

Stopping pattern and frequency optimization for multiple public transport services

Master Thesis TU Delft

Civil Engineering - Transport & Planning

Merlijn van Beurden

4204670

October, 2017



Stopping pattern and frequency optimization for multiple public transport services

Master of Science graduation thesis at TU Delft

Civil Engineering, track Transport and Planning

October, 2017

Author: Merlijn van Beurden
4204670
Delft University of Technology
Faculty of Civil Engineering and Geosciences
Department of Transport and Planning
Contact: merlijnvanbeurden@gmail.com

Thesis committee: *Chairman*
Prof. dr. ir. Serge Hoogendoorn
Delft University of Technology
Faculty of Civil Engineering and Geosciences
Department of Transport and Planning

Daily supervisor TU Delft
Dr. Oded Cats
Delft University of Technology
Faculty of Civil Engineering and Geosciences
Department of Transport and Planning

External supervisor TU Delft
Dr. Martijn Warnier
Delft University of Technology
Faculty of Technology, Policy and Management
Department of Multi-Actor Systems

Daily supervisor Royal HaskoningDHV
Ir. Barth Donners
Royal HaskoningDHV
Consultant Mobility - Rail & Public Transport

Supervisor Royal HaskoningDHV
Dr. ir. Geertje Hegeman
Royal HaskoningDHV
Associate Director Sustainable Mobility

Preface

This thesis is the finale of five years of study at Delft University of Technology, first as a Bachelor student Civil Engineering and now as a Master student in the track Transport and Planning. I can say with absolute certainty that the last eight months have been the greatest challenge of all these years, but I am proud of the result that now lies before you.

Of course, I could not have done this alone. First of all, I would like to thank Royal HaskoningDHV for providing me the opportunity to do this research with them. A special thank to the colleagues for welcoming me so warmly, helping me with anything I needed and teaching me everything there is to know about trains. I owe my thanks to DAT.Mobility for providing me the needed OmniTRANS licences and to Vervoerregio Amsterdam for VENOM, which were both crucial for the case study.

The first person I would like to thank personally is Oded Cats, who helped me from shaping the idea for this thesis to dotting the i's and crossing the t's. Your advice and feedback were invaluable to me. I can say the same of Barth Donners, your sharp and critical remarks have taken this thesis to a higher level. I thank the rest of my committee as well: Geertje Hegeman, for asking the right questions; Martijn Warnier, for helping me with the genetic algorithm; and Serge Hoogendoorn, for his useful comments.

I am very grateful to my girlfriend Jennifer, for coping with me these past few months and understanding my lack of time for her, especially when things did not go as well as I would have liked. I also want to thank her for roaming through Amsterdam together, looking for the perfect cover photo and succeeding perfectly.

Last but not least I would like to thank my family, friends and all those around me for providing me with the right conditions to complete my studies, not only during the last months but always. Without all of you, I would not be where I am today.

Merlijn van Beurden

October 2017

Executive summary

Growing mobility, especially in urban areas, is one of the challenges of the near future. One of many ways to tackle this problem is by shifting from private transport to public transport. However, public transport infrastructure is reaching its capacity as well and expanding this infrastructure is expensive.

Using the available infrastructure more efficiently can offer a solution. In rail transport, efficient infrastructure usage can be achieved with homogeneity: the more similar the services on a track are, the more these services can be run within a certain time. Having multiple services, which all differ in stopping pattern and stopping distance, thus has a negative effect on the efficiency of the network. This problem is studied using the following research question:

Given a public transport passenger demand, a network with multiple services and a set of possible lines, what are the optimal stopping patterns and frequencies for each line, while satisfying capacity constraints?

A model is developed to optimize stopping patterns and frequencies of a given set of lines, which can be a subset of the total public transport network. It is assumed that all tracks can be used by all vehicles in the network.

Literature review

The problem studied in this thesis falls within the field of Transit Network Planning, the process of generating and operating public transport lines. In this thesis, it is examined at the highest level of planning, the strategic level. At this level, lines and stops are defined and selected, as well as preliminary frequencies.

