± 1.83
Brisbane	15,120	10	10	240.12 ± 11.13	100.00 ± 0.00	<b>213.01</b> ± 4.79	237.41 ± 9.77	100.00 ± 0.00	214.24 ± 5.17

## B.4 Visualizations of Best SRH-4 Solutions

This section visualizes the best SRH-4 solutions listed in Table 4.11. In Figures B.21 to B.40, the lime dot represents the school, the red dots represent the visited bus stops, and the yellow dots represent the student addresses. Moreover, each route is shown in a different colour. Note that certain subroutes are not visible since they overlap with other subroutes. The following visualizations and further information on arrival times, dwell times, distances covered, and walking/riding times and distances for each student can be viewed from the interactive visualization files at (Sciortino, 2024b).

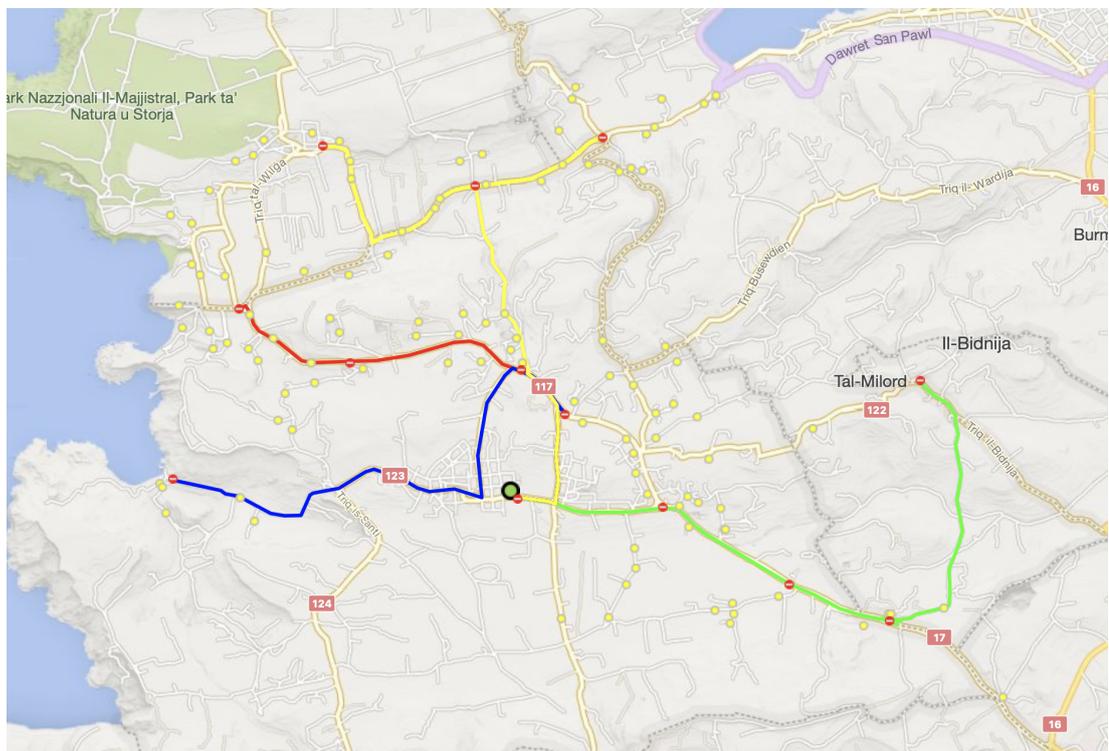


Figure B.21: Mgarr best SRH-4 solution with 4 routes (red, lime, blue, yellow), 190 students, average walk 6.05 minutes, and average journey time 13.53 minutes.

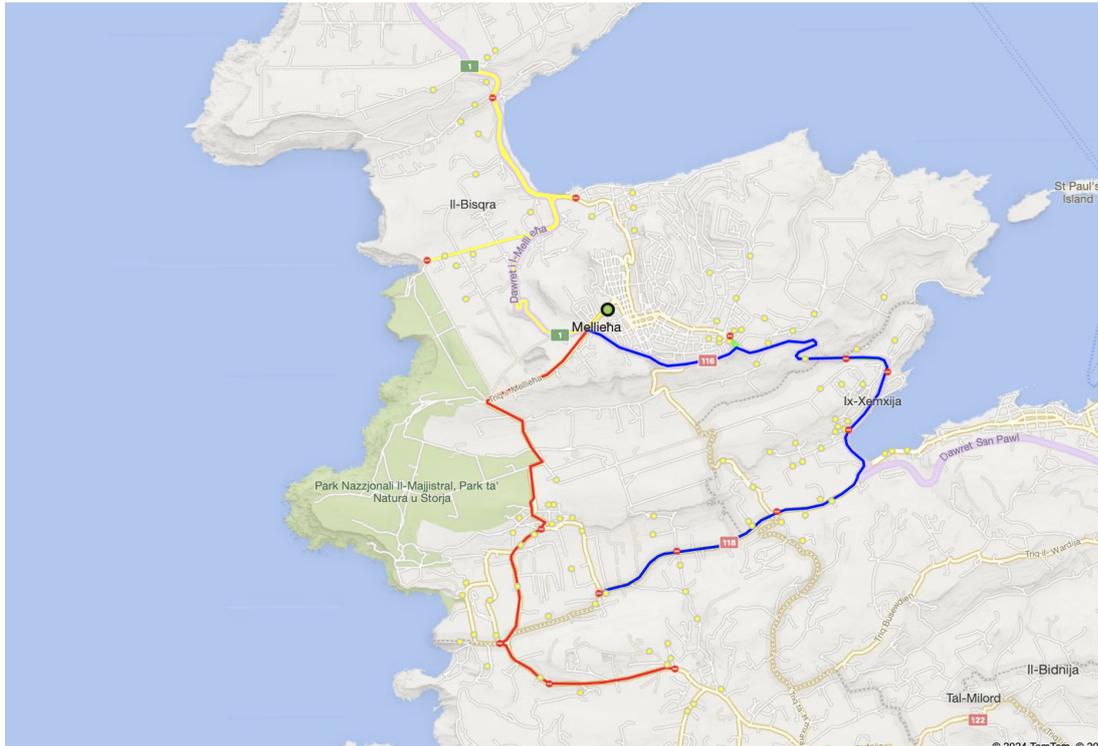


Figure B.22: Mellieħa best SRH-4 solution with 4 routes (red, lime, blue, yellow), 171 students, average walk 5.03 minutes, and average journey time 14.07 minutes.

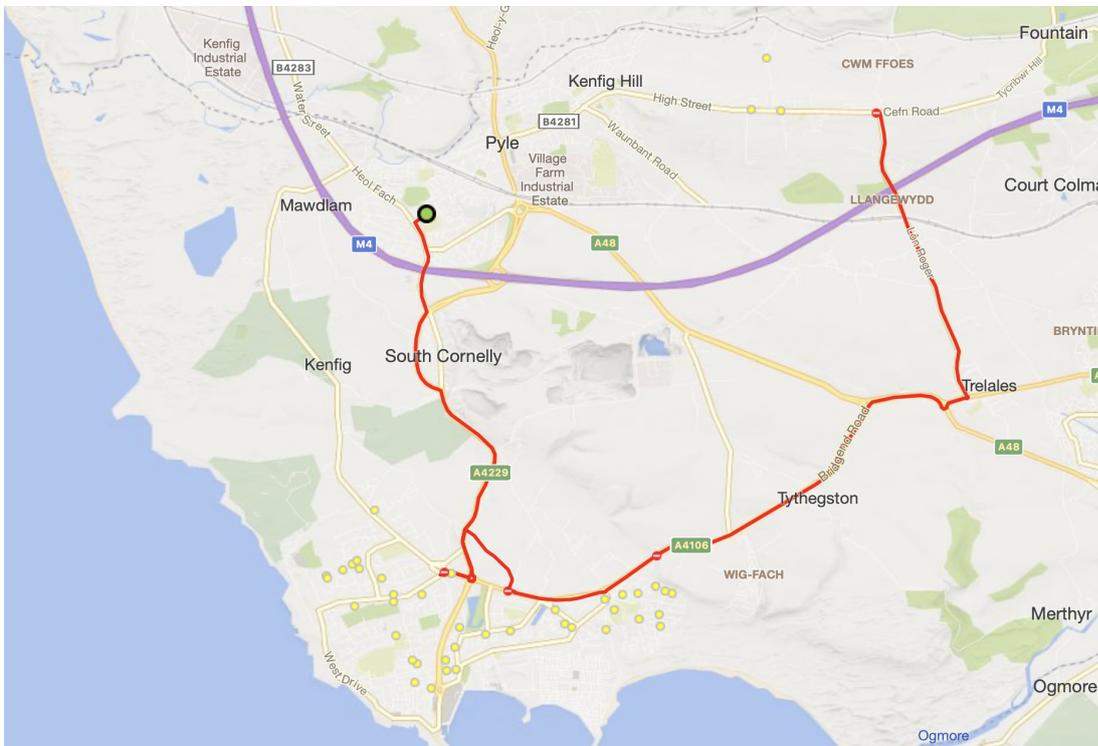


Figure B.23: Porthcawl best SRH-4 solution with 1 route (red), 66 students, average walk 14.26 minutes, and average journey time 26.87 minutes.



APPENDIX B. HEURISTIC ALGORITHM DETERMINISTIC SCENARIO RESULTS

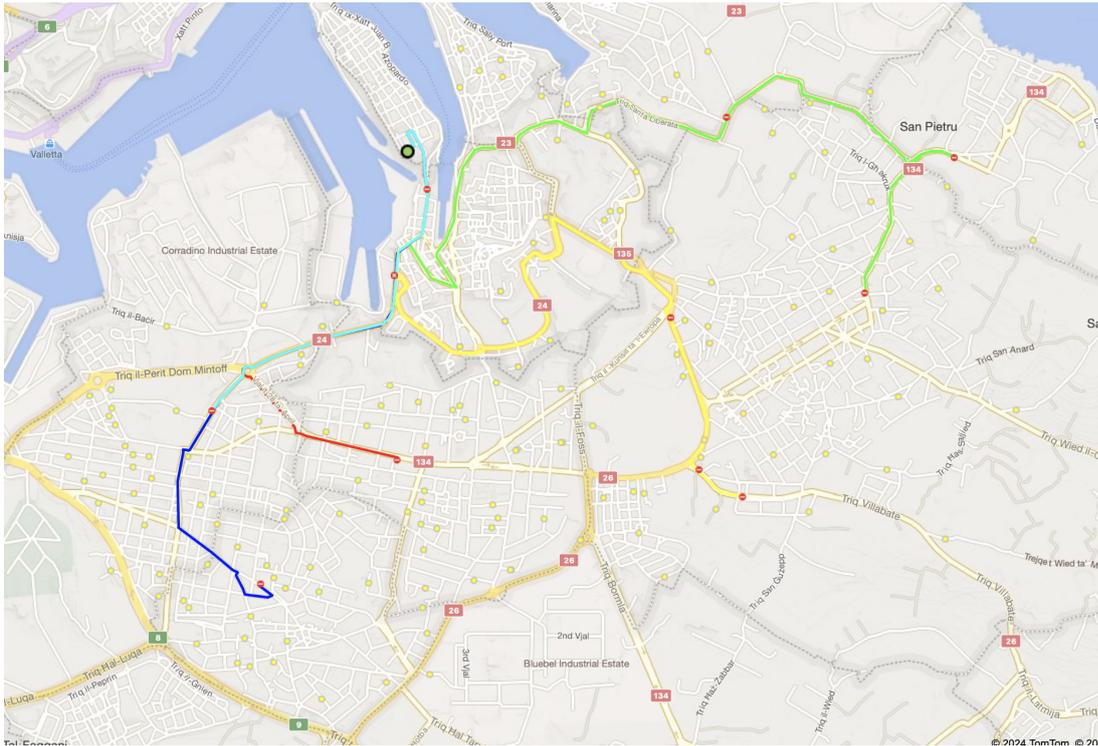


Figure B.26: Senglea best SRH-4 solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 266 students, average walk 6.72 minutes, and average journey time 11.89 minutes.

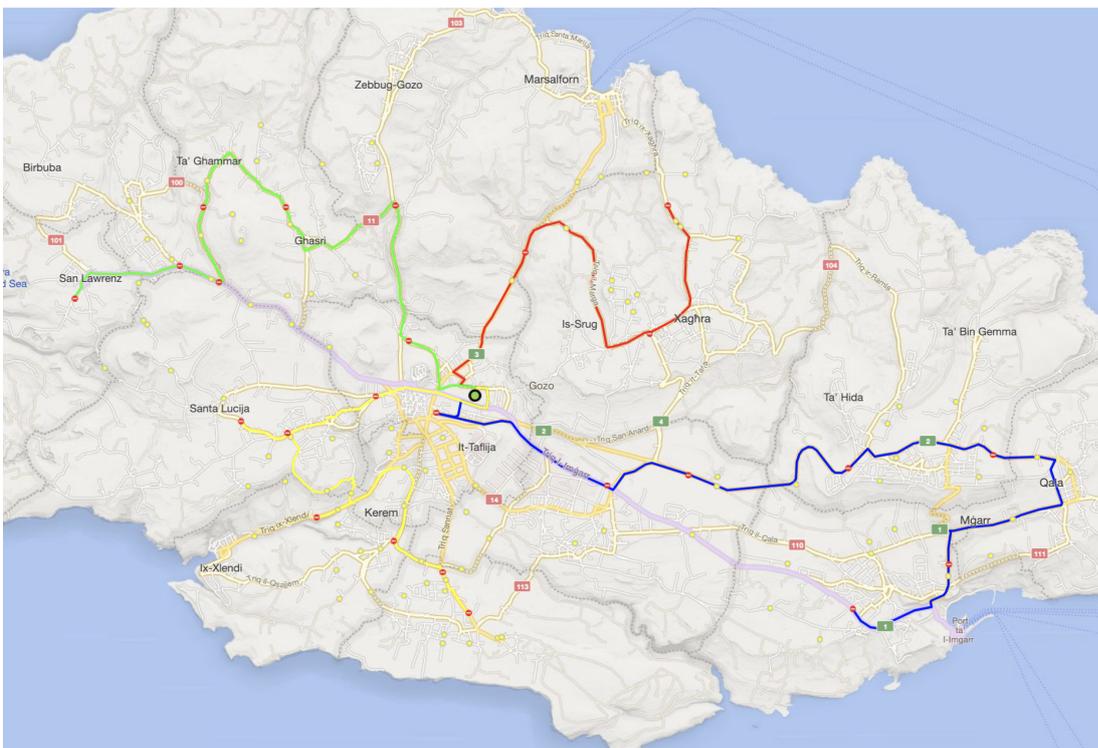


Figure B.27: Victoria best SRH-4 solution with 4 routes (red, lime, blue, yellow), 171 students, average walk 7.42 minutes, and average journey time 23.15 minutes.

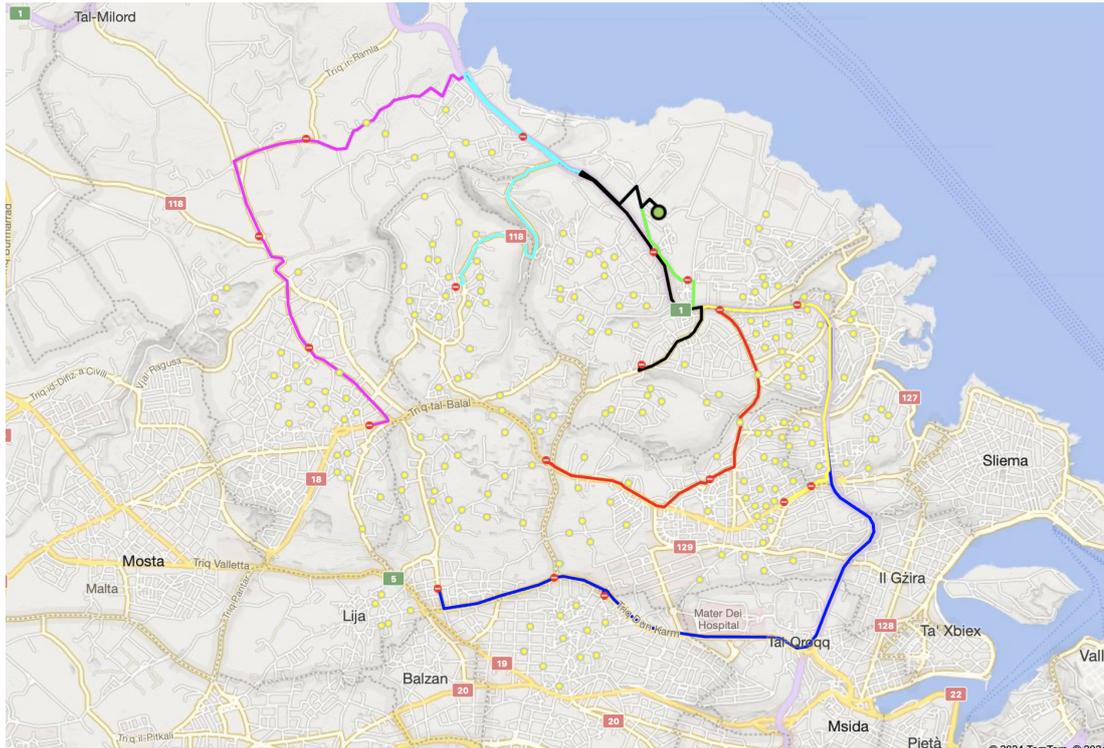


Figure B.28: Pembroke best SRH-4 solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 335 students, average walk 6.86 minutes, and average journey time 14.65 minutes.

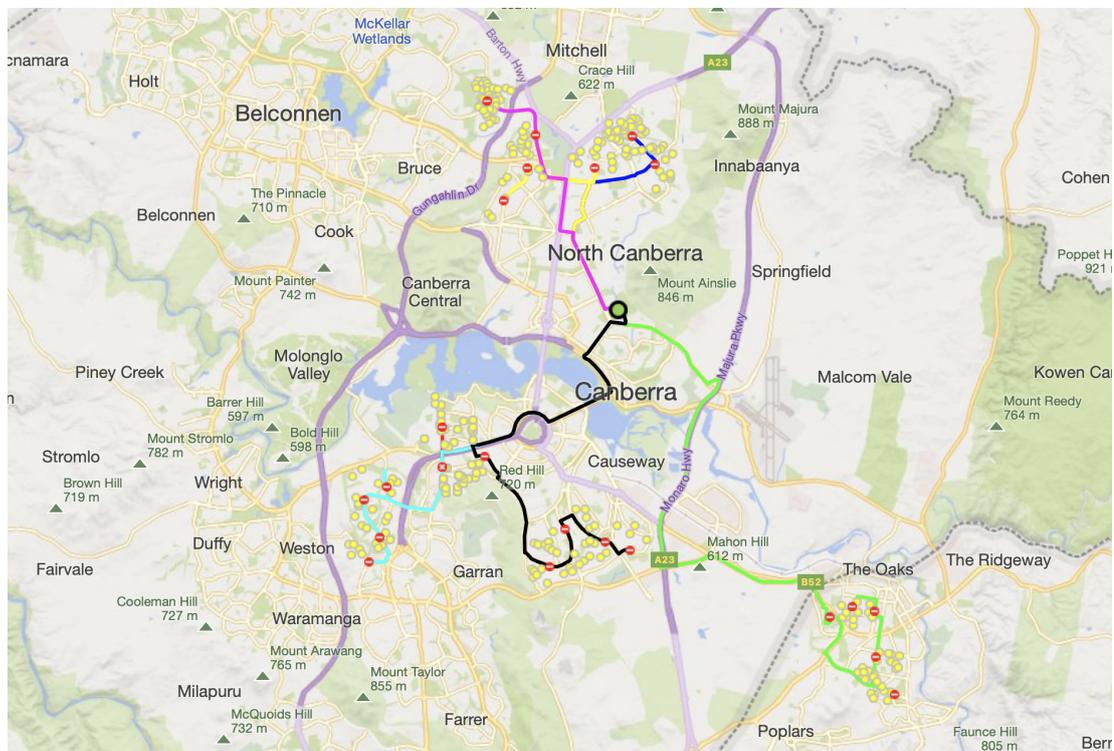


Figure B.29: Canberra best SRH-4 solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 499 students, average walk 7.19 minutes, and average journey time 25.45 minutes.