The literature on the Transit Network Problem is extensive, as there are many aspects that can be researched. Using a set of eight criteria, this thesis is compared to the literature, which shows that it is primarily the combination of research choices that is new and relevant. These criteria are: decision variables, passenger assignment, constraints, assessment criteria, services, optimization method, and application of the model.

When compared with previous studies, one of the most notable differences is that passengers' origins and destinations are not fixed to pre-determined in stations, but in zones. These zones are connected to one or more stations via a fixed network that is not part of the optimisation. Therefore, passengers can choose a station, or even use only the fixed network. This freedom in route choice allows passengers to choose alternatives to skipped stations, so the consequences of stopping patterns can be examined adequately.

The capacity of the infrastructure is taken into account as well: potential solutions have to fit within the network, before they are assessed using a passenger assignment model. The assignment model is connected to a genetic algorithm, which is capable of finding near-optimal solutions for problems that cannot be solved within reasonable computing time. The Transit Network Problem is such a problem and the genetic algorithm has been used extensively to find a good solution for it.

The model is applied to both a fictional network and the existing metro and train network of Amsterdam. This way, a comparison can be made between theory and practice to give context to the results of the model.

Stopping Pattern and Frequency Optimization Model

The framework of SPAFOM is shown in Figure 0.1. It shows the different components of the model.

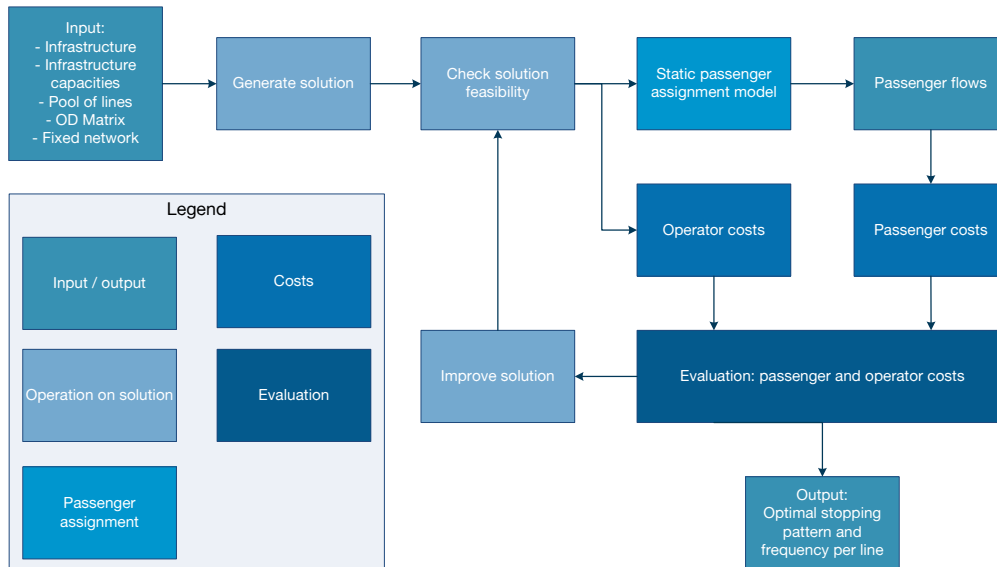


Figure 0.1: Stopping Pattern and Frequency Optimization Model framework.

The input consists of the network topology, the track and station infrastructure and its capacities, track alignments, a pool of lines to be optimized, potential stations per line, an origin-destination matrix, and a fixed underlying network. This pool of lines can be any subset of lines from the network, as long as the line is already in the network. Per line, it can also be determined which stations are part of the optimization and at which station it is fixed whether the line stops or not. The network, its infrastructure, and the travel demand for public transport is treated as given and fixed.

The optimization of the stopping patterns and frequencies take place using a genetic algorithm. A genetic algorithm follows the concept of evolution to improve solutions, in which the strong survive and the weak perish.

The process starts with an initial population, which is a set of solutions. These solutions are repaired to make sure that stopping patterns and frequencies are the same in both directions for all lines. Then, each solution is checked on whether it fits within the infrastructure capacity of the network. Solutions that do not fit are replaced with a new random solution until the population is filled.