APPENDIX B. HEURISTIC ALGORITHM DETERMINISTIC SCENARIO RESULTS

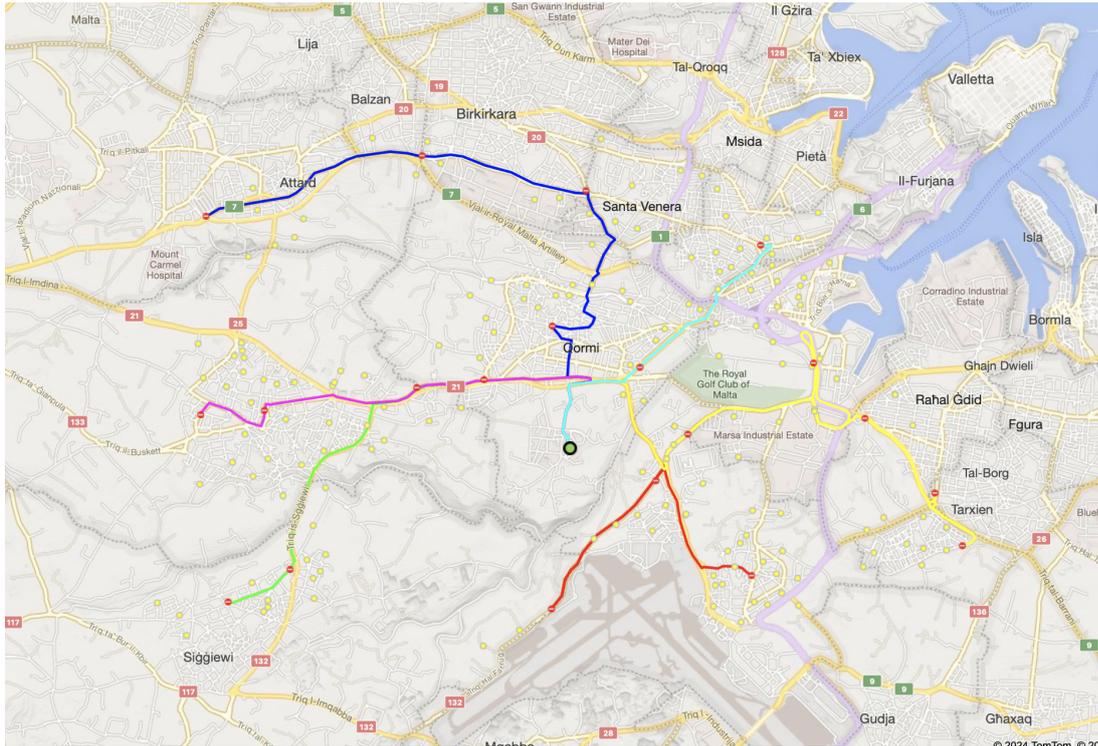


Figure B.30: Handaqa best SRH-4 solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 285 students, average walk 6.97 minutes, and average journey time 15.94 minutes.

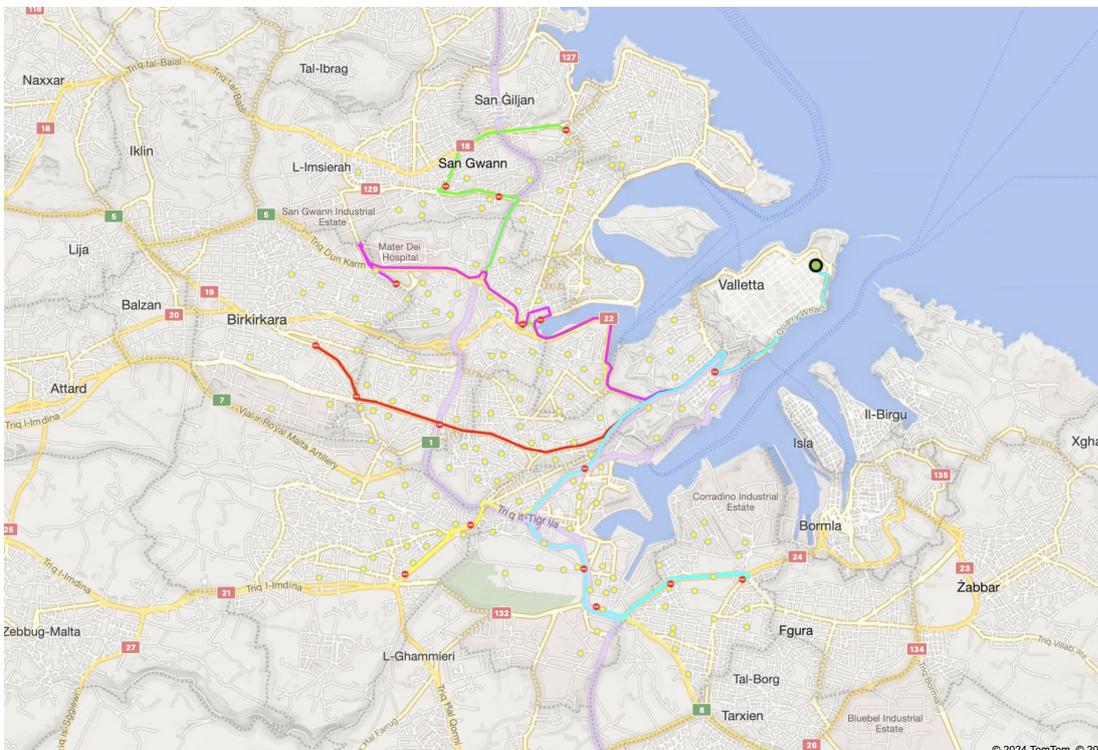


Figure B.31: Valletta best SRH-4 solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 268 students, average walk 6.82 minutes, and average journey time 17.06 minutes.

APPENDIX B. HEURISTIC ALGORITHM DETERMINISTIC SCENARIO RESULTS

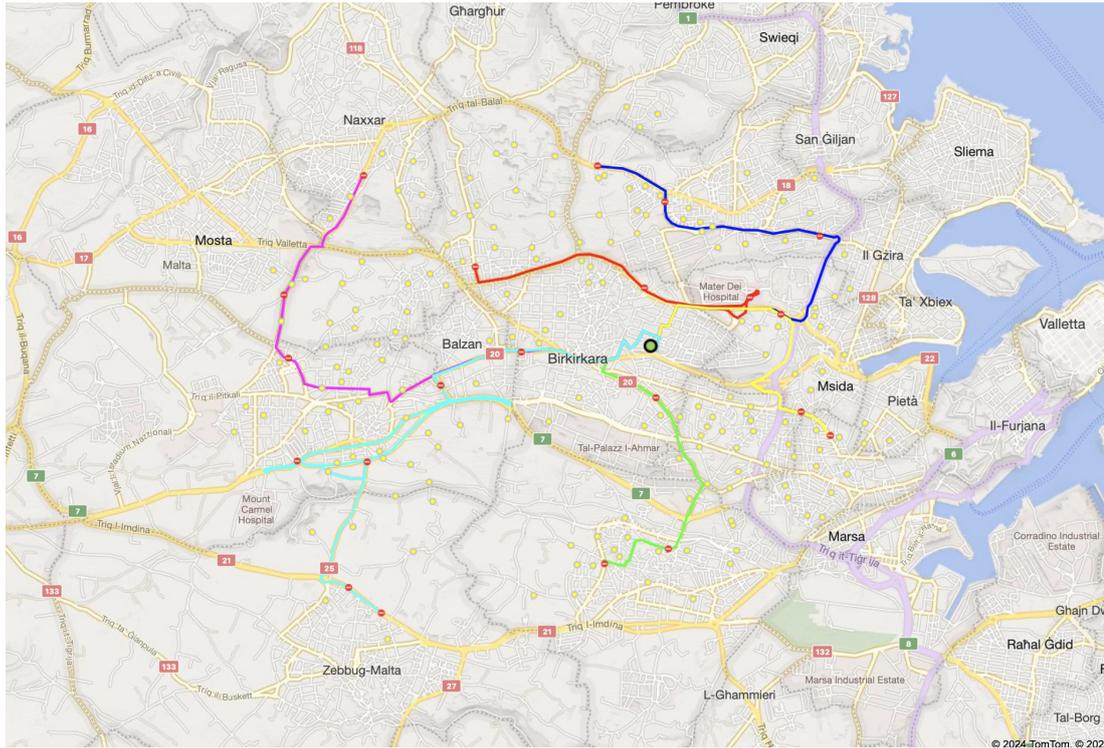


Figure B.32: Birkirkara best SRH-4 solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 306 students, average walk 7.17 minutes, and average journey time 15.65 minutes.

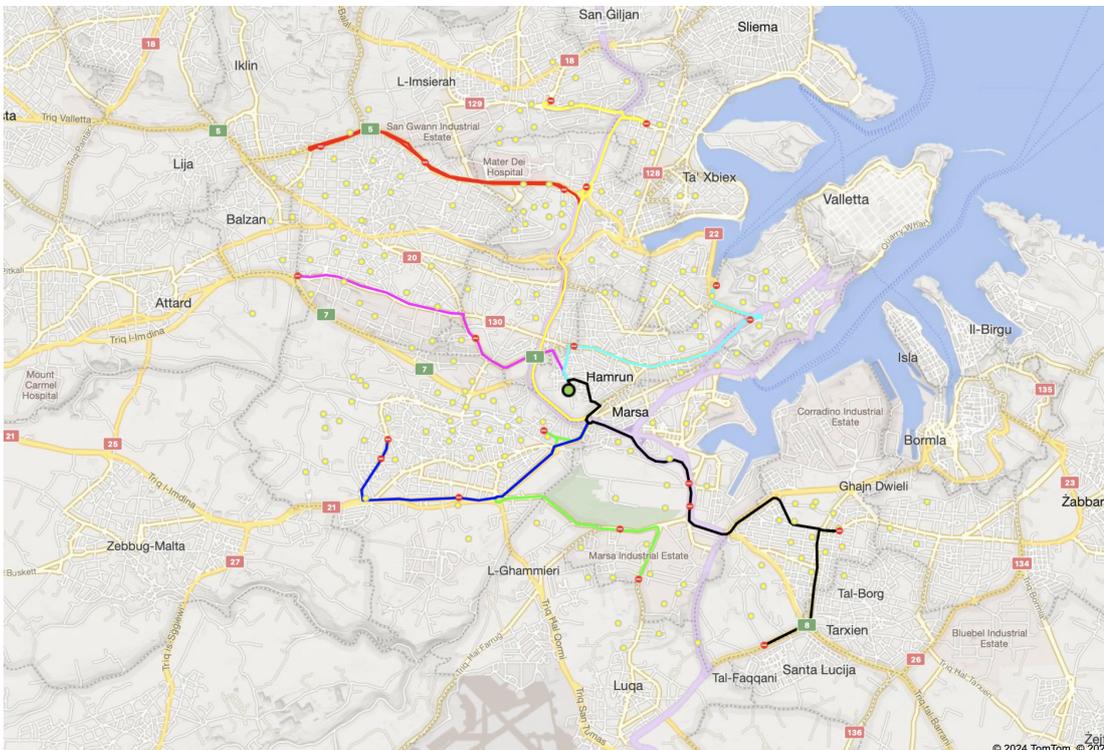


Figure B.33: Hamrun best SRH-4 solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 321 students, average walk 6.94 minutes, and average journey time 13.91 minutes.

## APPENDIX B. HEURISTIC ALGORITHM DETERMINISTIC SCENARIO RESULTS

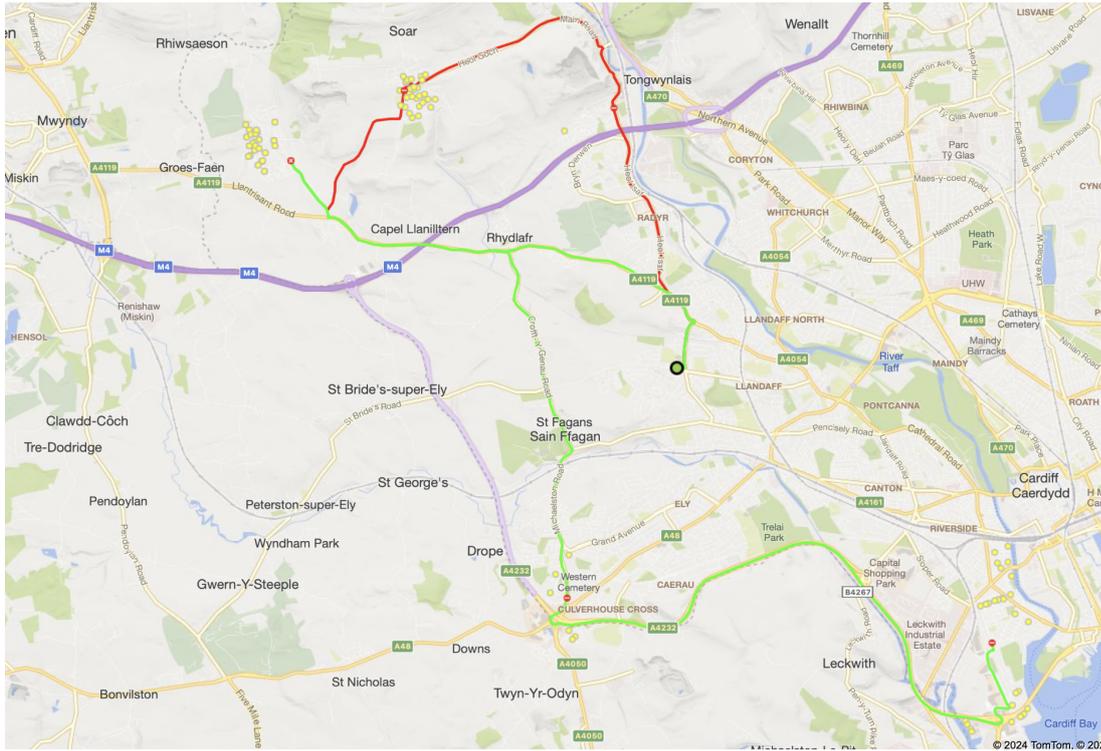


Figure B.34: Cardiff best SRH-4 solution with 2 routes (red, lime), 156 students, average walk 10.12 minutes, and average journey time 33.56 minutes.

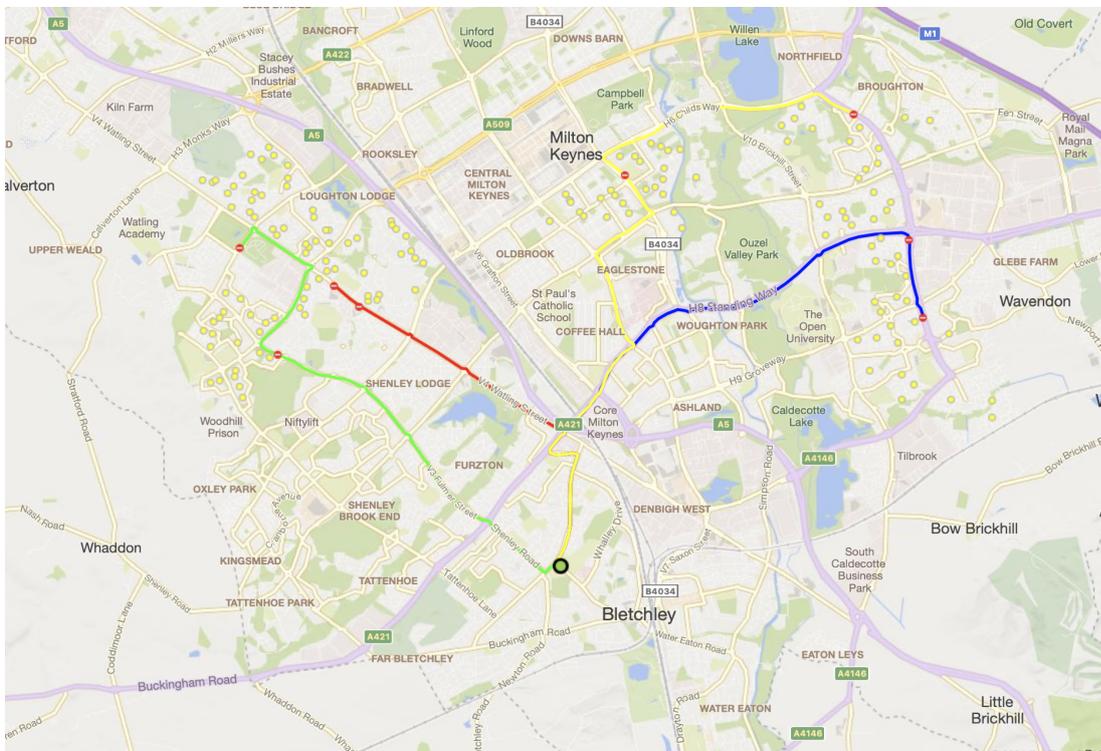


Figure B.35: Milton Keynes best SRH-4 solution with 4 routes (red, lime, blue, yellow), 274 students, average walk 11.41 minutes, and average journey time 14.76 minutes.

APPENDIX B. HEURISTIC ALGORITHM DETERMINISTIC SCENARIO RESULTS

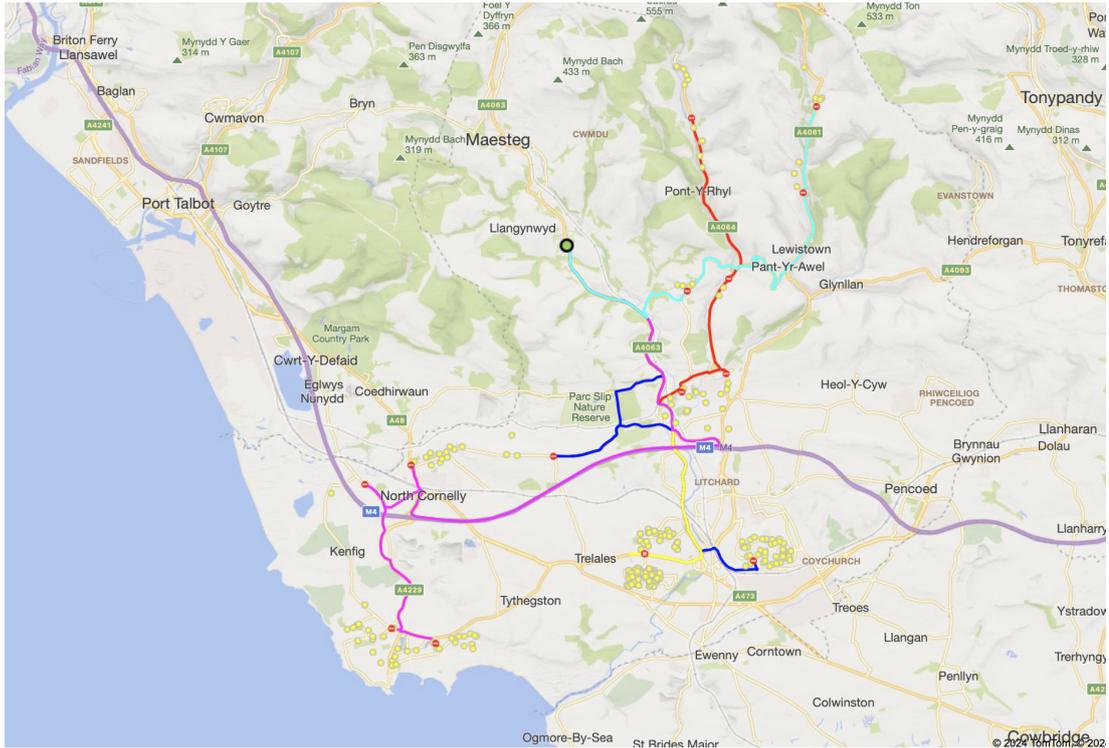


Figure B.36: Bridgend best SRH-4 solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 381 students, average walk 11.75 minutes, and average journey time 28.32 minutes.

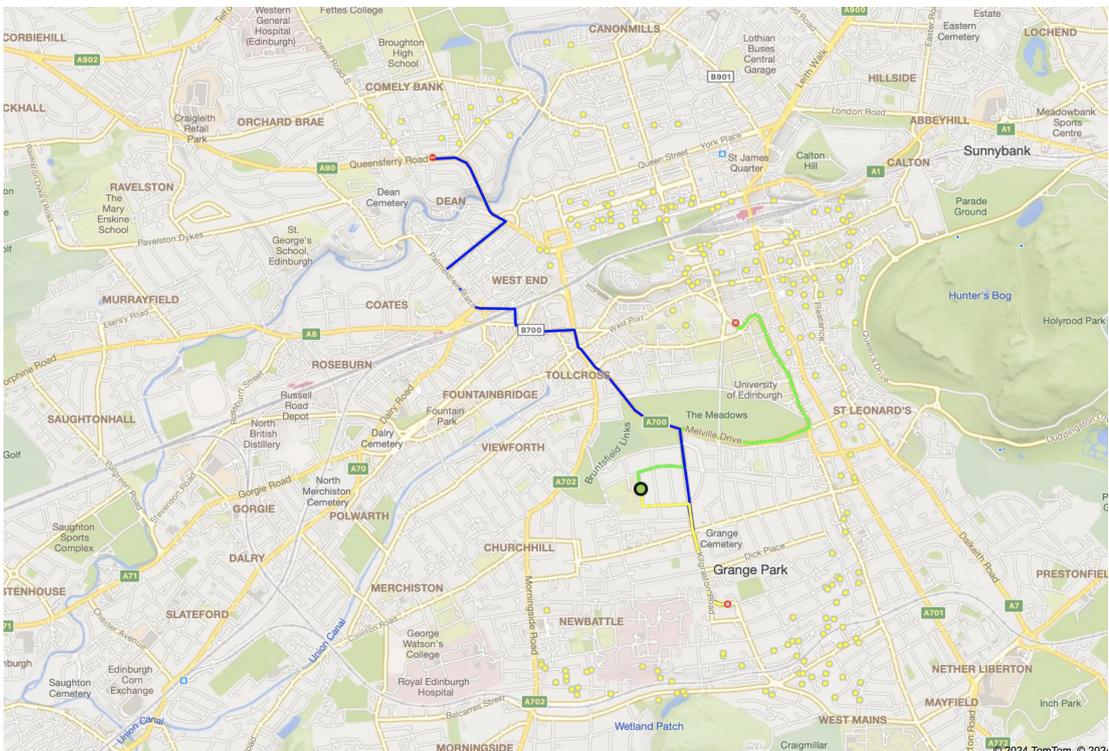


Figure B.37: Edinburgh-2 best SRH-4 solution with 4 routes (red, lime, blue, yellow), 320 students, average walk 11.10 minutes, and average journey time 14.78 minutes.