A static, deterministic All Or Nothing passenger assignment model is run for each solution. From the results of the passenger assignment model, the passenger costs are calculated. These costs consist of walking costs, waiting costs, in-vehicle costs, transfer costs and unsatisfied demand costs. The sum of passenger and operator costs is the fitness score of a solution.

Based on this score, a predetermined amount of solutions directly passes on as elites to a new population, called the next generation. Out of all solutions, two solutions are picked based on their fitness score. The solutions are combined into two new solutions and both new solutions have a change of being changed slightly. The new solutions are added to the new population. The crossovers, each time with a new parent selection, continues until the new population contains the right amount of solutions. From here, the whole process repeats itself until it is stopped when a pre-determined amount of generations have been run.

Numerical experiments

The numerical experiments have been carried out with the network depicted in Figure 0.2. It consists of two diagonal lines and eight angular lines, connecting a total of twenty zones. Each line and each station is part of the optimization. Zone 1 acts as the centre and is the largest zone in terms of production and attraction of passengers, the corner zones 2 up to 4 are medium zones and the other zones are small. The central station and corner stations have 4 platform tracks; the other stations have 2 platform tracks. All stations have an additional 2 through tracks, which cannot be used to board or alight a vehicle. This network is simple enough to easily explain the behaviour of the model, while still allowing for walking detours and stop or line skipping, for example.




Figure 0.2: Test network with zones

SPAFOM is applied to the network for six scenarios: the base case scenario, a high demand scenario with a ten-fold increase in transport demand, an asymmetric scenario in which one of the corner zones is the largest zone, a high operator cost with a hundred-fold increase in operator costs, a passenger costs only scenario, and a scenario without capacity constraints. A population of 26 solutions per generation was used, 10 of which were saved as elites for the next generation. Test experiments showed that the resulting diversity and run time were good.

From the results of the scenarios, it can be concluded that waiting time is the most important factor in reducing the total costs. In each infrastructure capacity constraint scenario, the central station was the bottleneck of the network. Stopping patterns were adapted in such a way, that the

frequencies of as many angular lines as possible could be increased at the expense of the shortened lines. This frequency increase made transferring less costly, due to the decrease in waiting time, which resulted in the operational costs of the diagonal line exceeding the corresponding passenger benefits. One or both of the diagonal lines were therefore in almost all scenarios removed from the network.

Amsterdam case study

SPAFOM was then applied to the metro and rail network of Amsterdam for the year 2040 and compared to the original situation, which is the network for 2040 as it is predicted in 2017. The bus and tram network were part of the passenger assignment model as well, but only as an underlying network. This means this network could be used by all passengers. Some changes to the pool of lines were adopted to include prevailing ideas that circulate in planning circles. The North-South metro line was extended to Hoofddorp and regional trains between Amsterdam Centraal and Amstel were routed through the metro tunnel between Amsterdam Centraal and Amstel, to observe the consequences of these plans. No capacity constraints were applied, as the numerical experiments showed that these significantly influenced the result. For the Amsterdam case study, the focus is on how the network should develop towards 2040. By looking this far towards the future, capacities become less relevant, as there is time to increase them if necessary. Using the current capacities would limit the outcome of the model more than necessary.

Due to the larger network, more solutions and elites per generations were used: 32 and 16 respectively. The resulting network after 669 generations is shown in Figure 0.3, which shows only the lines included in the final solution. The same trend can be observed as in the numerical experiments, namely that the frequencies of lines with many stops have increased the most to reduce passenger waiting time. Without the capacity constraints, the lines run a factor 2 more often than possible at this time. Another reason is that these lines, which are all metro lines, are shorter than the other lines, so the additional operator costs for increasing frequencies are relatively low compared to regional and national lines. Reasoning the other way around, this is also the reason the frequencies of most train lines did not increase as much as the metro lines.

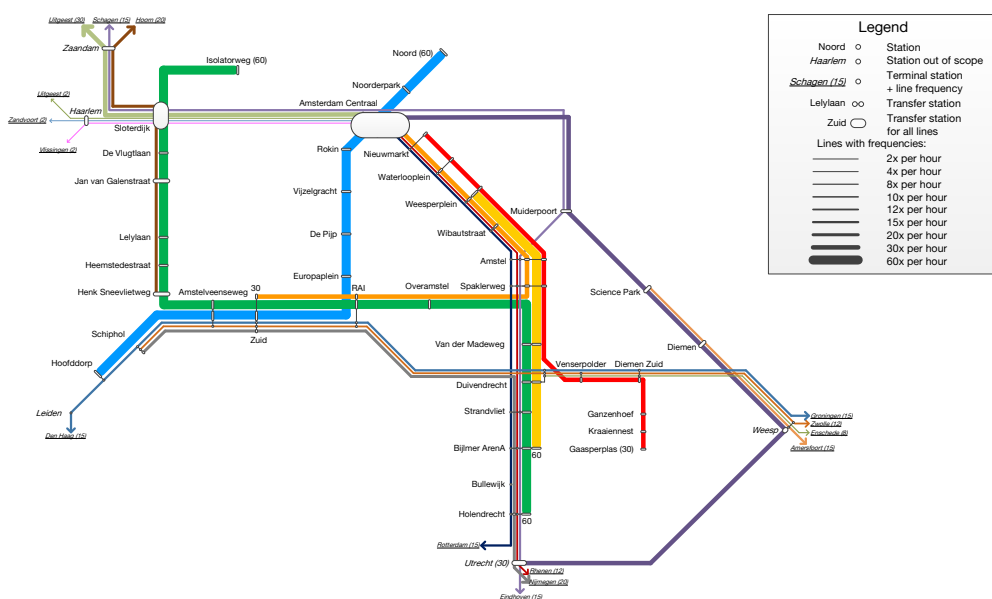


Figure 0.3: Final solution Amsterdam case study

The final solution does not show a stop of the North-South metro line at Schiphol Airport. This is proof that an optimal solution cannot be guaranteed, as the same solution with a stop at Schiphol resulted in a decrease in all costs, except for a negligible increase in operator costs. As the extension to Hoofddorp is already part of the final solution, it can be concluded that the benefits for passengers outweigh the extra operator costs. However, fixed costs such as building costs are not taken into account and neither is capacity.

The regional train lines through the metro tunnel between Amsterdam Centraal and Amstel make some of the metro lines there redundant to a certain extent. Two of the three lines are cut short before reaching the original end point of Amsterdam Centraal, as this station is mostly used to transfer between regional and national train lines. Passengers between the centre of Amsterdam and Amsterdam Centraal can take a regional train or the remaining metro line. This shows that there are certainly parts of Amsterdam where a different type of service can improve service performance for passengers.

Limitations

One of the largest limitations of this research is the calculation of passenger waiting times, which was done using a uniform arrival pattern assumption. Consequently, the effect of increasing frequencies, thereby reducing average waiting time, might be overestimated. The frequencies in the final solution should therefore not be taken too literally, but merely as an indication of which lines are most important. In a capacity constrained model run, these lines would most likely still have the highest frequency, but not as high as now.

Conclusions

The Stopping Pattern and Frequency Optimization Model (SPAFOM) has proved to be capable of systematically improving stopping patterns and frequencies for a predetermined set of lines within a public transport network, taking into account both passenger and operator costs and capacity constraints.

From the results of the model, it can be concluded that it is better to have a few lines with many stops and high frequencies, than many lines with different stopping patterns and low frequencies. Especially when lines run through densely populated areas, which metro lines almost always do, they are very important for that area. If multiple lines run through a bottleneck, some of them should be shortened so other lines' frequencies can be increased. As a result of the high frequencies, transferring becomes less of an issue as waiting time decreases. Passengers do not mind a slightly longer in-vehicle time, if this means they have to walk less and/or wait for a shorter time. The increased frequencies also increase seating capacity on a line, making the in-vehicle time more comfortable and thus less costly.

The model itself is applicable to existing networks. Despite the longer run time for larger and more complex network, it can handle any network and any number of lines, no matter what kind of vehicle serves the transit lines. The results can be used as a starting point for the joint design of public transport services, which were up to now separated, into one integrated public transport network.