APPENDIX B. HEURISTIC ALGORITHM DETERMINISTIC SCENARIO RESULTS

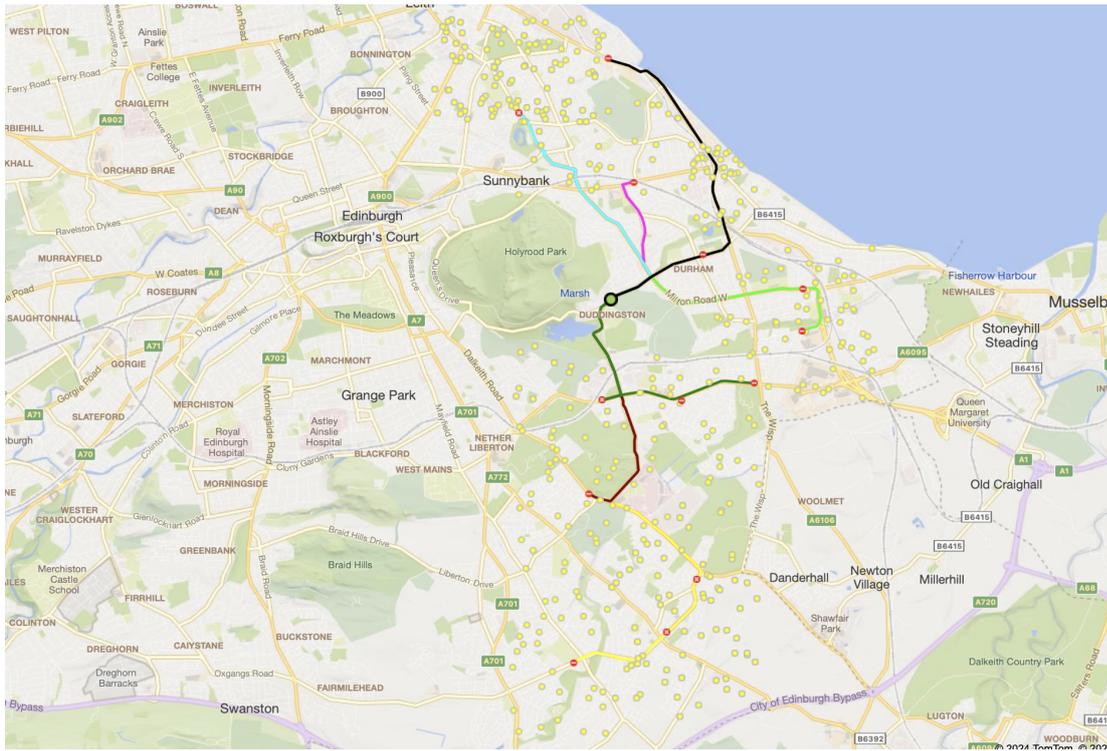


Figure B.38: Edinburgh-1 best SRH-4 solution with 9 routes (red, lime, blue, yellow, magenta, cyan, black, brown, green), 680 students, average walk 10.28 minutes, and average journey time 15.64 minutes.

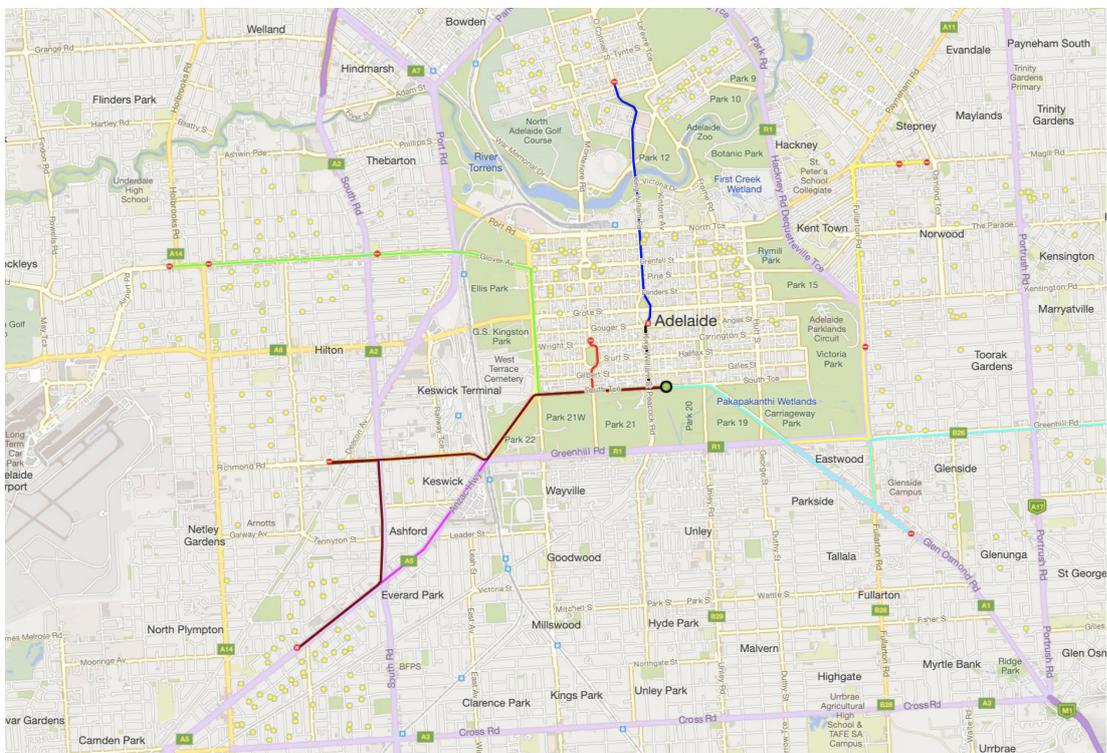


Figure B.39: Adelaide best SRH-4 solution with 8 routes (red, lime, blue, yellow, magenta, cyan, black, brown), 565 students, average walk 12.01 minutes, and average journey time 15.80 minutes.

APPENDIX B. HEURISTIC ALGORITHM DETERMINISTIC SCENARIO RESULTS

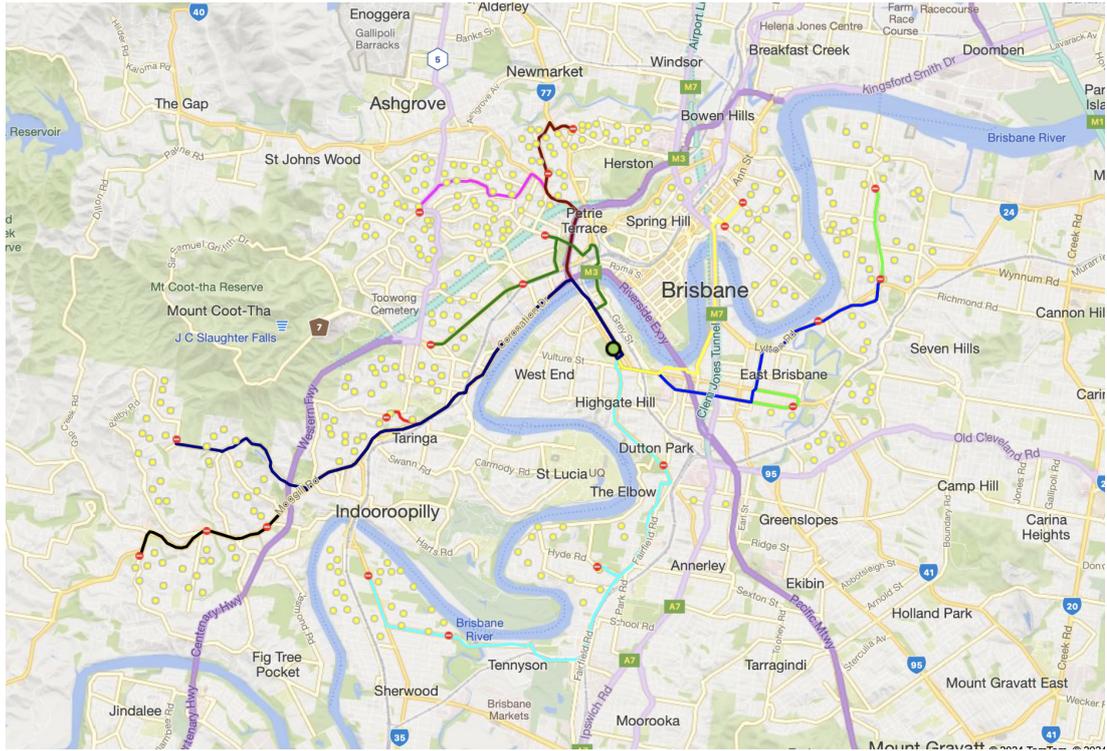


Figure B.40: Brisbane best SRH-4 solution with 10 routes (red, lime, blue, yellow, magenta, cyan, black, brown, green, navy), 757 students, average walk 11.85 minutes, and average journey time 20.70 minutes.



# Appendix C

## Heuristic Algorithm Independent Scenario Results

### C.1 Results for $\kappa = 0.95$

All results presented in this section are averaged across 25 runs. TMT and TPT stand for Total Mean Time and Total 95<sup>th</sup> Percentile Time (minutes), presented as mean  $\pm$  standard deviation. Comparative results for the deterministic scenario are also presented, with the column IR indicating the number of runs for which the deterministic solution is infeasible with respect to the on-time arrival chance constraints. A hyphen indicates infeasibility.

Table C.1: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.2, \beta = 0.5, \kappa = 0.95$ .

Location	Deterministic Scenario				Independent Scenario		
	$k$	TMT	IR	TPT	$k$	TMT	TPT
Suffolk	3	116.61 $\pm$ 0.90	25	-	4	116.01 $\pm$ 0.87	154.03 $\pm$ 1.30
Canberra	7	179.11 $\pm$ 0.53	25	-	8	190.27 $\pm$ 0.96	260.17 $\pm$ 1.28
Bridgend	6	170.55 $\pm$ 0.53	25	-	7	179.09 $\pm$ 0.76	248.76 $\pm$ 0.69
Brisbane	10	213.41 $\pm$ 4.31	0	287.16 $\pm$ 6.21	10	213.42 $\pm$ 3.66	285.74 $\pm$ 4.82

Table C.2: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.2, \beta = 0.75, \kappa = 0.95$ .

Location	Deterministic Scenario				Independent Scenario		
	$k$	TMT	IR	TPT	$k$	TMT	TPT
Suffolk	3	116.61 $\pm$ 0.90	25	-	4.96	115.60 $\pm$ 0.89	163.71 $\pm$ 1.22
Canberra	7	179.11 $\pm$ 0.53	25	-	8.00	190.05 $\pm$ 1.21	268.67 $\pm$ 1.63
Cardiff	2	67.12 $\pm$ 0.00	25	-	3.00	57.52 $\pm$ 0.00	81.86 $\pm$ 0.00
Bridgend	6	170.55 $\pm$ 0.53	25	-	7.00	180.43 $\pm$ 0.83	260.09 $\pm$ 1.16
Brisbane	10	213.41 $\pm$ 4.31	0	294.49 $\pm$ 6.59	10.00	212.20 $\pm$ 4.52	292.21 $\pm$ 6.72

APPENDIX C. HEURISTIC ALGORITHM INDEPENDENT SCENARIO  
RESULTS

Table C.3: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.1, \beta = 0.5, \kappa = 0.95$ .

Location	Deterministic Scenario				Independent Scenario		
	$k$	TMT	IR	TPT	$k$	TMT	TPT
Bridgend	6	170.55 $\pm$ 0.53	20	-	6.68	173.90 $\pm$ 4.22	213.00 $\pm$ 5.47

Table C.4: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.1, \beta = 0.75, \kappa = 0.95$ .

Location	Deterministic Scenario				Independent Scenario		
	$k$	TMT	IR	TPT	$k$	TMT	TPT
Canberra	7	179.11 $\pm$ 0.53	25	-	7.00	190.43 $\pm$ 0.95	230.54 $\pm$ 1.24
Bridgend	6	170.55 $\pm$ 0.53	20	-	6.76	176.47 $\pm$ 7.96	215.26 $\pm$ 10.38

Table C.5: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.3, \beta = 0.5, \kappa = 0.95$ .

Location	Deterministic Scenario				Independent Scenario		
	$k$	TMT	IR	TPT	$k$	TMT	TPT
Suffolk	3	116.61 $\pm$ 0.90	25	-	4.04	116.89 $\pm$ 0.75	159.74 $\pm$ 0.84
Canberra	7	179.11 $\pm$ 0.53	25	-	8.00	193.06 $\pm$ 0.63	277.41 $\pm$ 0.87
Cardiff	2	67.12 $\pm$ 0.00	25	-	3.00	57.53 $\pm$ 0.00	84.96 $\pm$ 0.00
Bridgend	6	170.55 $\pm$ 0.53	25	-	7.96	188.85 $\pm$ 3.25	281.63 $\pm$ 6.03
Brisbane	10	213.41 $\pm$ 4.31	0	304.84 $\pm$ 6.56	10.00	214.27 $\pm$ 3.49	301.99 $\pm$ 5.20

Table C.6: Results averaged across 25 runs for the independent scenario with parameters  $\alpha = 0.3, \beta = 0.75, \kappa = 0.95$ .

Location	Deterministic Scenario				Independent Scenario		
	$k$	TMT	IR	TPT	$k$	TMT	TPT
Suffolk	3	116.61 $\pm$ 0.90	25	-	5	115.25 $\pm$ 0.00	175.41 $\pm$ 0.00
Pembroke	7	103.49 $\pm$ 0.68	0	156.89 $\pm$ 1.31	7	104.22 $\pm$ 0.64	155.10 $\pm$ 1.05
Canberra	7	179.11 $\pm$ 0.53	25	-	9	196.09 $\pm$ 0.52	307.61 $\pm$ 0.48
Valletta	6	104.50 $\pm$ 0.78	0	163.34 $\pm$ 1.92	6	105.70 $\pm$ 1.07	162.50 $\pm$ 1.99
Birkirkara	6	97.49 $\pm$ 2.64	0	144.52 $\pm$ 5.57	6	97.34 $\pm$ 1.91	142.42 $\pm$ 3.35
Cardiff	2	67.12 $\pm$ 0.00	25	-	3	57.53 $\pm$ 0.00	90.76 $\pm$ 0.00
Bridgend	6	170.55 $\pm$ 0.53	25	-	8	194.97 $\pm$ 0.68	314.10 $\pm$ 0.89
Edinburgh-1	9	143.71 $\pm$ 1.56	0	208.43 $\pm$ 2.75	9	144.51 $\pm$ 1.45	208.76 $\pm$ 2.41
Brisbane	10	213.41 $\pm$ 4.31	5	-	10	212.85 $\pm$ 5.03	320.82 $\pm$ 7.60

## C.2 Computational Times and Feasibility Rates

All results presented in this section are averaged across 25 runs and presented as mean  $\pm$  standard deviation. Time gives the computational time (seconds), and FR stands for Feasibility Rate (percentage of iterations yielding a solution with all route percentile journey times being at most 45 minutes).

Table C.7: Computational times and feasibility rates of the independent scenario’s runs with  $\alpha = 0.2, \beta = 0.5$ .

Location	$\kappa$	Time (s)		FR (%)	
Mgarr	0.99	285.62 $\pm$	2.88	100.00 $\pm$	0.00
Mellieħa	0.99	321.07 $\pm$	22.97	100.00 $\pm$	0.00
Porthcawl	0.99	1,037.32 $\pm$	12.29	92.11 $\pm$	0.27
Qrendi	0.99	304.85 $\pm$	12.18	98.44 $\pm$	0.25
Suffolk	0.95	494.27 $\pm$	7.90	27.92 $\pm$	5.87
Senglea	0.99	314.79 $\pm$	20.93	100.00 $\pm$	0.00
Victoria	0.99	437.51 $\pm$	197.46	40.26 $\pm$	40.57
Pembroke	0.99	326.91 $\pm$	27.19	99.94 $\pm$	0.03
Canberra	0.95	477.52 $\pm$	28.70	71.18 $\pm$	4.86
Handaq	0.99	434.76 $\pm$	58.52	99.58 $\pm$	0.19
Valetta	0.99	381.76 $\pm$	39.57	66.61 $\pm$	12.81
Birkirkara	0.99	577.85 $\pm$	23.27	87.21 $\pm$	4.03
Hamrun	0.99	569.21 $\pm$	49.25	99.99 $\pm$	0.01
Cardiff	0.99	1,062.73 $\pm$	66.30	23.17 $\pm$	0.09
Milton Keynes	0.99	523.18 $\pm$	85.10	95.52 $\pm$	0.73
Bridgend	0.95	936.63 $\pm$	18.65	55.16 $\pm$	4.18
Edinburgh-2	0.99	568.90 $\pm$	3.26	65.04 $\pm$	0.25
Edinburgh-1	0.99	619.53 $\pm$	12.78	72.57 $\pm$	5.61
Adelaide	0.99	602.12 $\pm$	13.10	53.40 $\pm$	4.81
Brisbane	0.95	623.17 $\pm$	22.06	83.25 $\pm$	4.11

## C.3 Visualizations of Best Solutions

This section visualizes the best solutions listed in Table 5.4. In Figures C.1 to C.20, the lime dot represents the school, the red dots represent the visited bus stops, and the yellow dots represent the student addresses. Moreover, each route is shown in a different colour. Note that certain subroutes are not visible since they overlap with other subroutes. The following visualizations and further information on arrival times, dwell times, distances covered, and walking/riding times and distances for each student can be viewed from the interactive visualization files at (Sciortino, 2025c).

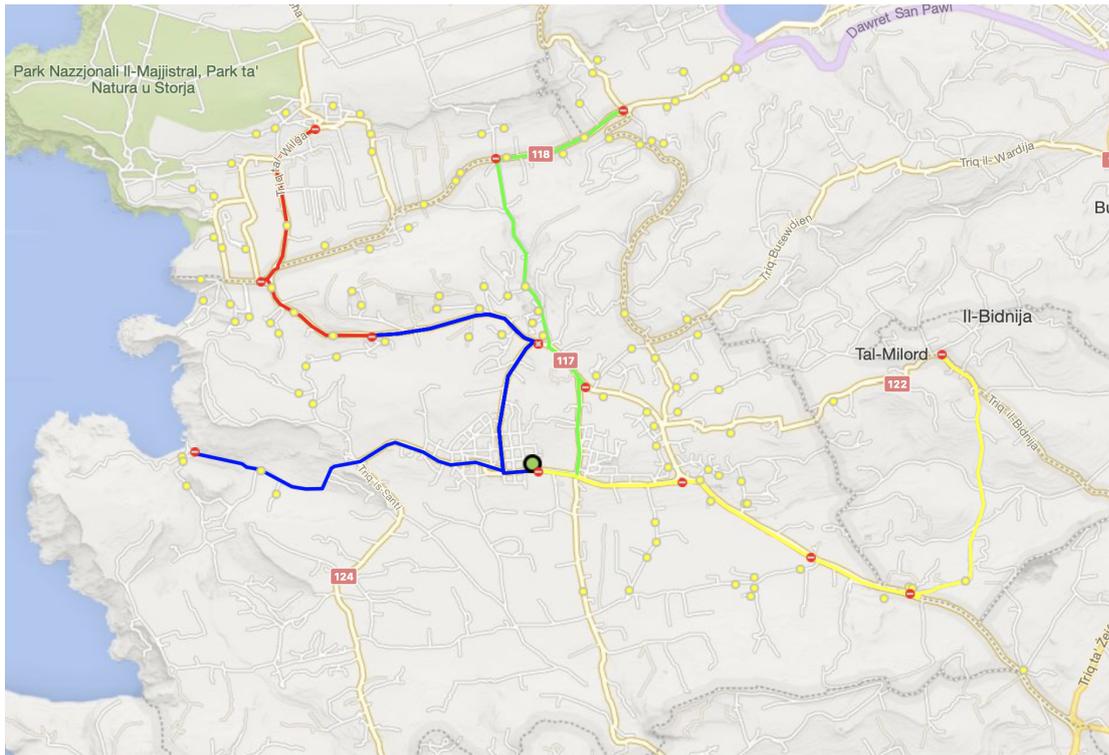


Figure C.1: Mgarr best independent solution with 4 routes (red, lime, blue, yellow), 190 students, average walk 6.02 minutes, average journey time 13.63 minutes, and average percentile time 24.49 minutes.

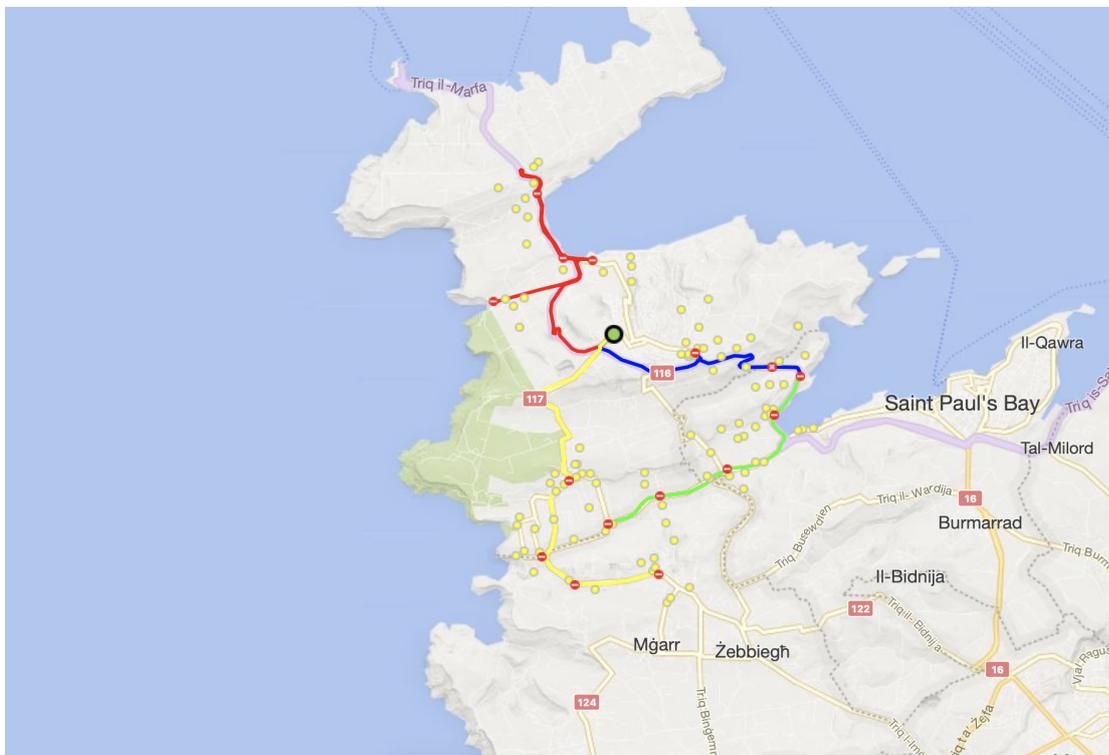


Figure C.2: Mellieħa best independent solution with 4 routes (red, lime, blue, yellow), 171 students, average walk 4.98 minutes, average journey time 14.19 minutes, and average percentile time 24.94 minutes.

APPENDIX C. HEURISTIC ALGORITHM INDEPENDENT SCENARIO RESULTS

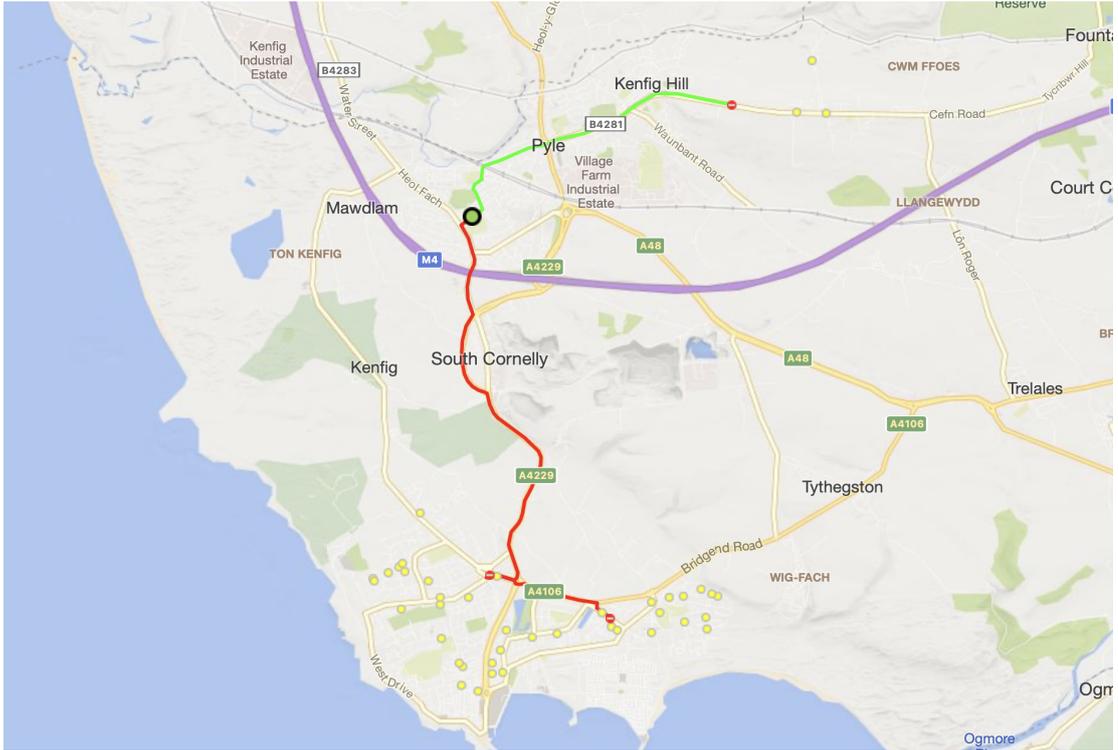


Figure C.3: Porthcawl best independent solution with 2 routes (red, lime), 66 students, average walk 12.38 minutes, average journey time 12.14 minutes, and average percentile time 26.39 minutes.

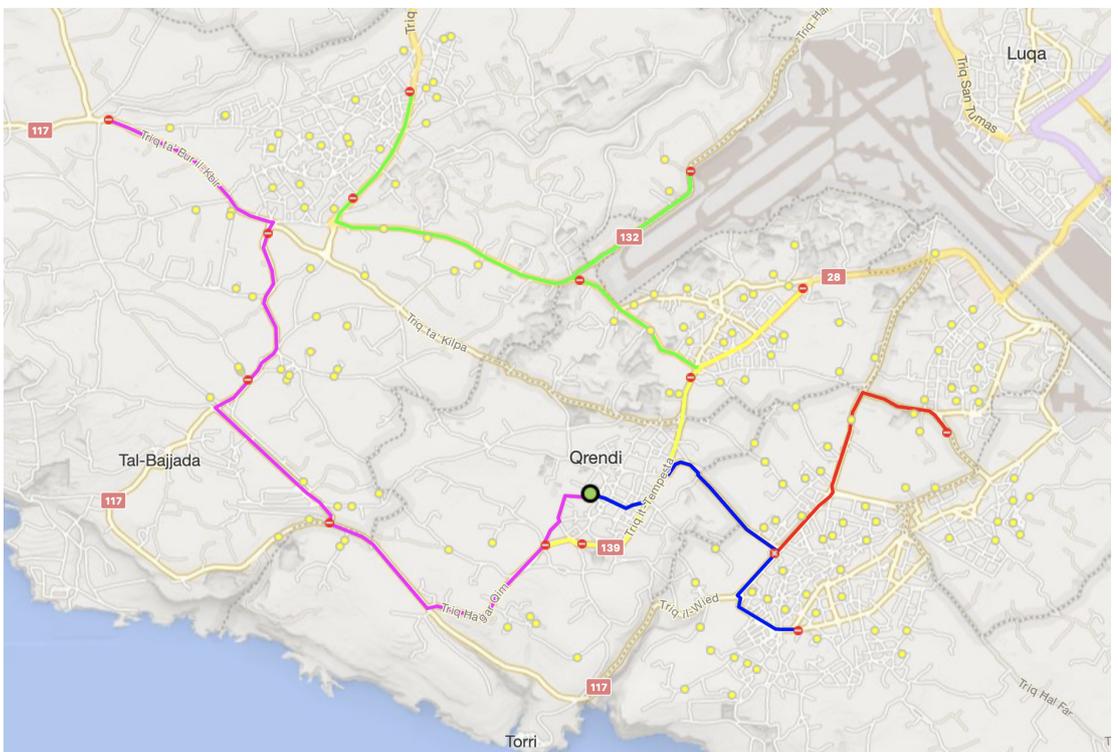


Figure C.4: Qrendi best independent solution with 5 routes (red, lime, blue, yellow, magenta), 255 students, average walk 6.67 minutes, average journey time 15.15 minutes, and average percentile time 27.21 minutes.



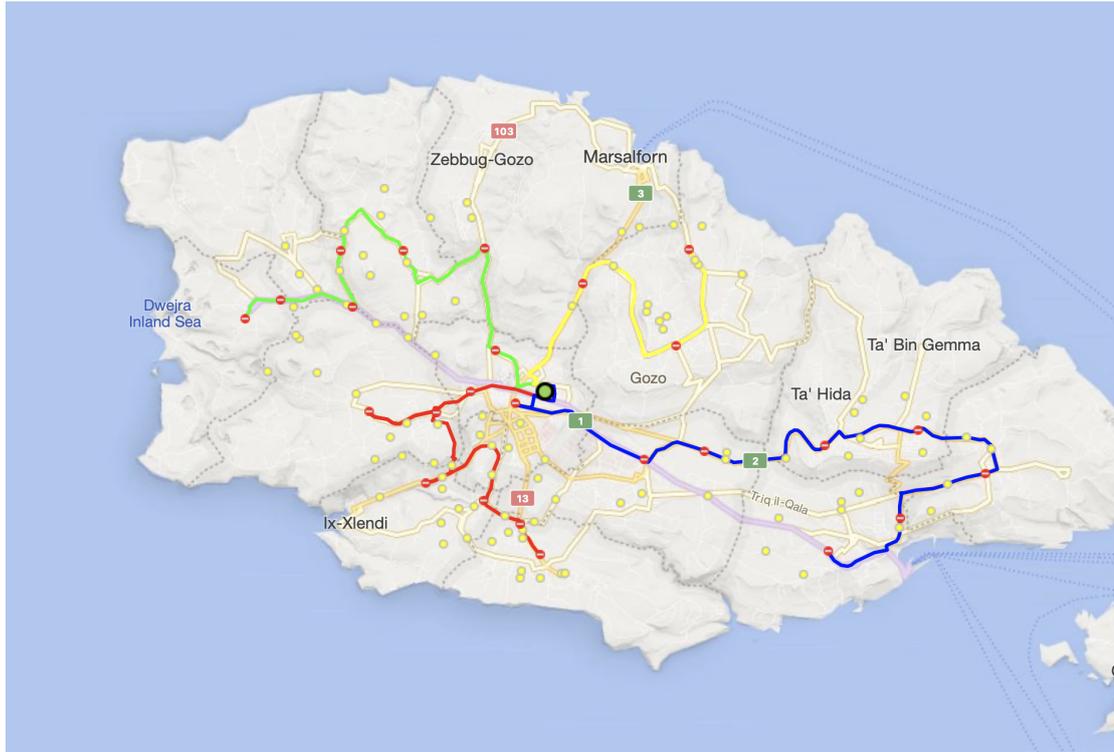


Figure C.7: Victoria best independent solution with 4 routes (red, lime, blue, yellow), 171 students, average walk 7.32 minutes, average journey time 23.26 minutes, and average percentile time 38.21 minutes.

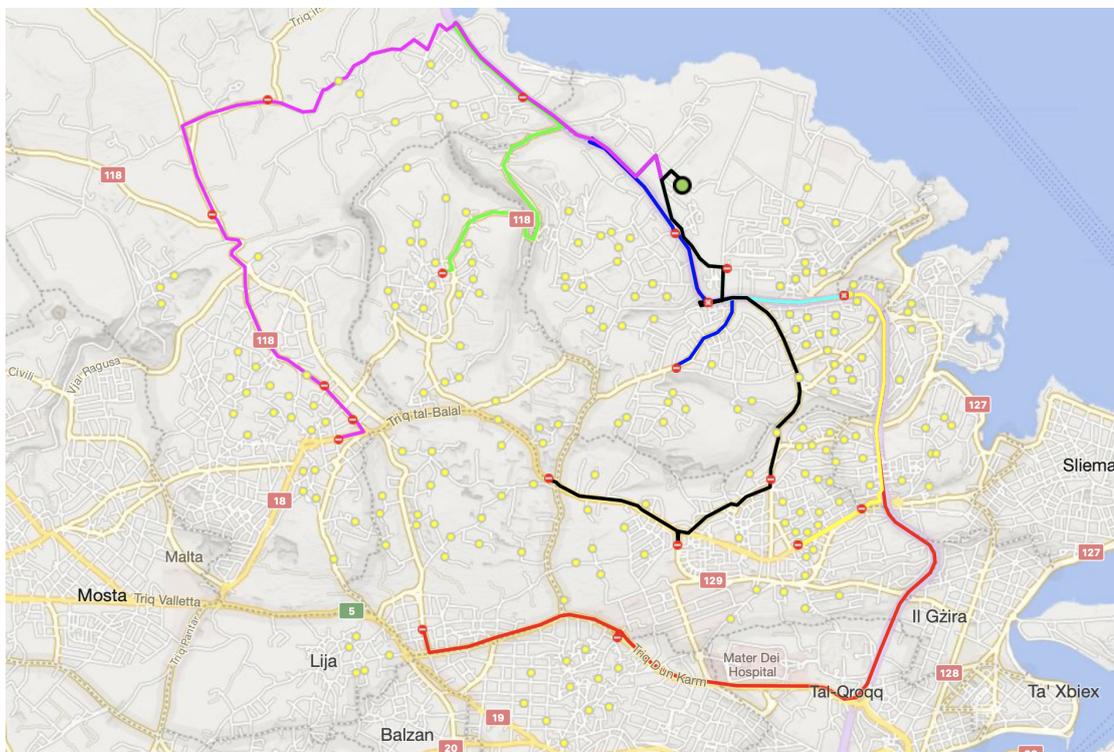


Figure C.8: Pembroke best independent solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 335 students, average walk 6.78 minutes, average journey time 14.94 minutes, and average percentile time 27.30 minutes.

APPENDIX C. HEURISTIC ALGORITHM INDEPENDENT SCENARIO RESULTS

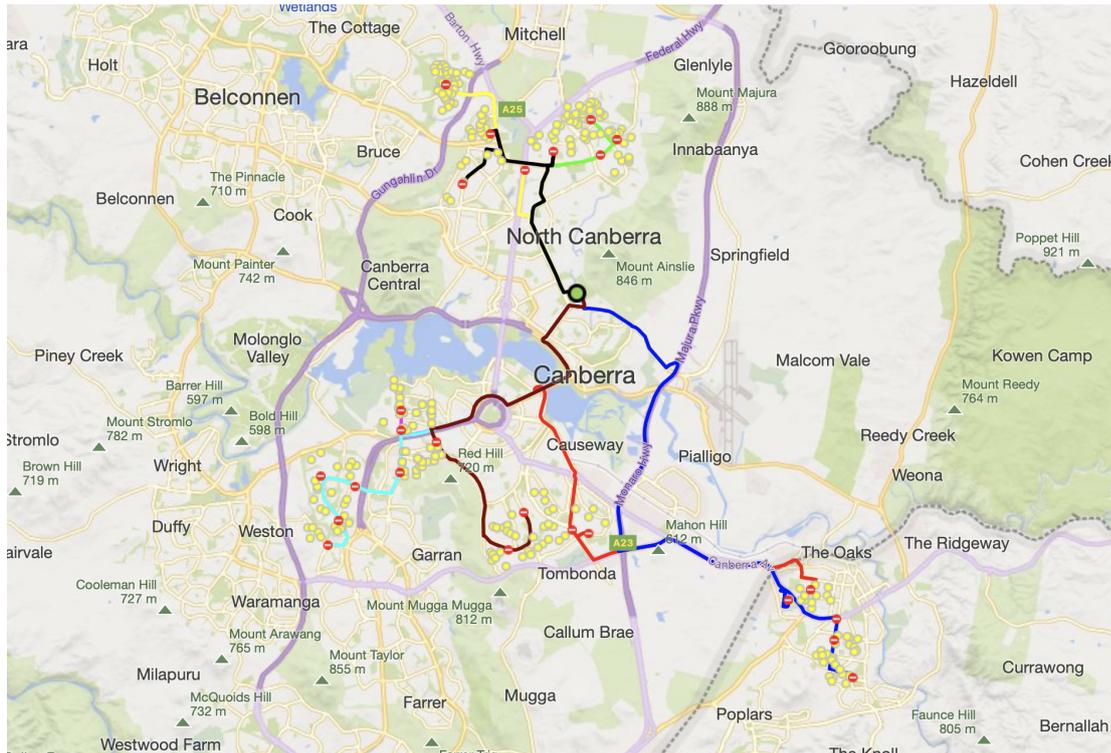


Figure C.9: Canberra best independent solution with 8 routes (red, lime, blue, yellow, magenta, cyan, black, brown), 499 students, average walk 6.97 minutes, average journey time 23.60 minutes, and average percentile time 32.26 minutes.

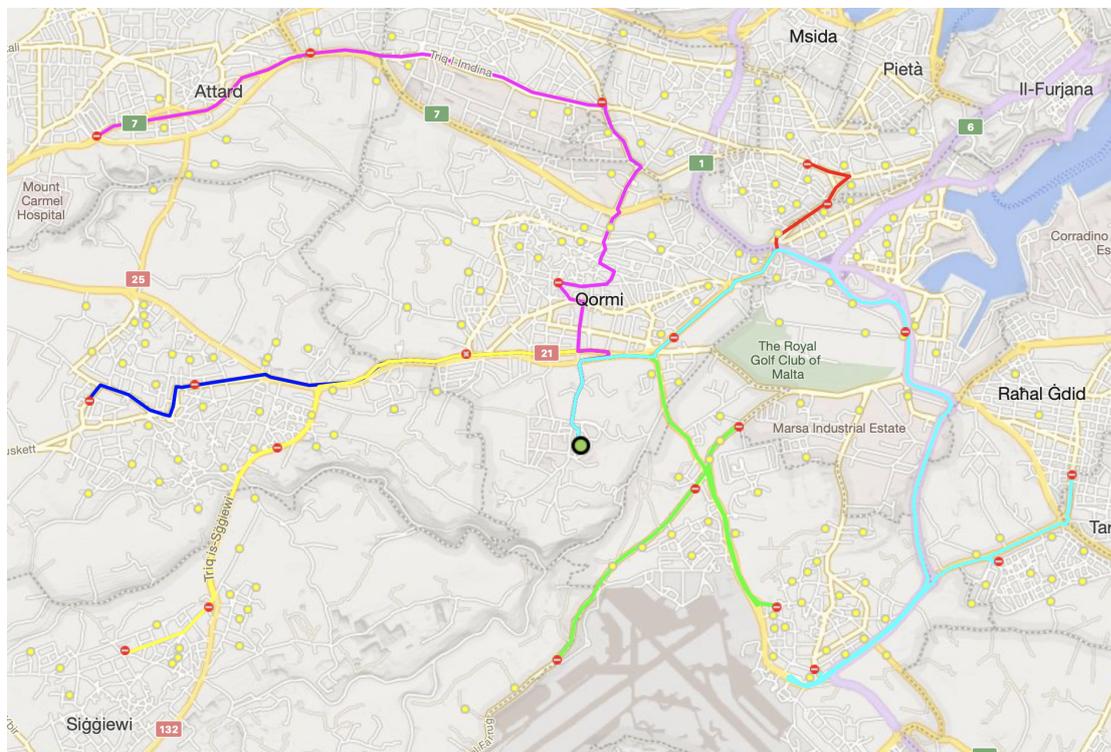


Figure C.10: Handaq best independent solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 285 students, average walk 6.80 minutes, average journey time 16.19 minutes, and average percentile time 28.45 minutes.

APPENDIX C. HEURISTIC ALGORITHM INDEPENDENT SCENARIO RESULTS

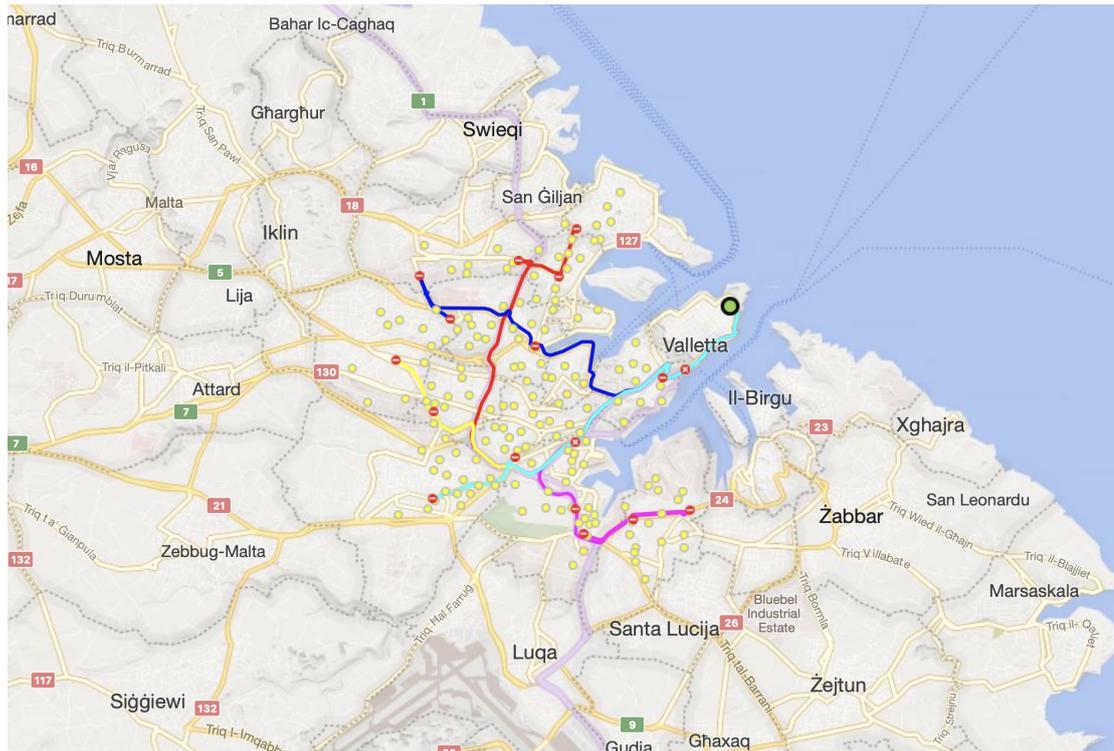


Figure C.11: Valletta best independent solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 268 students, average walk 6.67 minutes, average journey time 17.67 minutes, and average percentile time 34.06 minutes.