Table of Contents

Preface	I
Executive summary	II
List of figures.....	1
List of tables.....	4
1 Introduction.....	5
1.1 Problem definition.....	6
1.2 Research objectives.....	7
1.3 Research scope	7
1.4 Stakeholders and actors	8
1.5 Reading guide	8
2 Background and literature review	9
2.1 Background on Transit Network Planning	9
2.2 Literature overview.....	10
2.2.1 Decision variables.....	13
2.2.2 Passenger assignment.....	13
2.2.3 Constraints.....	13
2.2.4 Assessment criteria.....	14
2.2.5 Services	14
2.2.6 Optimization method	14
2.2.7 Application.....	16
2.3 Relevance.....	17
2.3.1 Scientific contributions	17
2.3.2 Practical contributions.....	17
2.4 Conclusion of the literature review.....	19
3 Modelling framework	20
3.1 Framework.....	20
3.2 Components of SPAFOM	22
3.2.1 Input.....	22
3.2.2 Generate solutions	22
3.2.3 Check solution feasibility	22
3.2.4 Assignment model.....	23
3.2.5 Passenger costs	23
3.2.6 Operator costs	24

3.2.7	Evaluation	24
3.2.8	Improve solution	24
3.3	SPAFOM information flow chart.....	24
3.4	Conclusion of the framework.....	26
4	Model formulation	27
4.1	Notations	27
4.2	Cost functions.....	28
4.2.1	Input.....	28
4.2.2	Passenger costs	28
4.2.3	Operator costs	31
4.2.4	Objective function	32
4.3	Constraints.....	32
4.3.1	Stopping pattern and frequency constraints	32
4.3.2	Node capacity	33
4.3.3	Link capacity	34
4.3.4	Demand satisfaction.....	34
4.4	Conclusion of the model formulation.....	34
5	Model implementation.....	35
5.1	Introduction to genetic algorithms.....	35
5.2	Introduction to OmniTRANS databases	37
5.3	SPAFOM implementation	37
5.3.1	Input.....	39
5.3.2	Generate solutions	40
5.3.3	Check solution feasibility	40
5.3.4	Implement solution in OmniTRANS.....	42
5.3.5	Passenger costs	42
5.3.6	Operator costs	43
5.3.7	Evaluation	44
5.3.8	Improve solution	44
5.3.9	Output	46
5.3.10	Model verification and validation	47
5.4	Conclusion of the model implementation.....	47
6	SPAFOM numerical experiments.....	48
6.1	Model specifications.....	48

6.1.1	Genetic algorithm	48
6.1.2	Other parameters	49
6.2	Numerical experiments	53
6.2.1	Network.....	53
6.2.2	Scenarios.....	56
6.2.3	Results.....	59
6.2.4	Sensitivity analysis	79
6.3	Limitations of the numerical experiments	82
6.4	Conclusions of the numerical experiments	83
7	Case study: Amsterdam.....	84
7.1	Model input.....	85
7.2	Optimization process	92
7.3	Results.....	95
7.3.1	Costs.....	95
7.3.2	Network.....	97
7.3.3	Service performance	102
7.3.4	Network performance	106
7.3.5	Transfers	108
7.3.6	Walking.....	111
7.4	Feasibility of the results	112
7.5	Limitations.....	113
7.6	Conclusions of the Amsterdam case study.....	115
8	Conclusions and recommendations	116
8.1	Conclusions.....	116
8.2	Practical implications.....	119
8.3	Limitations.....	120
8.4	Future research.....	123
8.4.1	Model improvements	123
8.4.2	Usage of results.....	124
	Bibliography	125
	Appendix A: Definitions	129
	Appendix B: Solution data current and final solution.....	130
	Appendix C: Change in transfers per station, sorted.....	131