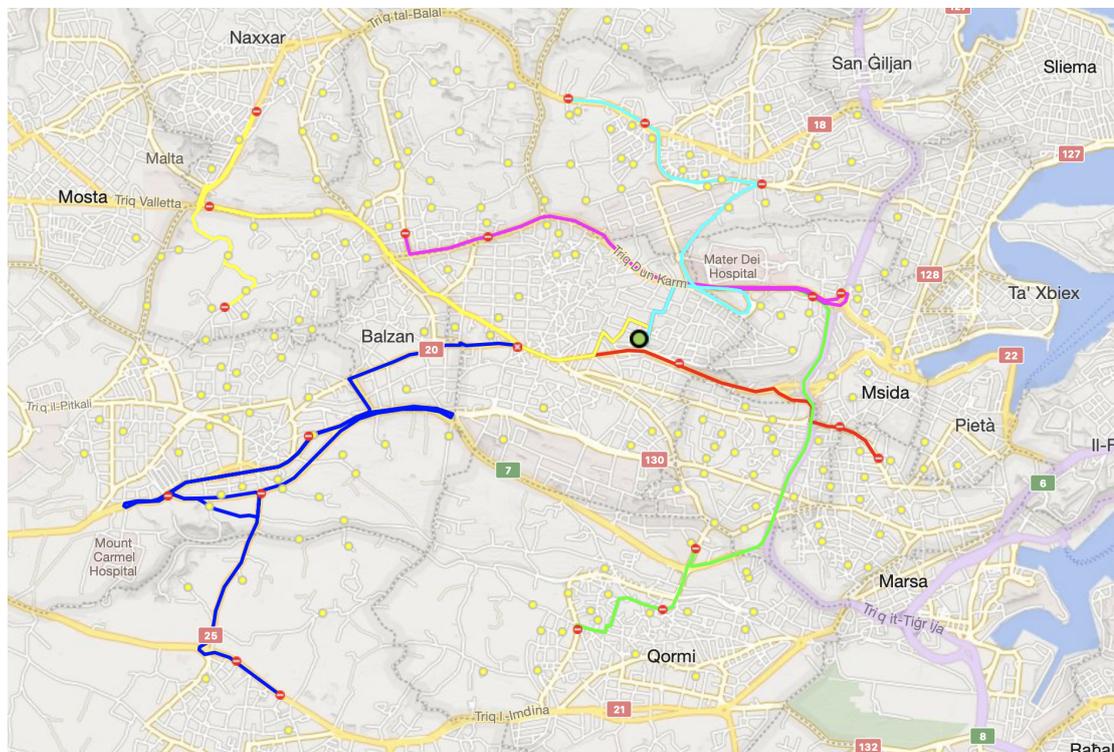


Figure C.12: Birkirkara best independent solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 306 students, average walk 6.54 minutes, average journey time 15.44 minutes, and average percentile time 26.84 minutes.

APPENDIX C. HEURISTIC ALGORITHM INDEPENDENT SCENARIO RESULTS

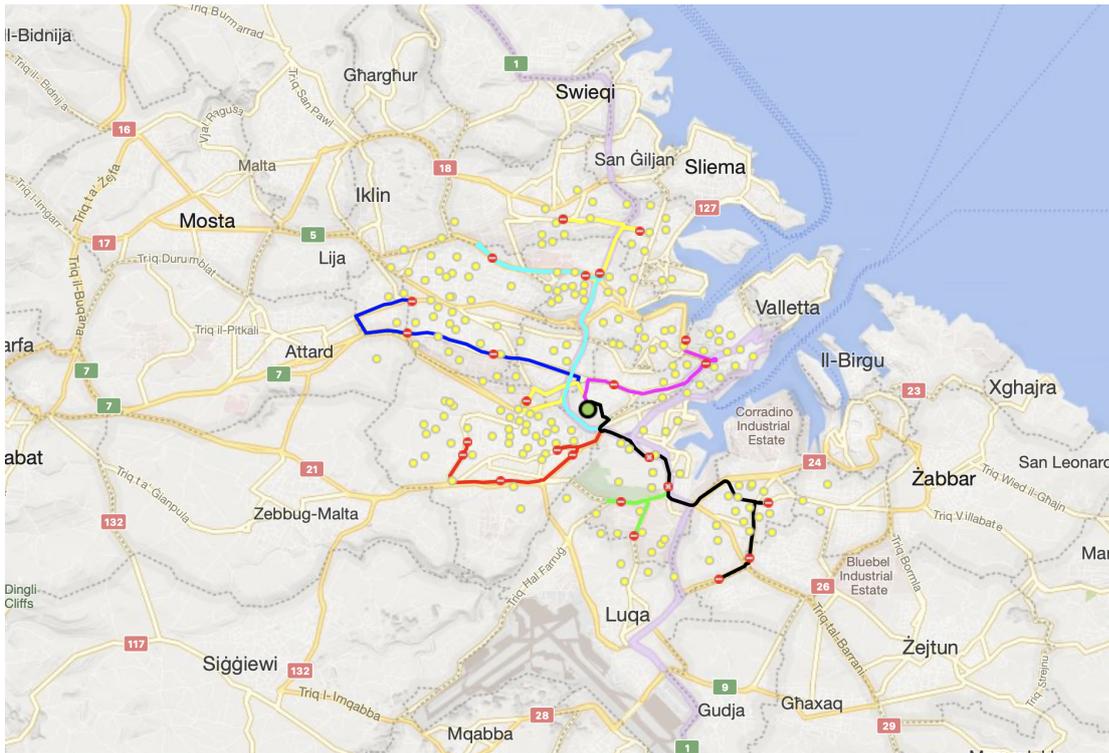


Figure C.13: Hamrun best independent solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 321 students, average walk 6.66 minutes, average journey time 14.14 minutes, and average percentile time 25.04 minutes.

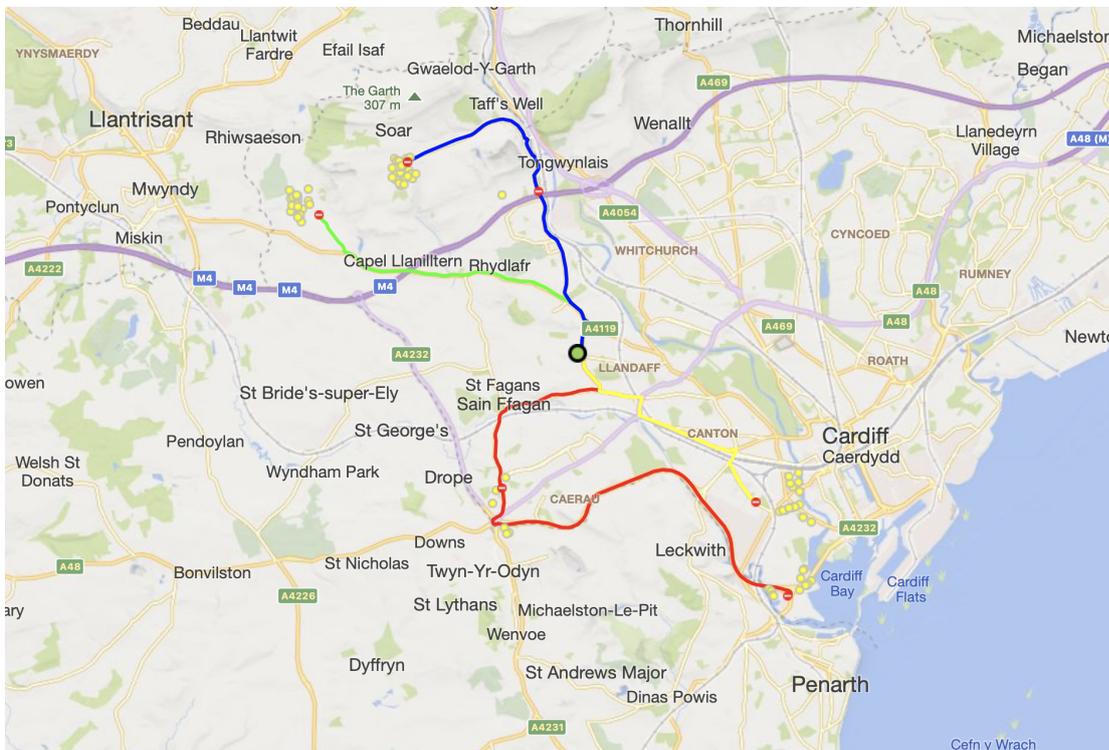


Figure C.14: Cardiff best independent solution with 4 routes (red, lime, blue, yellow), 156 students, average walk 10.06 minutes, average journey time 16.98 minutes, and average percentile time 37.64 minutes.

APPENDIX C. HEURISTIC ALGORITHM INDEPENDENT SCENARIO RESULTS

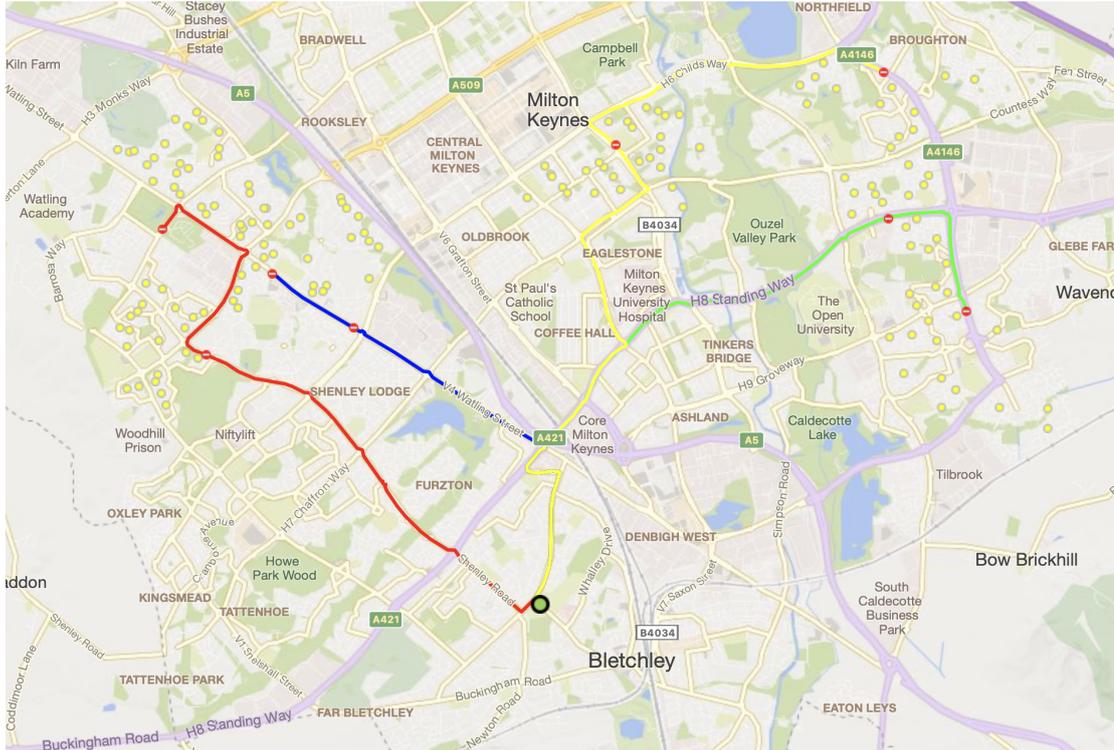


Figure C.15: Milton Keynes best independent solution with 4 routes (red, lime, blue, yellow), 274 students, average walk 11.44 minutes, average journey time 14.76 minutes, and average percentile time 27.97 minutes.

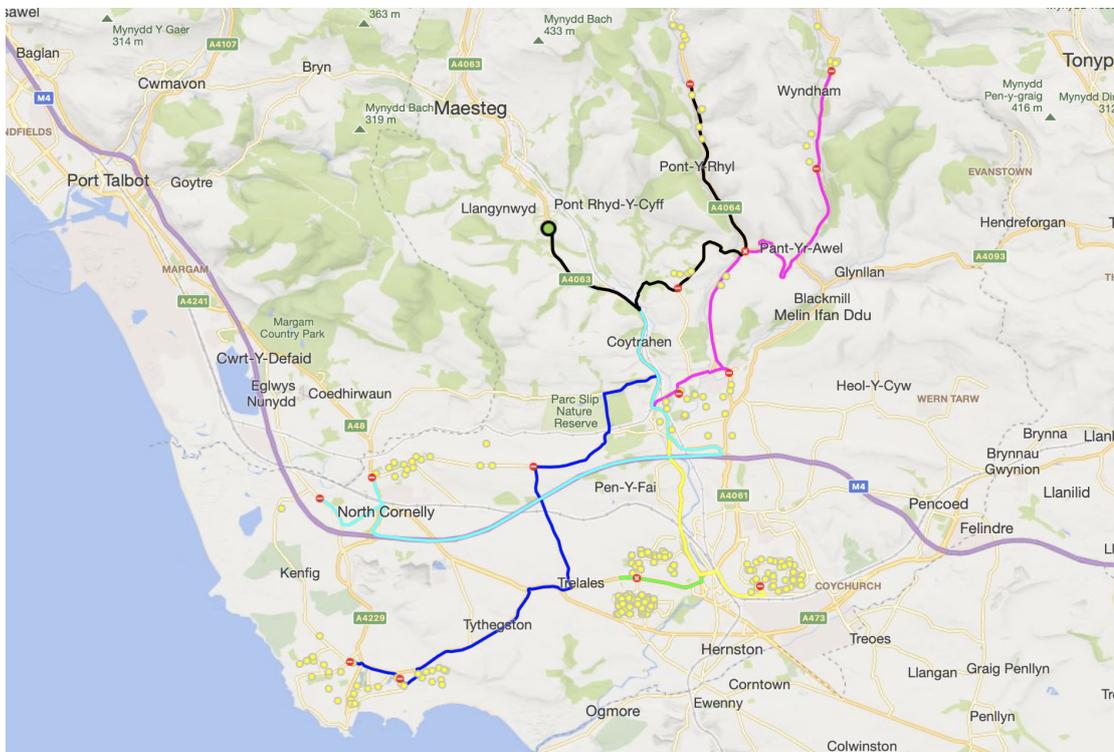


Figure C.16: Bridgend best independent solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 381 students, average walk 11.82 minutes, average journey time 25.39 minutes, and average percentile time 35.40 minutes.

APPENDIX C. HEURISTIC ALGORITHM INDEPENDENT SCENARIO RESULTS

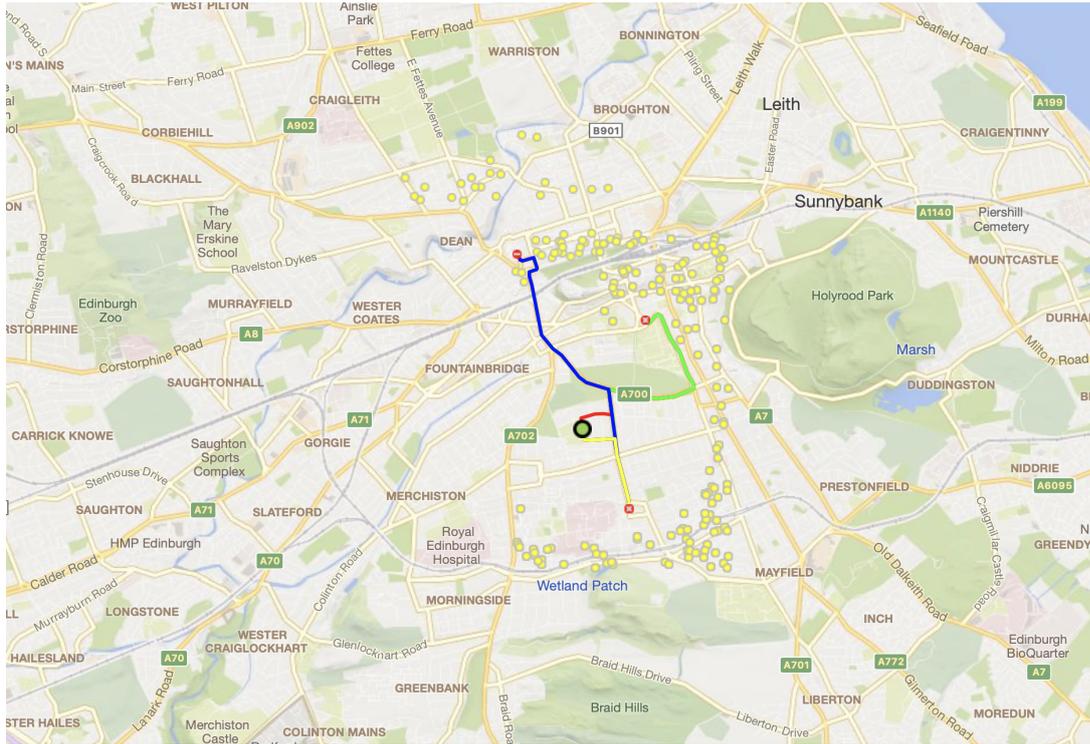


Figure C.17: Edinburgh-2 best independent solution with 4 routes (red, lime, blue, yellow), 320 students, average walk 10.81 minutes, average journey time 15.03 minutes, and average percentile time 27.52 minutes.

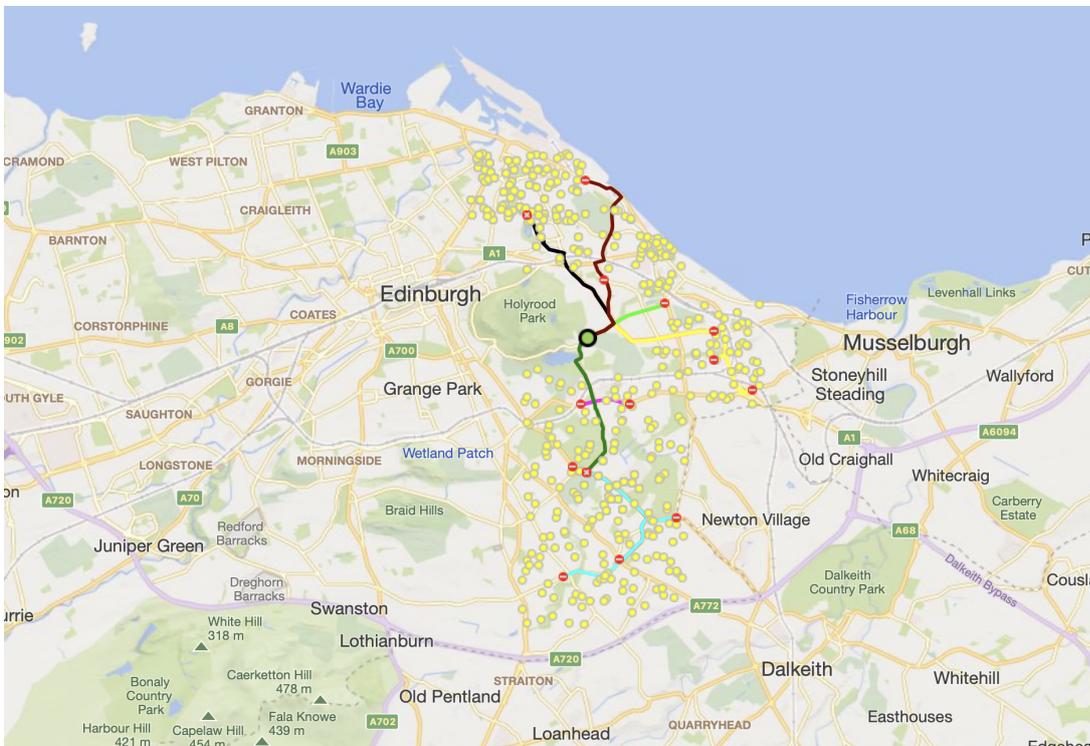


Figure C.18: Edinburgh-1 best independent solution with 9 routes (red, lime, blue, yellow, magenta, cyan, black, brown, green), 680 students, average walk 10.28 minutes, average journey time 15.62 minutes, and average percentile time 28.51 minutes.

APPENDIX C. HEURISTIC ALGORITHM INDEPENDENT SCENARIO RESULTS

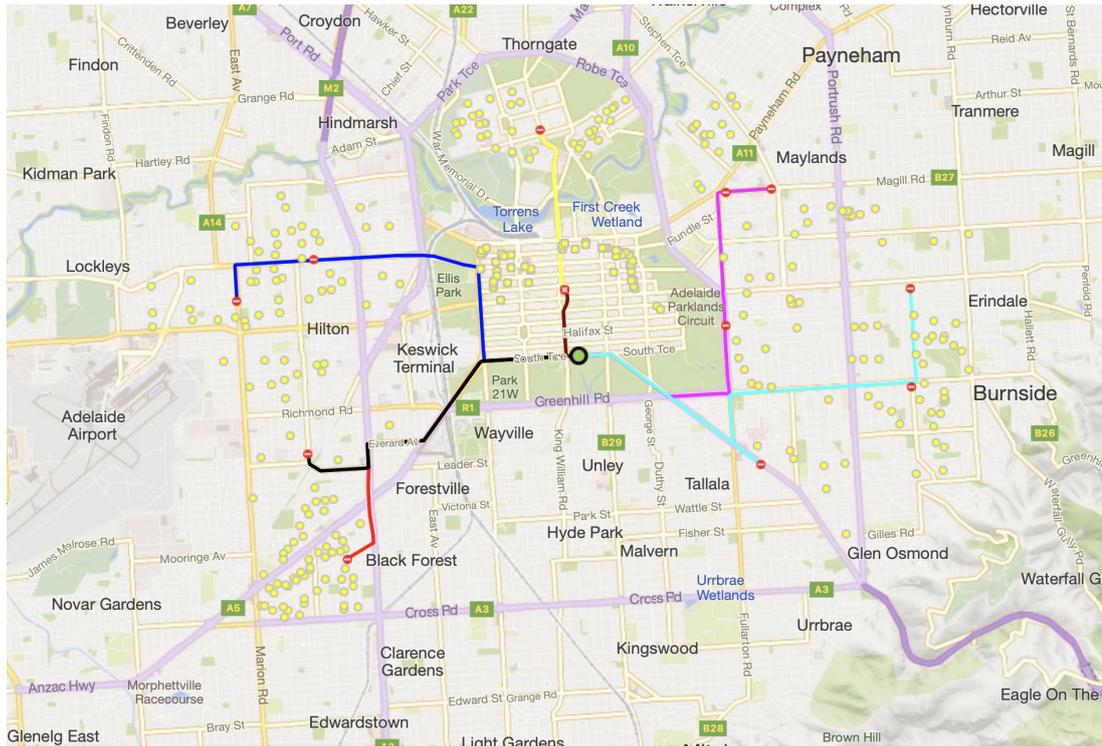


Figure C.19: Adelaide best independent solution with 8 routes (red, lime, blue, yellow, magenta, cyan, black, brown), 565 students, average walk 12.12 minutes, average journey time 15.90 minutes, and average percentile time 30.32 minutes.

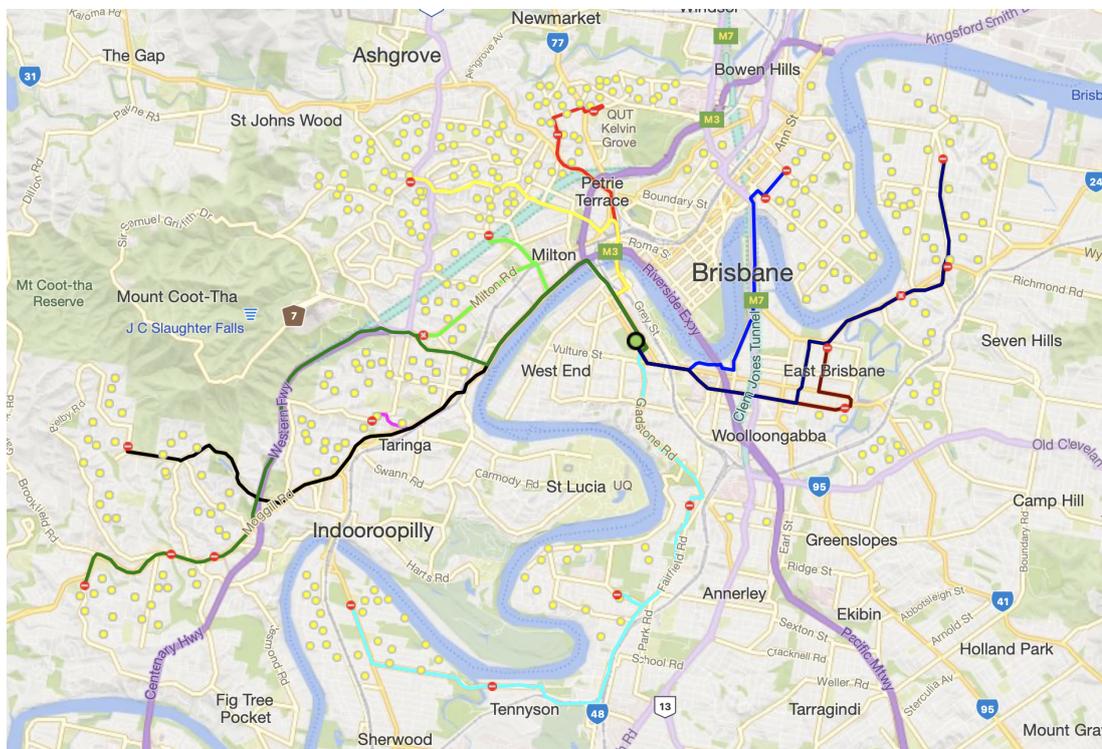


Figure C.20: Brisbane best independent solution with 10 routes (red, lime, blue, yellow, magenta, cyan, black, brown, green, navy), 757 students, average walk 11.85 minutes, average journey time 20.91 minutes, and average percentile time 27.97 minutes.



# Appendix D

## Heuristic Algorithm Correlated Scenario Results

### D.1 Computational Times and Feasibility Rates

All results presented in this section are averaged across 25 runs and presented as mean  $\pm$  standard deviation. Time gives the computational time (minutes), and FR stands for Feasibility Rate (percentage of iterations yielding a solution with all route percentile journey times being at most 45 minutes).

Table D.1: Computational times and feasibility rates of the correlated scenario's runs with  $\alpha = 0.2, \beta = 0.5$ .

Location	$\kappa$	Time (mins.)	FR (%)
Mġarr	0.99	97.22 $\pm$ 2.02	100.00 $\pm$ 0.00
Mellieħa	0.99	78.31 $\pm$ 5.58	100.00 $\pm$ 0.00
Porthcawl	0.99	28.35 $\pm$ 0.39	83.84 $\pm$ 0.08
Qrendi	0.99	73.47 $\pm$ 4.79	95.25 $\pm$ 0.49
Suffolk	0.95	86.60 $\pm$ 0.82	18.69 $\pm$ 4.83
Senglea	0.99	18.05 $\pm$ 3.04	99.99 $\pm$ 0.01
Victoria	0.99	960.98 $\pm$ 48.23	90.59 $\pm$ 1.64
Pembroke	0.99	75.08 $\pm$ 11.48	99.51 $\pm$ 0.20
Canberra	0.95	147.72 $\pm$ 3.99	72.32 $\pm$ 4.05
Handaq	0.99	127.68 $\pm$ 27.69	97.31 $\pm$ 0.87
Valletta	0.99	123.54 $\pm$ 25.03	50.31 $\pm$ 8.92
Birkirkara	0.99	499.06 $\pm$ 81.73	69.34 $\pm$ 7.71
Hamrun	0.99	404.78 $\pm$ 38.66	99.92 $\pm$ 0.04
Cardiff	0.99	14.60 $\pm$ 0.38	12.04 $\pm$ 0.25
Milton Keynes	0.99	8.63 $\pm$ 1.66	91.52 $\pm$ 1.62
Bridgend	0.95	540.72 $\pm$ 20.46	45.39 $\pm$ 7.10
Edinburgh-2	0.99	5.34 $\pm$ 0.46	43.00 $\pm$ 0.23
Edinburgh-1	0.99	243.73 $\pm$ 11.03	42.52 $\pm$ 6.88
Adelaide	0.99	190.35 $\pm$ 13.06	21.69 $\pm$ 7.37
Brisbane	0.95	143.77 $\pm$ 10.08	74.66 $\pm$ 6.82

## D.2 Visualizations of Best Solutions

This section visualizes the best solutions listed in Table 6.6. In Figures D.1 to D.20, the lime dot represents the school, the red dots represent the visited bus stops, and the yellow dots represent the student addresses. Moreover, each route is shown in a different colour. Note that certain subroutes are not visible since they overlap with other subroutes. The following visualizations and further information on arrival times, dwell times, distances covered, and walking/riding times and distances for each student can be viewed from the interactive visualization files at Sciortino (2025b).

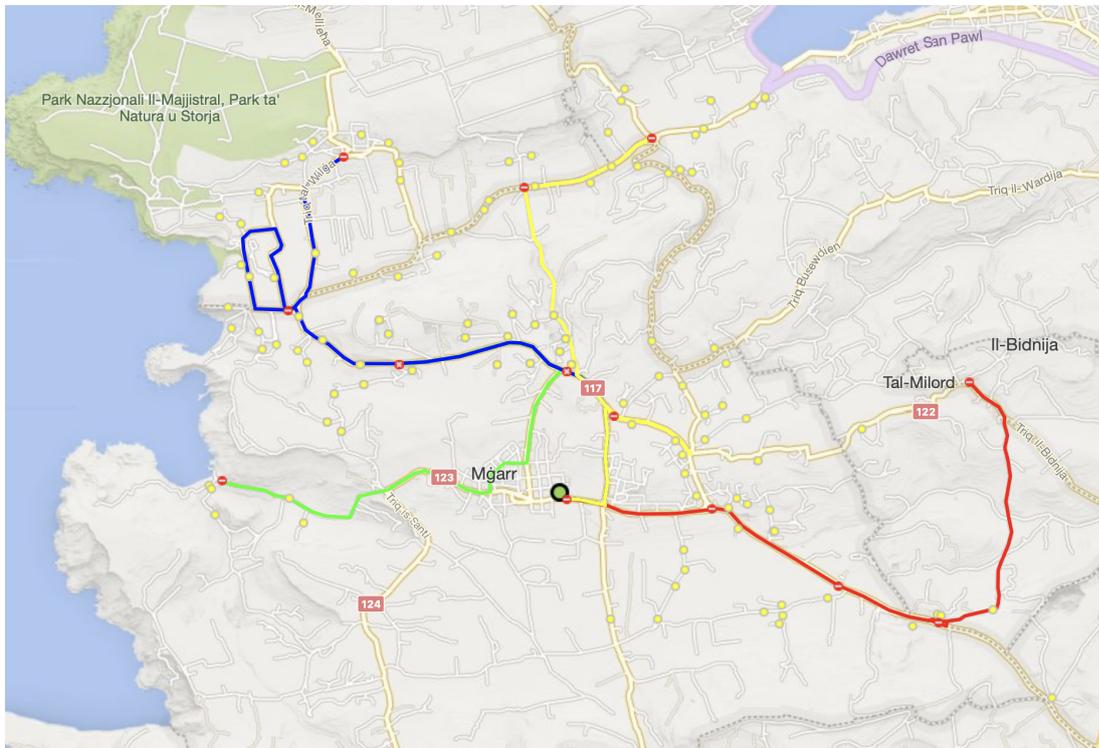


Figure D.1: Mġarr best correlated solution with 4 routes (red, lime, blue, yellow), 190 students, average walk 6.01 minutes, average journey time 13.68 minutes, and average percentile time 24.82 minutes.

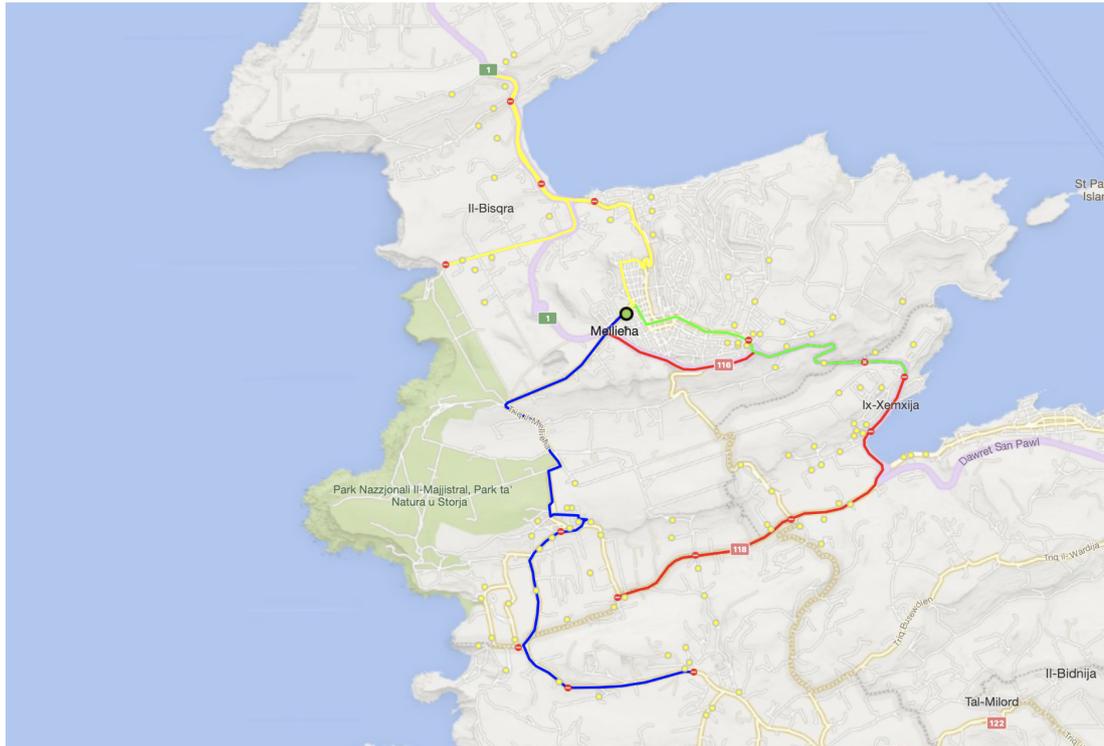


Figure D.2: Mellieħa best correlated solution with 4 routes (red, lime, blue, yellow), 171 students, average walk 5.04 minutes, average journey time 14.21 minutes, and average percentile time 25.70 minutes.

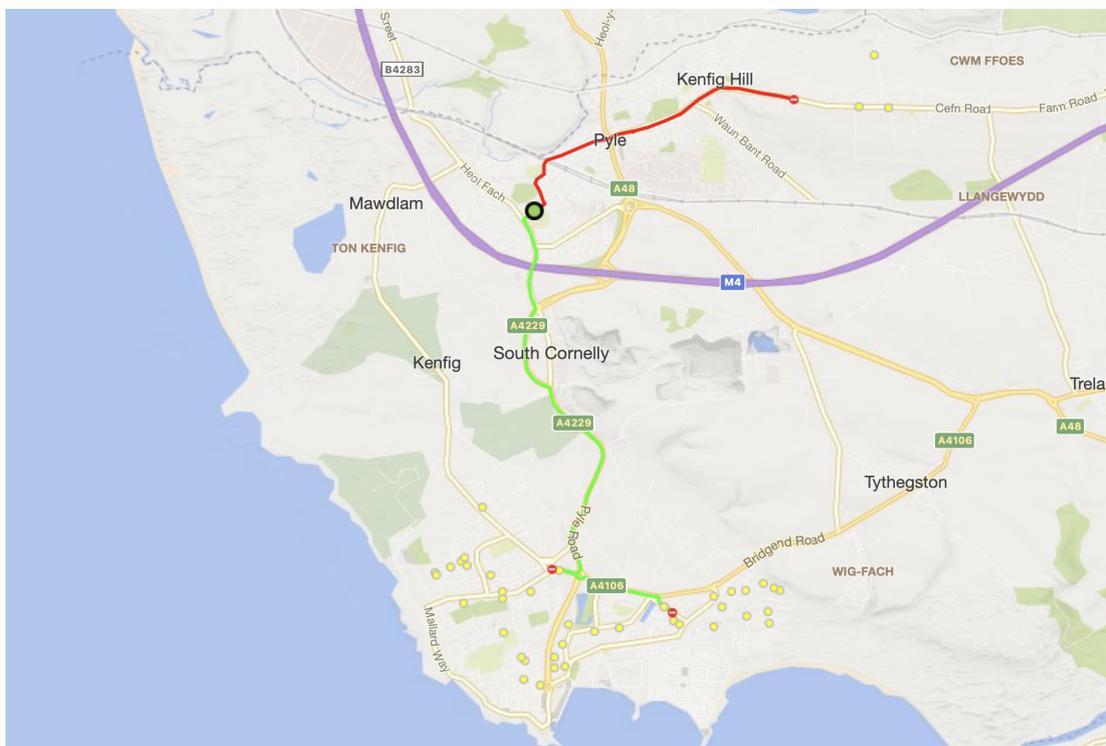


Figure D.3: Porthcawl best correlated solution with 2 routes (red, lime), 66 students, average walk 12.38 minutes, average journey time 12.14 minutes, and average percentile time 26.31 minutes.

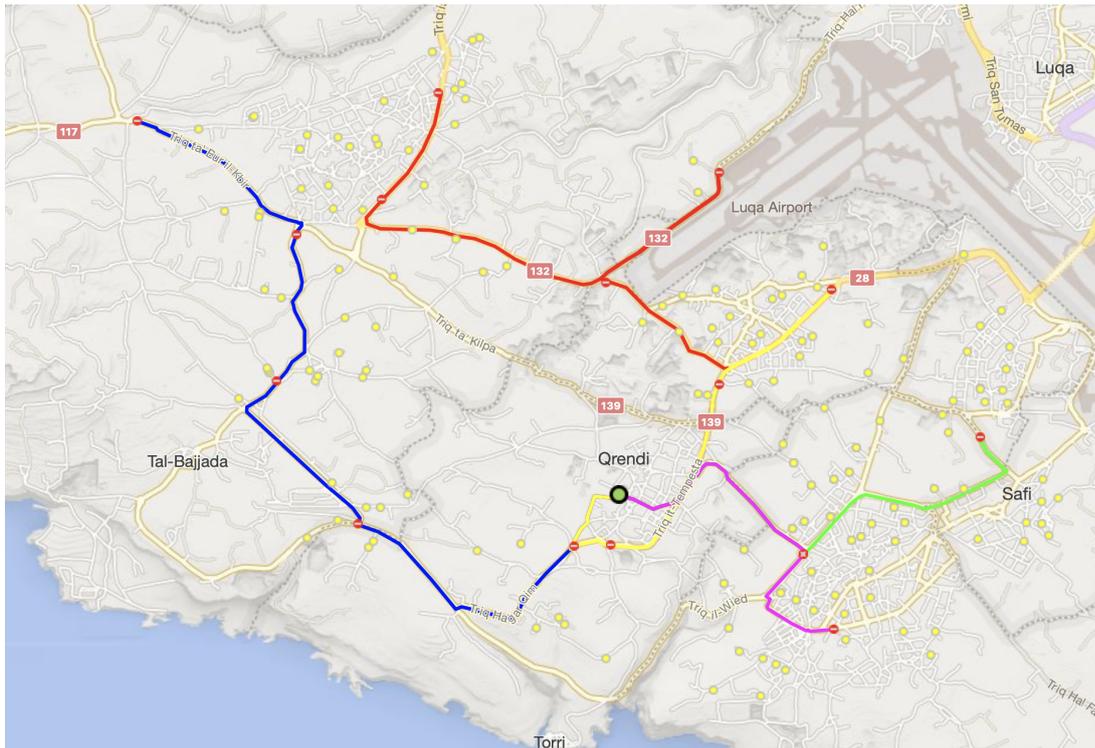


Figure D.4: Qrendi best correlated solution with 5 routes (red, lime, blue, yellow, magenta), 255 students, average walk 6.70 minutes, average journey time 15.18 minutes, and average percentile time 27.42 minutes.

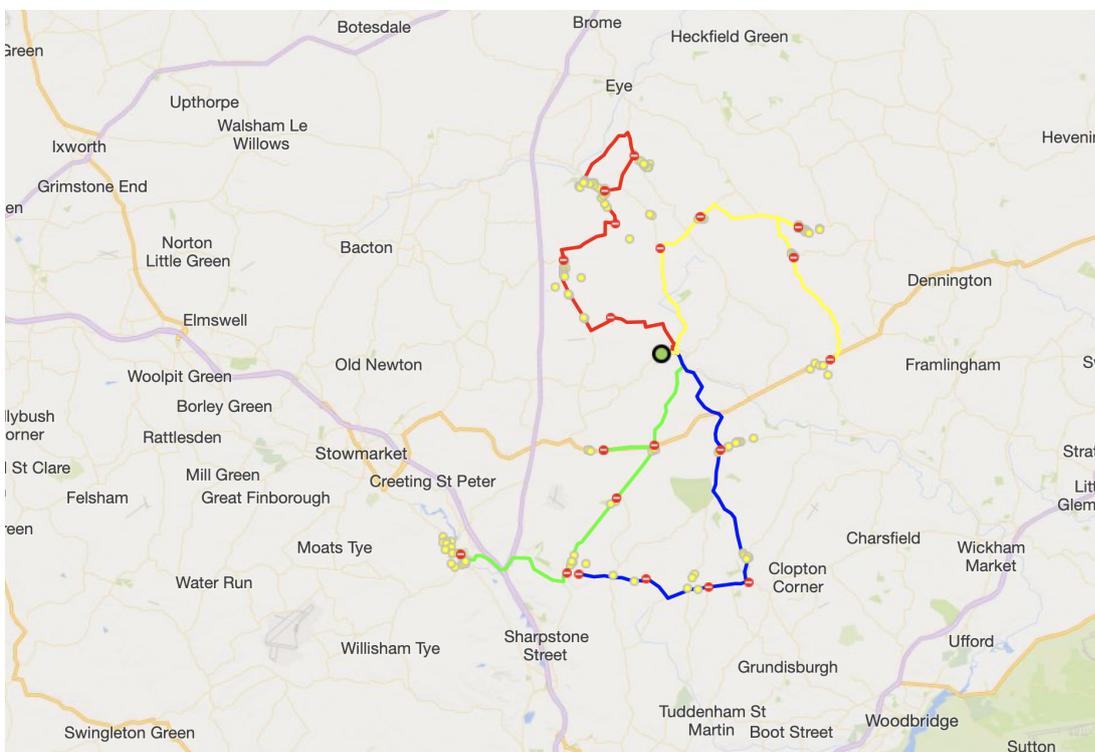


Figure D.5: Suffolk best correlated solution with 4 routes (red, lime, blue, yellow), 209 students, average walk 7.69 minutes, average journey time 29.17 minutes, and average percentile time 39.21 minutes.