List of figures

FIGURE 0.1: STOPPING PATTERN AND FREQUENCY OPTIMIZATION MODEL FRAMEWORK.	III
FIGURE 0.2: TEST NETWORK WITH ZONES.....	IV
FIGURE 0.3: FINAL SOLUTION AMSTERDAM CASE STUDY.....	V
FIGURE 1.1: SERVICES (ADAPTED FROM (GOVERS, 2017)).....	5
FIGURE 1.2: STOP SPACING, LINE SPACING AND FREQUENCY (VAN NES, 2015).....	6
FIGURE 2.1: TRANSIT PLANNING PROCESS (IBARRA-ROJAS <i>ET AL.</i> , 2015).....	9
FIGURE 3.1: STOPPING PATTERN AND FREQUENCY OPTIMIZATION MODEL FRAMEWORK.	21
FIGURE 3.2: SCHEMATIC REPRESENTATION OF A JOURNEY OF A PASSENGER	23
FIGURE 3.3: STOPPING PATTERN AND FREQUENCY OPTIMIZATION MODEL INFORMATION FLOW CHART	25
FIGURE 5.1: GENETIC ALGORITHM	35
FIGURE 5.2: GENERATIONS, CHROMOSOMES AND GENES	36
FIGURE 5.3: SOLUTION REPRESENTATION.....	36
FIGURE 5.4: FREQUENCY SOLUTION REPRESENTATION	37
FIGURE 5.5: MODEL IMPLEMENTATION FLOW CHART	38
FIGURE 5.6: NODE AND LINK CAPACITY	41
FIGURE 5.7: COMPOSITION OF A NEW GENERATION WITH ELITES (E) AND Crossover (C).....	44
FIGURE 5.8: EXAMPLE ROULETTE WHEEL SELECTION	45
FIGURE 5.9: EXAMPLE UNIFORM Crossover.....	45
FIGURE 6.1: PARAMETER IN-VEHICLE TIME	50
FIGURE 6.2: NS VIRM TRAIN	51
FIGURE 6.3: TEST NETWORK. DISTANCE IN KM.....	53
FIGURE 6.4: PASSENGER FLOWS IN BASE CASE SCENARIO.....	55
FIGURE 6.5: DESIGNED SOLUTIONS FOR THE TEST NETWORK SCENARIOS	57
FIGURE 6.6: MOVEMENT PATTERN BETWEEN CENTROIDS	58
FIGURE 6.7: EVOLUTION BASE CASE SCENARIO	59
FIGURE 6.8: COST COMPARISON START AND FINAL SOLUTION BASE CASE SCENARIO.....	60
FIGURE 6.9: FINAL SOLUTION BASE CASE SCENARIO AND PASSENGER FLOWS PER LINK	61
FIGURE 6.10: EVOLUTION HIGH DEMAND SCENARIO	63
FIGURE 6.11: COST COMPARISON FIRST AND FINAL SOLUTION HIGH DEMAND SCENARIO (OPERATOR COSTS ON SEPARATE AXIS).....	64
FIGURE 6.12: FINAL SOLUTION HIGH DEMAND SCENARIO.....	65
FIGURE 6.13: EVOLUTION ASYMMETRIC DEMAND SCENARIO.....	66
FIGURE 6.14: COST COMPARISON FIRST AND FINAL SOLUTION ASYMMETRIC SCENARIO.....	66
FIGURE 6.15: FINAL SOLUTION ASYMMETRIC SCENARIO AND PASSENGER FLOWS PER LINK	67
FIGURE 6.16: EVOLUTION HIGH OPERATOR COSTS SCENARIO	68
FIGURE 6.17: COSTS COMPARISON FIRST AND FINAL SOLUTION HIGH OPERATOR COSTS SCENARIO	69
FIGURE 6.18: FINAL SOLUTION HIGH OPERATOR COSTS SCENARIO AND PASSENGER FLOWS PER LINK	70
FIGURE 6.19: EVOLUTION PASSENGER COSTS ONLY SCENARIO.....	71
FIGURE 6.20: COST COMPARISON PASSENGER COSTS ONLY SCENARIO	71
FIGURE 6.21: FINAL SOLUTION PASSENGER COSTS ONLY SCENARIO AND PASSENGER FLOWS PER LINK.....	73
FIGURE 6.22: EVOLUTION UNLIMITED INFRASTRUCTURE CAPACITY SCENARIO	74
FIGURE 6.23: COSTS COMPARISON FIRST AND FINAL SOLUTION UNLIMITED INFRASTRUCTURE CAPACITY SCENARIO .	74