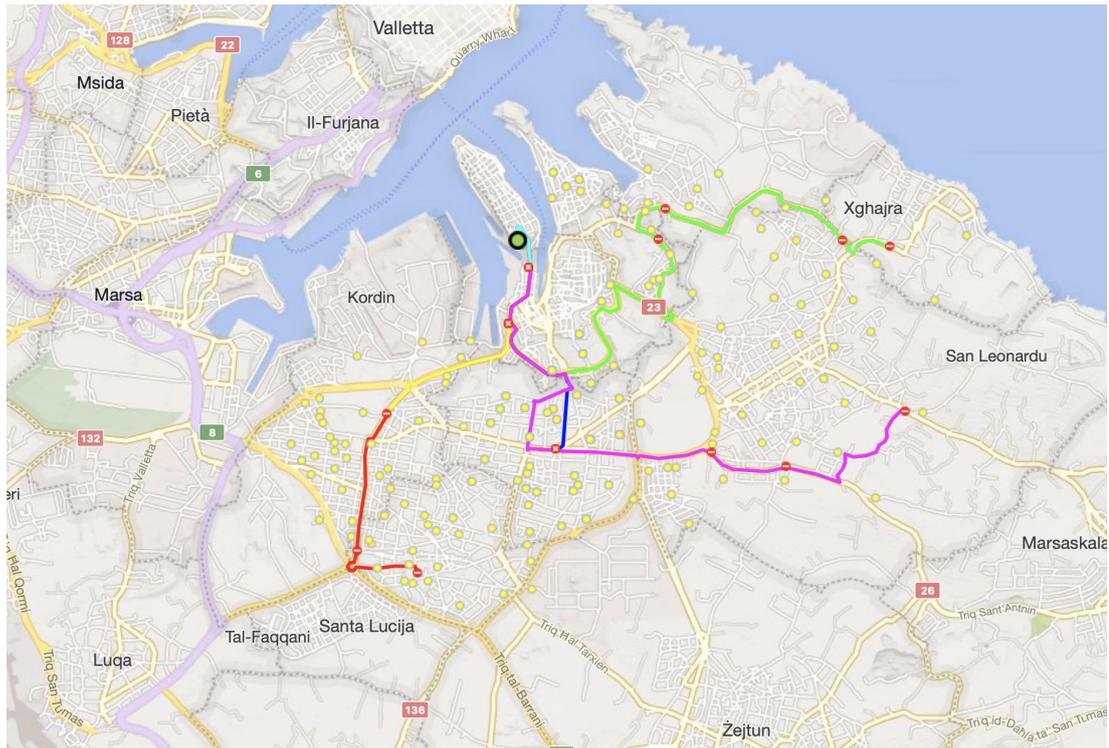


Figure D.6: Senglea best correlated solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 266 students, average walk 6.49 minutes, average journey time 12.36 minutes, and average percentile time 22.46 minutes.

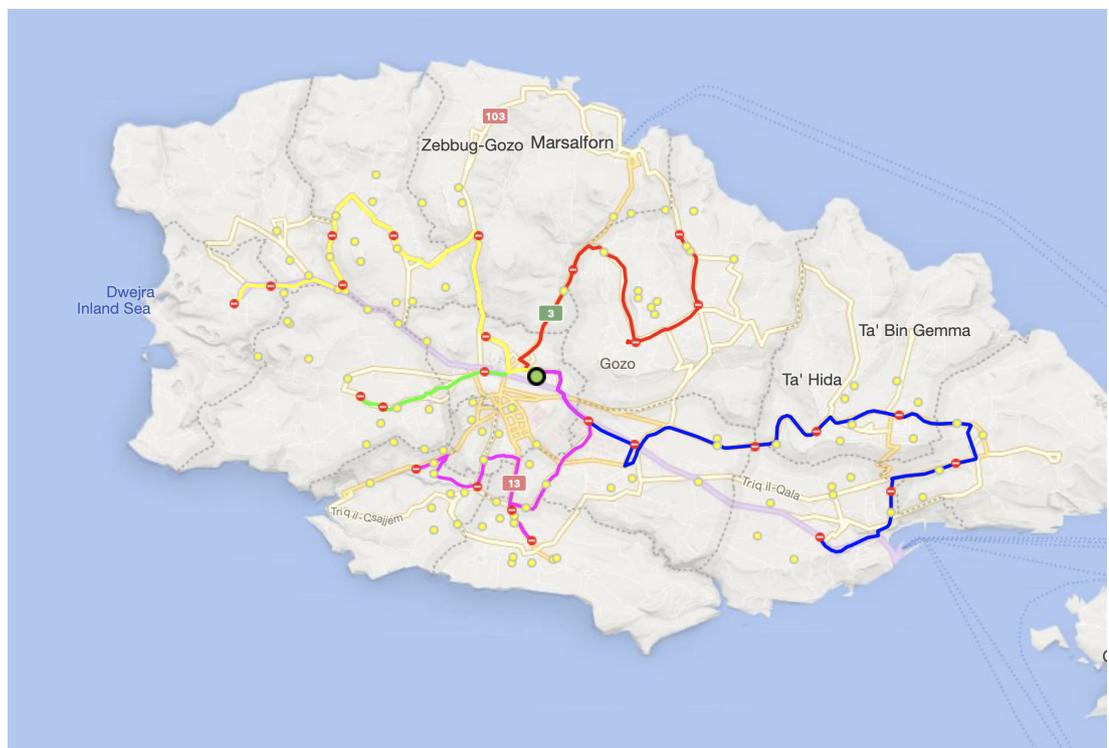


Figure D.7: Victoria best correlated solution with 5 routes (red, lime, blue, yellow, magenta), 171 students, average walk 7.29 minutes, average journey time 18.51 minutes, and average percentile time 31.94 minutes.

APPENDIX D. HEURISTIC ALGORITHM CORRELATED SCENARIO RESULTS

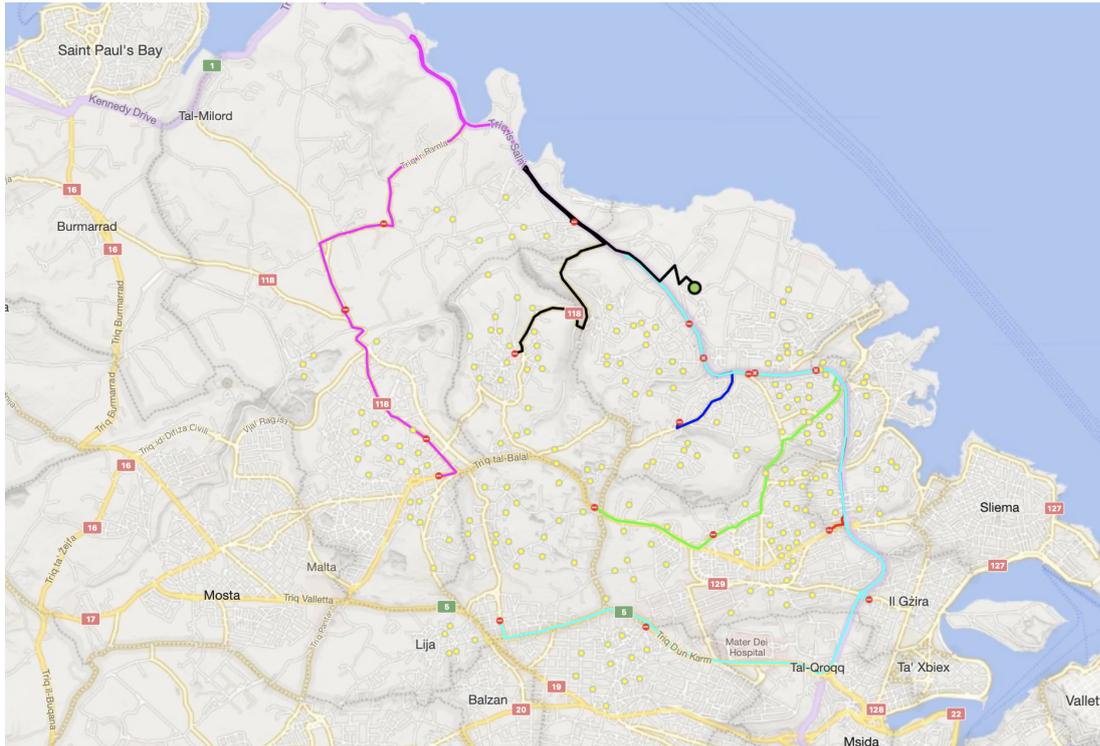


Figure D.8: Pembroke best correlated solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 335 students, average walk 7.21 minutes, average journey time 14.96 minutes, and average percentile time 27.99 minutes.

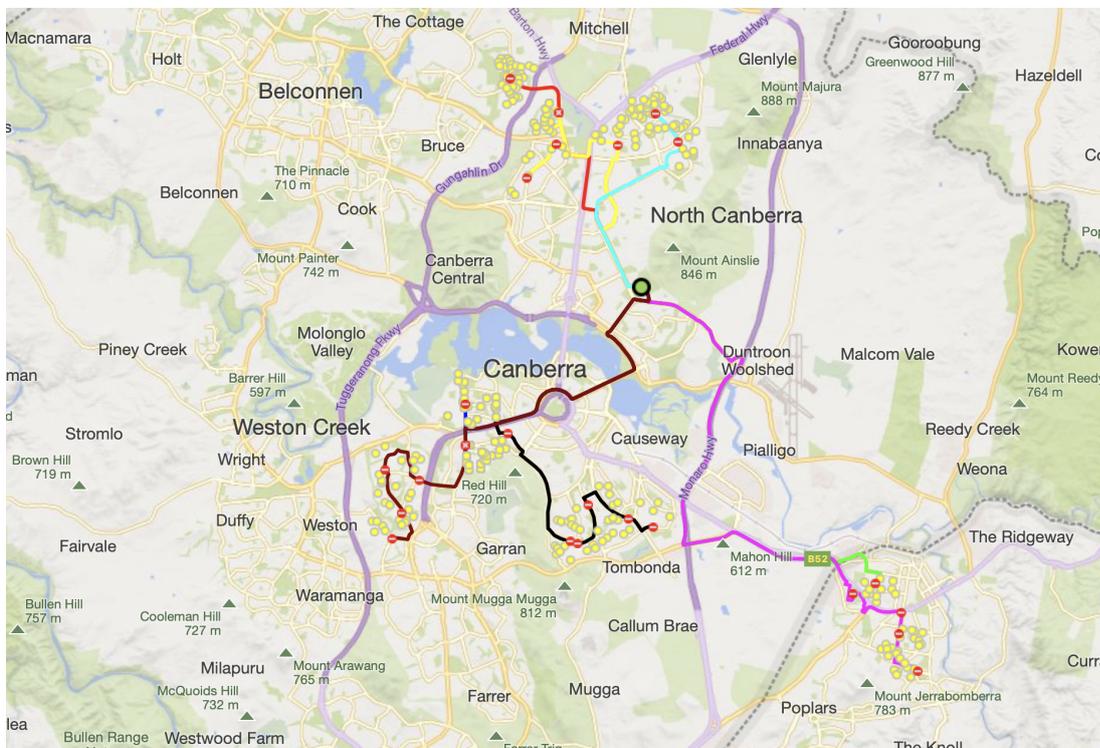


Figure D.9: Canberra best correlated solution with 8 routes (red, lime, blue, yellow, magenta, cyan, black, brown), 499 students, average walk 7.21 minutes, average journey time 23.66 minutes, and average percentile time 32.44 minutes.

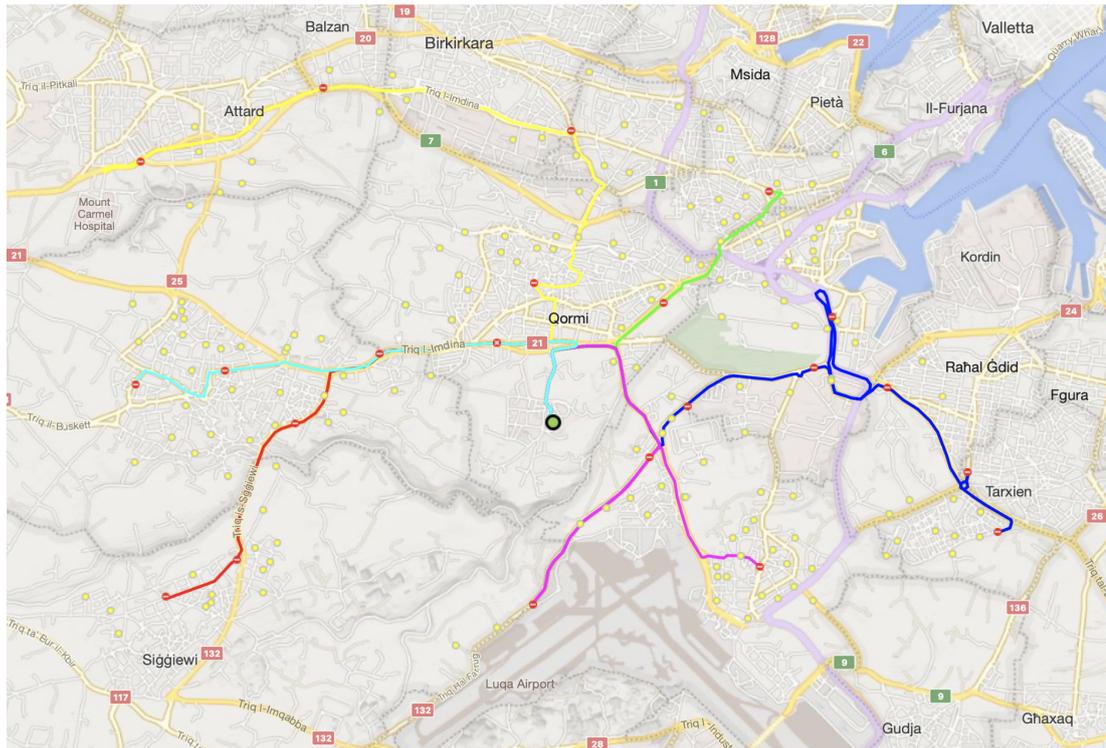


Figure D.10: Handaq best correlated solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 285 students, average walk 6.79 minutes, average journey time 16.06 minutes, and average percentile time 28.58 minutes.

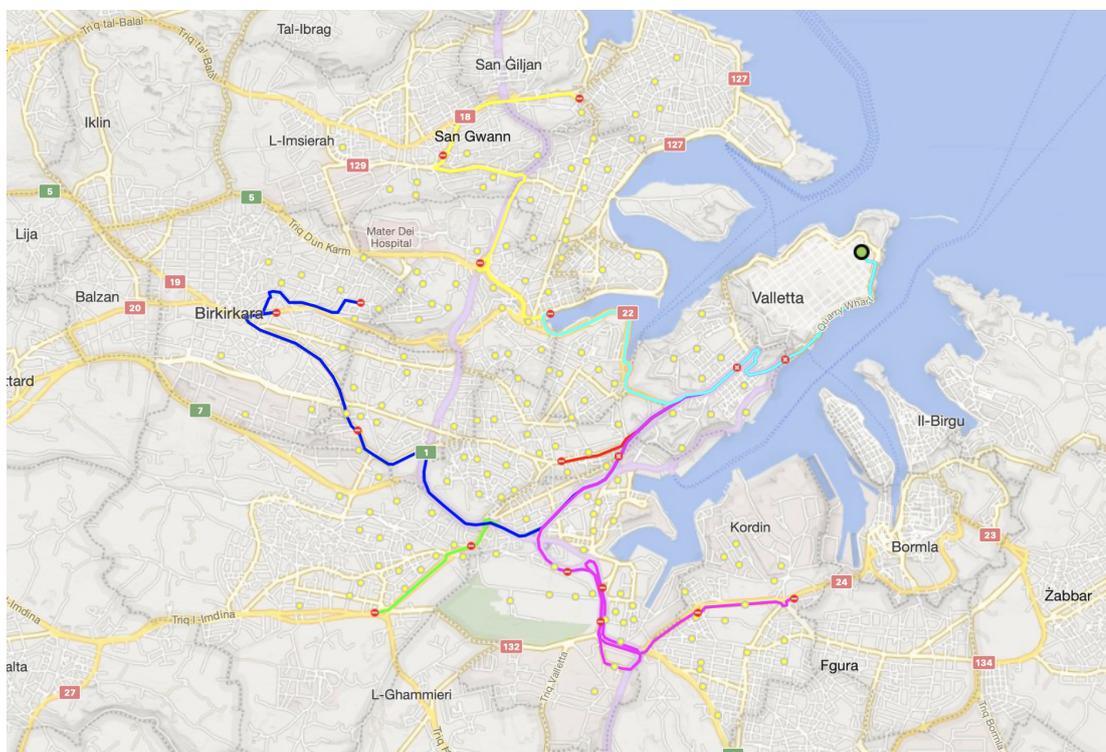


Figure D.11: Valletta best correlated solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 268 students, average walk 6.84 minutes, average journey time 18.07 minutes, and average percentile time 34.15 minutes.

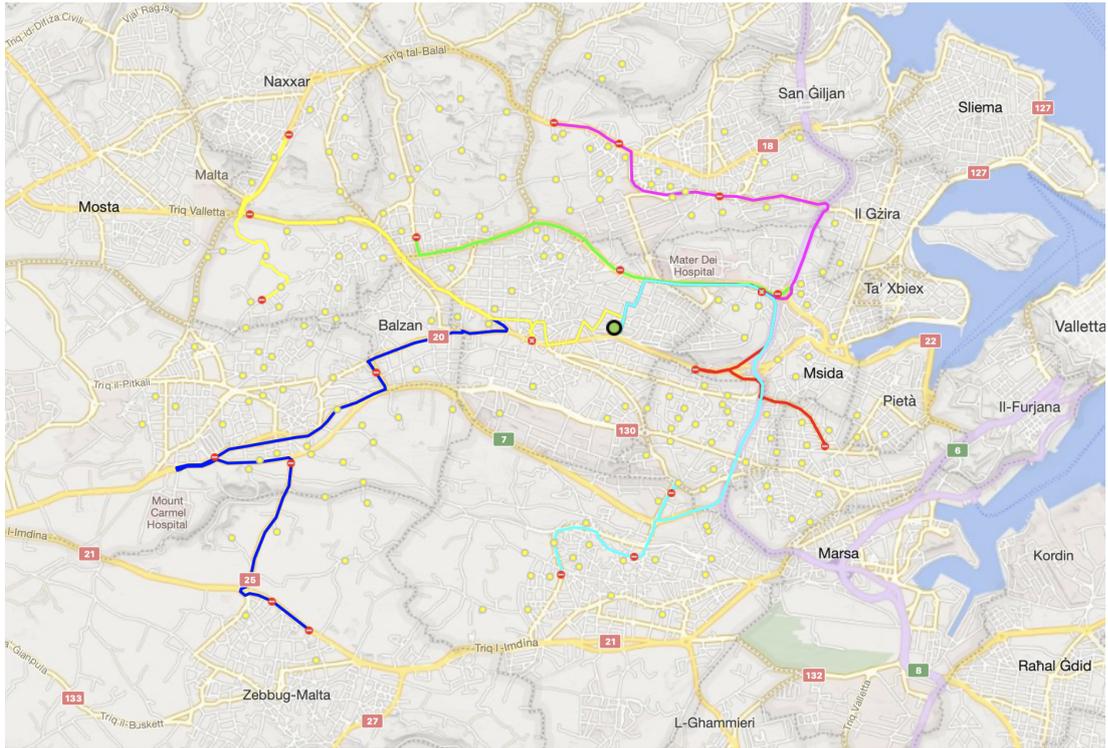


Figure D.12: Birkirkara best correlated solution with 6 routes (red, lime, blue, yellow, magenta, cyan), 306 students, average walk 6.62 minutes, average journey time 15.62 minutes, and average percentile time 27.82 minutes.

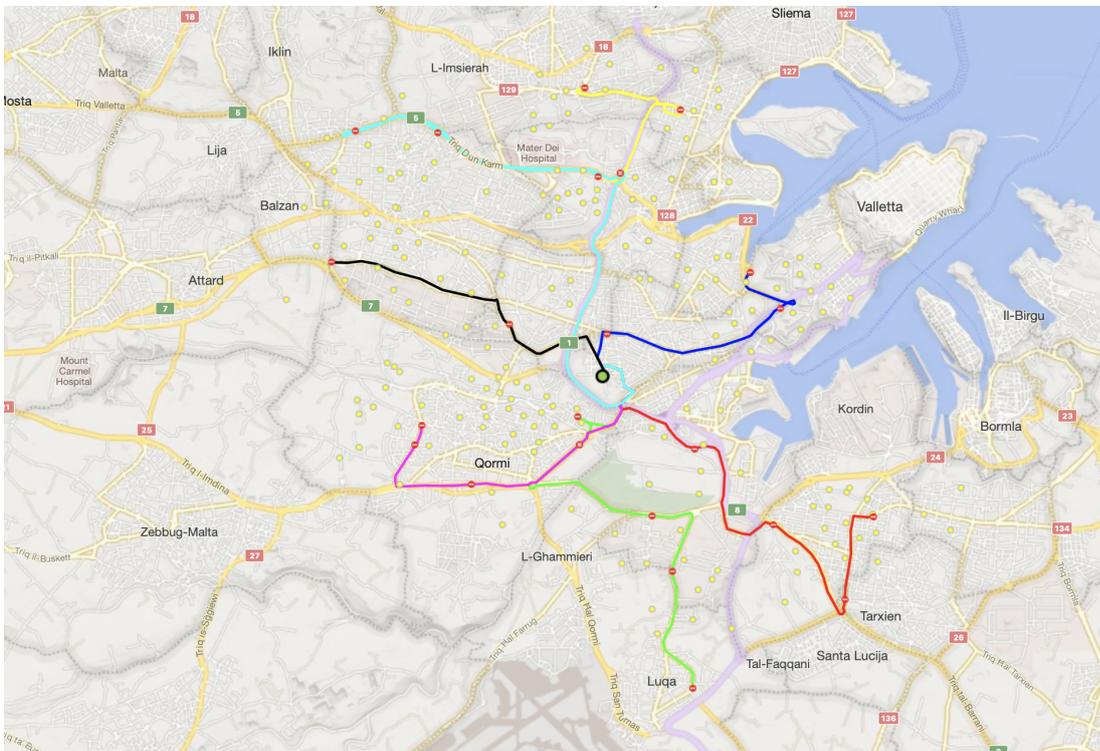


Figure D.13: Hamrun best correlated solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 321 students, average walk 6.99 minutes, average journey time 14.16 minutes, and average percentile time 25.44 minutes.