FIGURE 6.24: FINAL SOLUTION UNLIMITED INFRASTRUCTURE CAPACITY SCENARIO AND PASSENGER FLOW PER LINK	75
FIGURE 6.25: COMPARISON OF PASSENGER AND OPERATOR COSTS PER PASSENGER PER SCENARIO	76
FIGURE 6.26: COMPARISON OF COSTS PER PASSENGER	77
FIGURE 6.27: COMPARISON OF OPERATOR COSTS PER PASSENGER.....	78
FIGURE 6.28: RESULTS SENSITIVITY ANALYSIS -25% UNSATISFIED DEMAND COSTS.....	79
FIGURE 6.29: COST COMPARISON FIRST AND FINAL SOLUTION SENSITIVITY ANALYSIS -25% UNSATISFIED DEMAND COSTS.....	80
FIGURE 6.30: FINAL SOLUTION SENSITIVITY ANALYSIS -25%	80
FIGURE 6.31: RESULTS SENSITIVITY ANALYSIS +25% UNSATISFIED DEMAND COSTS.....	81
FIGURE 6.32: COST COMPARISON FIRST AND FINAL SOLUTION SENSITIVITY ANALYSIS +25% UNSATISFIED DEMAND COSTS.....	81
FIGURE 6.33: FINAL SOLUTION SENSITIVITY ANALYSIS +25% UNSATISFIED DEMAND COSTS	82
FIGURE 7.1: LOCAL (S-BAHN) TRAIN IN KARLSRUHE AT SHARED TRAM AND TRAIN STATION	84
FIGURE 7.2: ALL TRANSIT LINES IN STUDY AREA	86
FIGURE 7.3: AVAILABLE LINKS AND NODES FOR AMSTERDAM CASE STUDY	88
FIGURE 7.4: CURRENT STOPPING PATTERN AND FREQUENCIES OF LINES TO BE OPTIMIZED	89
FIGURE 7.5: NUMBER OF PASSENGERS PRODUCED PER ZONE WITHIN STUDY AREA FOR MORNING PEAK HOUR, INCLUDING LINES TO OPTIMIZE	90
FIGURE 7.6: NUMBER OF PASSENGERS ATTRACTED PER ZONE WITHIN STUDY AREA FOR MORNING PEAK HOUR, INCLUDING LINES TO OPTIMIZE	91
FIGURE 7.7: RESULTS AMSTERDAM CASE STUDY	95
FIGURE 7.8: COST COMPONENTS COMPARISON CURRENT SITUATION, FIRST SOLUTION AND FINAL SOLUTION AMSTERDAM CASE STUDY.....	96
FIGURE 7.9: TOTAL COSTS COMPARISON CURRENT SITUATION, FIRST SOLUTION AND FINAL SOLUTION AMSTERDAM CASE STUDY.....	97
FIGURE 7.10: NETWORK FINAL SOLUTION AMSTERDAM CASE STUDY.....	98
FIGURE 7.11: MAGNIFICATION AMSTERDAM CENTRAAL - AMSTERDAM AMSTEL: SOME LINES DO NOT STOP AT AMSTERDAM CENTRAAL ANYMORE	99
FIGURE 7.12: MAGNIFICATION PRODUCTION (L) AND ATTRACTION AMSTERDAM CENTRAAL - AMSTERDAM AMSTEL FOR MORNING PEAK, WITH A CONCENTRATION OF ATTRACTION IN THE CITY CENTRE	100
FIGURE 7.13: MAGNIFICATION GEIN, WHERE THE METRO LINES DO NO REACH THE GEIN AREA ANYMORE.....	101
FIGURE 7.14: COMPARISON OF COSTS FINAL AND ADJUSTED FINAL SOLUTION	102
FIGURE 7.15: VEHICLE KILOMETRES BY MODE IN ORIGINAL SITUATION AND FINAL SOLUTION FOR ALL LINES.....	104
FIGURE 7.16: PASSENGER KILOMETRES BY MODE IN ORIGINAL SITUATION AND FINAL SOLUTION FOR ALL LINES.....	104
FIGURE 7.17: PASSENGER HOURS BY MODE IN ORIGINAL SITUATION AND FINAL SOLUTION FOR ALL LINES	104
FIGURE 7.18: MODAL SPLIT ORIGINAL SITUATION AND FINAL SOLUTION IN TERMS OF PASSENGER KILOMETRES AND PASSENGER HOURS.....	105
FIGURE 7.19: DIFFERENCE IN PASSENGER LINK LOADS BETWEEN THE FINAL SOLUTION AND THE ORIGINAL SITUATION	106
FIGURE 7.20: PASSENGER LOAD COMPARISON AMSTERDAM CENTRE	106
FIGURE 7.21: PASSENGER LOAD COMPARISON GEIN AREA	107
FIGURE 7.22: PASSENGER LOAD COMPARISON AMSTERDAM SOUTH WEST	107
FIGURE 7.23: COMPARISON OF PASSENGER FLOWS IN THE ORIGINAL SITUATION (LEFT) AND FINAL SOLUTION, WHICH SHOWS A SHIFT FROM TRAIN TO METRO.....	108