APPENDIX D. HEURISTIC ALGORITHM CORRELATED SCENARIO RESULTS

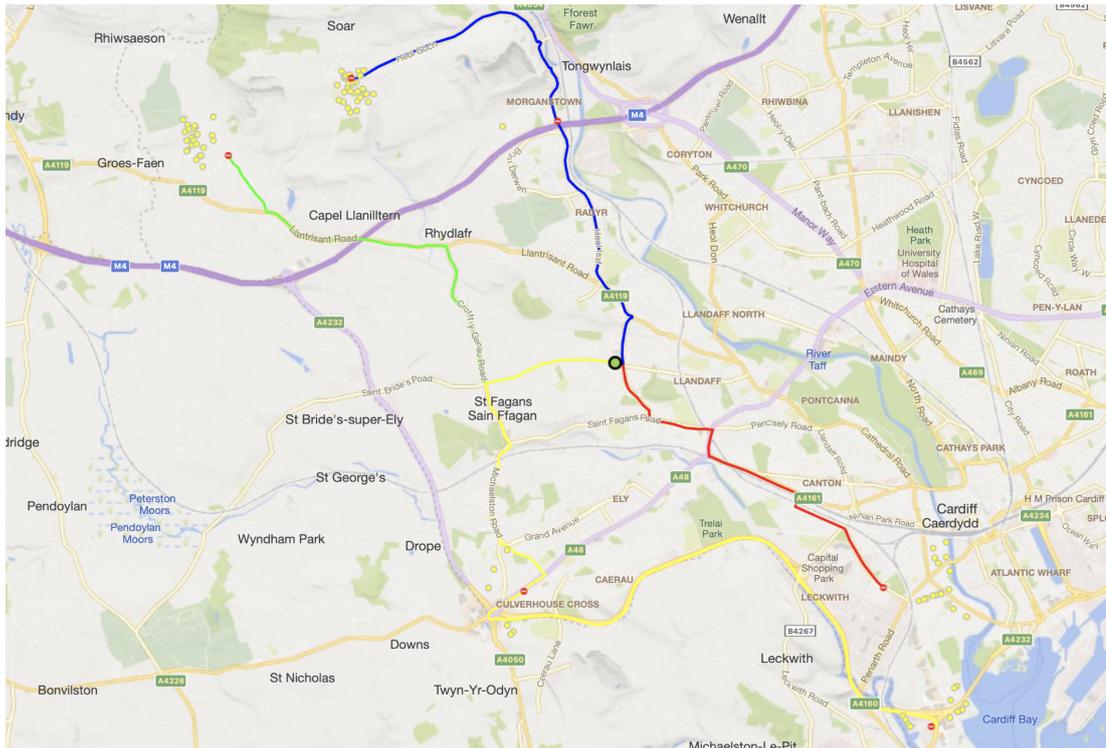


Figure D.14: Cardiff best correlated solution with 4 routes (red, lime, blue, yellow), 156 students, average walk 9.65 minutes, average journey time 17.41 minutes, and average percentile time 38.57 minutes.

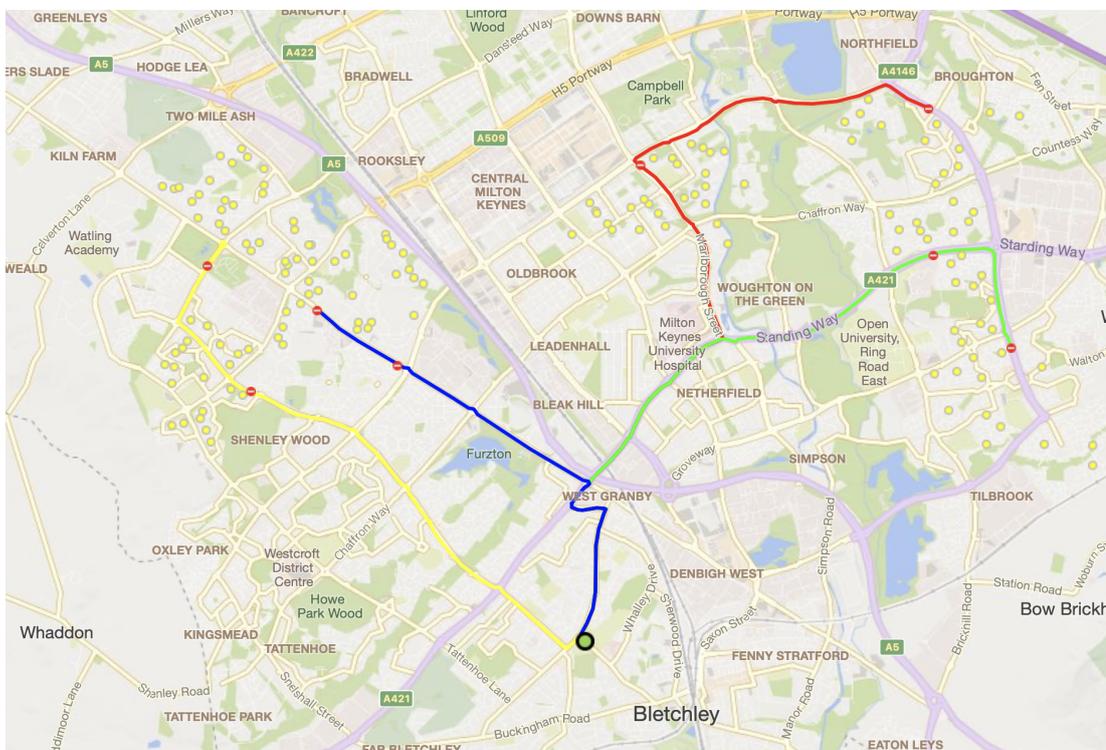


Figure D.15: Milton Keynes best correlated solution with 4 routes (red, lime, blue, yellow), 274 students, average walk 11.54 minutes, average journey time 14.76 minutes, and average percentile time 28.31 minutes.

APPENDIX D. HEURISTIC ALGORITHM CORRELATED SCENARIO RESULTS

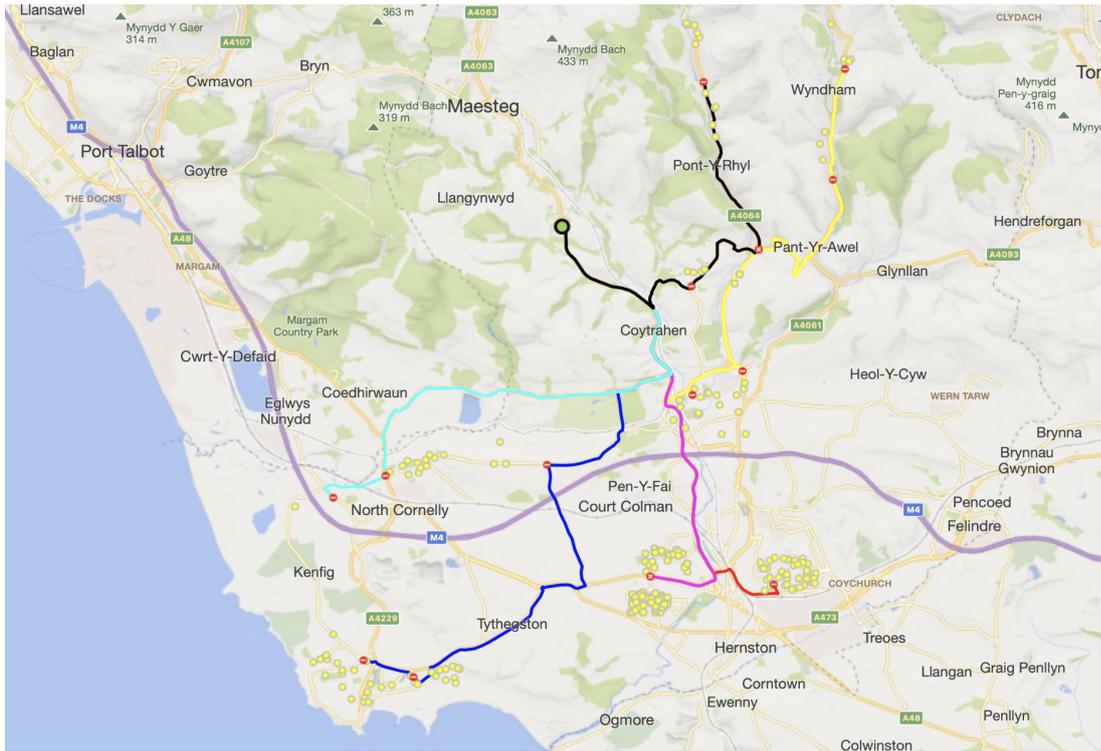


Figure D.16: Bridgend best correlated solution with 7 routes (red, lime, blue, yellow, magenta, cyan, black), 381 students, average walk 11.83 minutes, average journey time 25.43 minutes, and average percentile time 35.52 minutes.

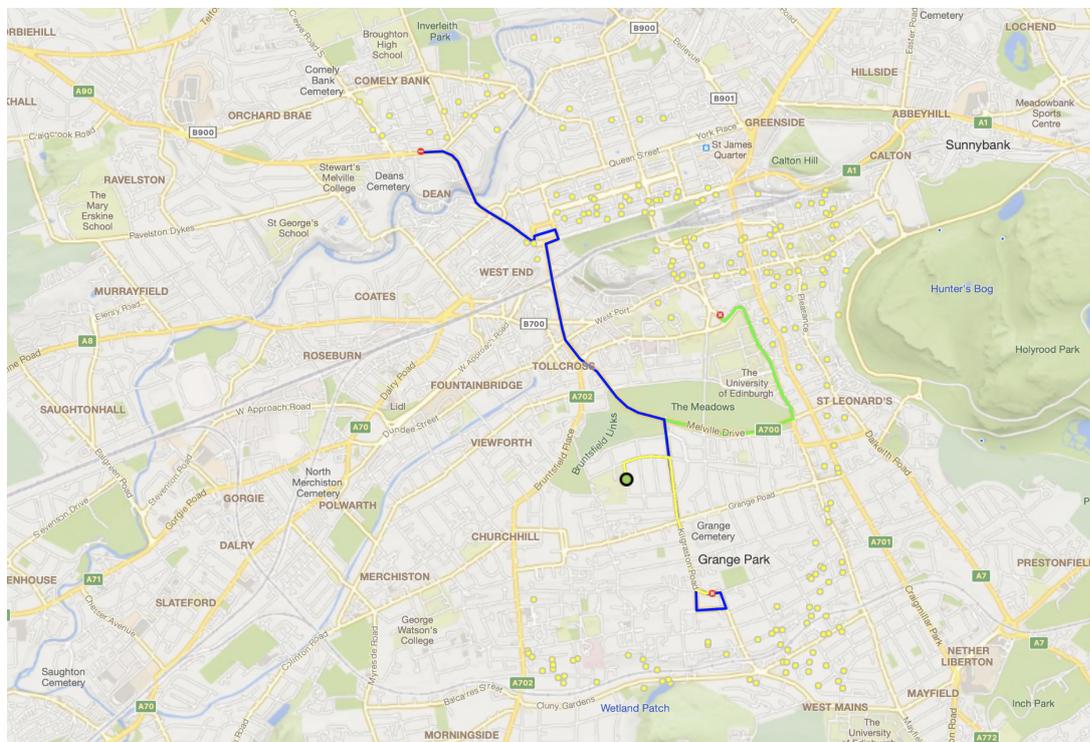


Figure D.17: Edinburgh-2 best correlated solution with 4 routes (red, lime, blue, yellow), 320 students, average walk 11.10 minutes, average journey time 14.78 minutes, and average percentile time 27.69 minutes.

APPENDIX D. HEURISTIC ALGORITHM CORRELATED SCENARIO RESULTS

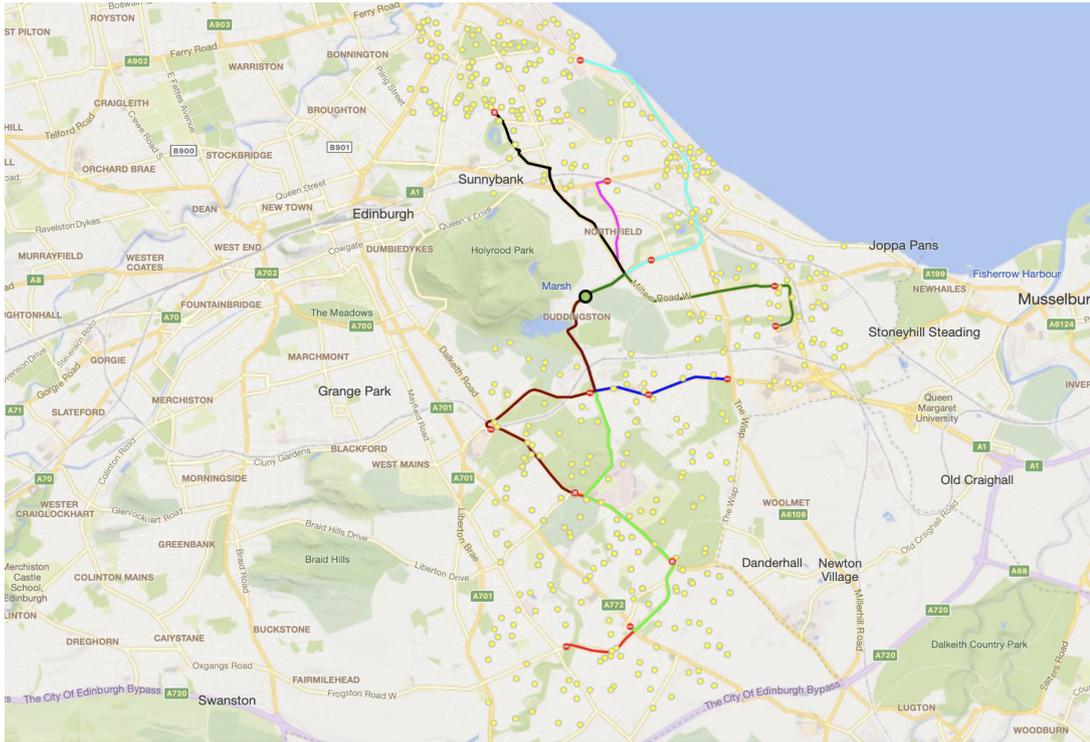


Figure D.18: Edinburgh-1 best correlated solution with 9 routes (red, lime, blue, yellow, magenta, cyan, black, brown, green), 680 students, average walk 10.24 minutes, average journey time 15.62 minutes, and average percentile time 29.24 minutes.

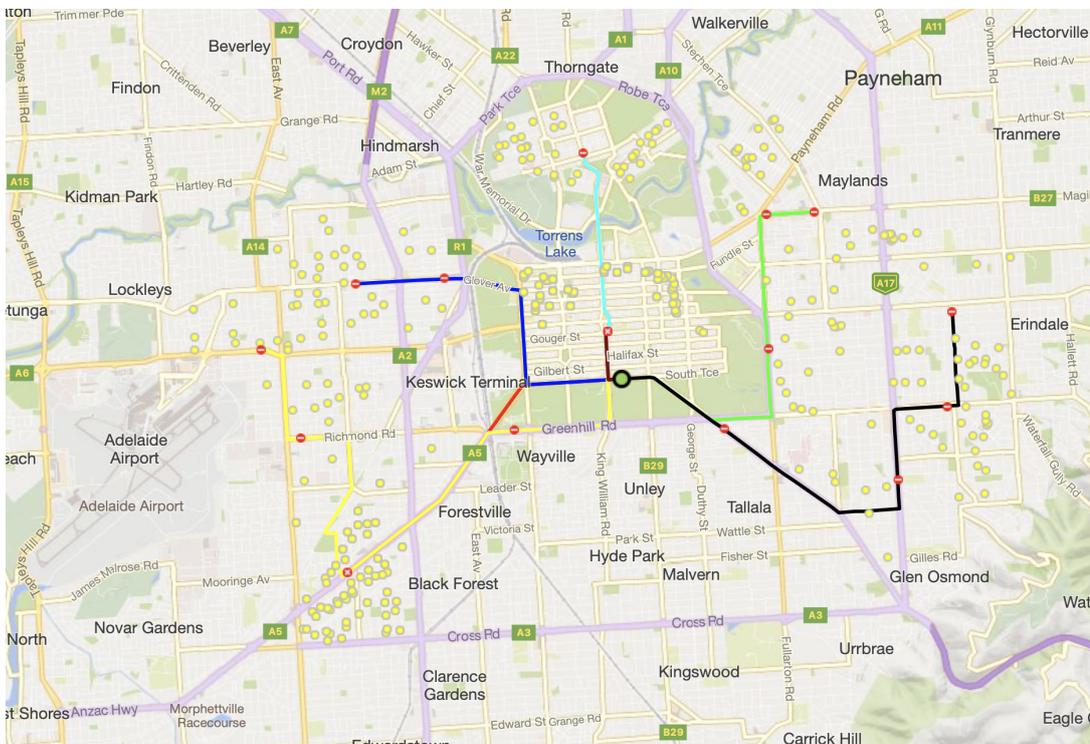


Figure D.19: Adelaide best correlated solution with 8 routes (red, lime, blue, yellow, magenta, cyan, black, brown), 565 students, average walk 12.23 minutes, average journey time 16.04 minutes, and average percentile time 28.82 minutes.

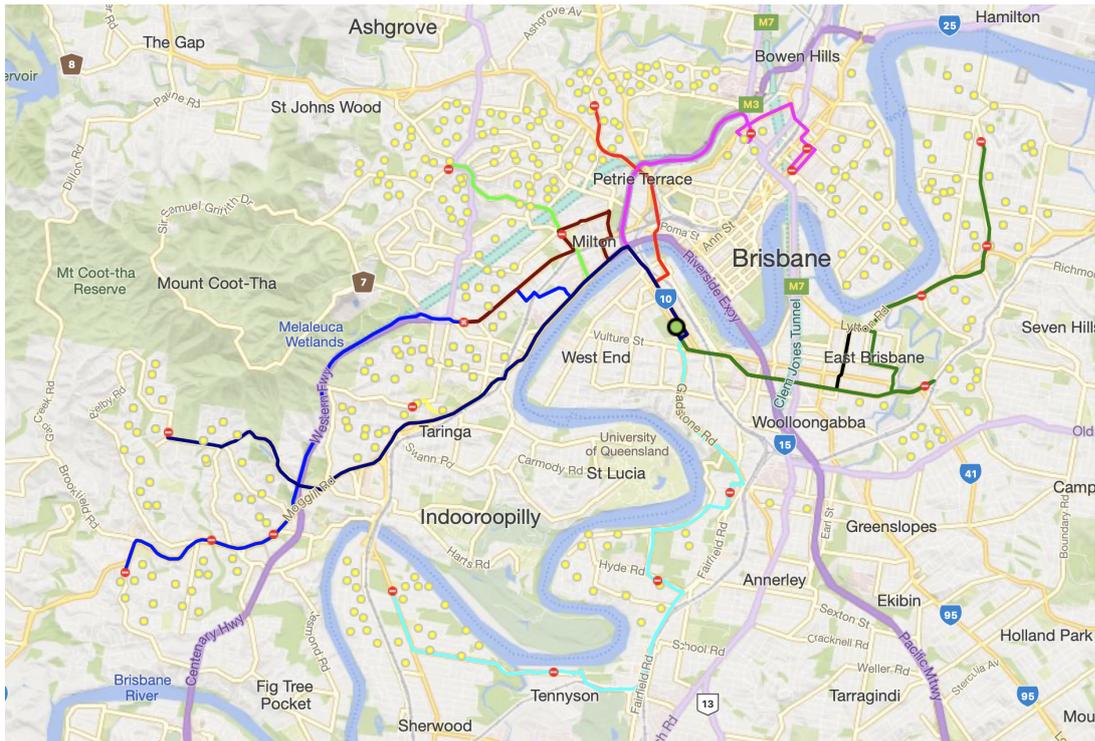


Figure D.20: Brisbane best correlated solution with 10 routes (red, lime, blue, yellow, magenta, cyan, black, brown, green, navy), 757 students, average walk 12.02 minutes, average journey time 20.53 minutes, and average percentile time 27.58 minutes.