FIGURE 7.24: NUMBER OF TRANSFERS PER STATION IN ORIGINAL SITUATION AND FINAL SOLUTION. AMSTERDAM CENTRAAL AND SLOTERDIJK ARE MODELLED SEPARATELY FOR TRAIN AND METRO AND ARE SUMMED IN THE TABLES.	109
FIGURE 7.25: NUMBER OF TRANSFERS PER STATION IN ORIGINAL SITUATION AND FINAL SOLUTION FOR AMSTERDAM ZUID AREA.....	110
FIGURE 7.26: NUMBER OF TRANSFERS IN ORIGINAL SITUATION AND FINAL SOLUTION FOR AMSTERDAM CENTRAAL	110
FIGURE 7.27: DIFFERENCE WALKING IN CENTRE BETWEEN ORIGINAL SITUATION AND FINAL SOLUTION	111

List of tables

TABLE 2.1: THESIS IN SCIENTIFIC CONTEXT	11
TABLE 4.1: SYMBOLS	27
TABLE 4.2: DECISION VARIABLES	27
TABLE 4.3: PASSENGER COST PARAMETERS	28
TABLE 4.4: OPERATOR COST PARAMETERS.....	28
TABLE 6.1: OVERVIEW OF MODEL INPUT FOR NUMERICAL EXPERIMENTS	52
TABLE 7.1: FACTS AND FIGURES AMSTERDAM CASE STUDY	86
TABLE 7.2: INPUT AMSTERDAM CASE STUDY.....	93
TABLE 7.3: COMPARISON OF INDICATORS BETWEEN ORIGINAL SITUATION AND FINAL SOLUTION	103
TABLE 7.4: TOP 10 CHANGES IN NUMBER OF TRANSFERS	109
TABLE 8.1: LIMITATIONS	120

1 Introduction

Mobility has always been growing and is expected to keep growing (Schafer *et al.*, 2000). This growth will especially take place in urbanized areas. In the Netherlands, the train is expected to be the mode that will grow the most in terms of traveller kilometres (Centraal Planbureau and Planbureau Voor De Leefomgeving, 2015; Ministerie van Infrastructuur en Milieu, 2017). To accommodate this growth, there are two options: expansion of the rail infrastructure or more efficient use of the available infrastructure.

In rail transport, efficiency depends mostly on the different characteristics of the vehicles on the tracks. In a perfectly homogeneous system, like a metro system, headways can be fairly short, so a lot of vehicles can use the same track. The metro, however, is an exception. In most cases, different services (urban, regional, national and international) share the same tracks. As a result of their difference in function, the stopping patterns of the services differ. This corresponds to the travel distance the service is used for (see Figure 1.1). A larger travel distance usually means that the distance between consecutive stops is larger as well, increasing the operational speed of a line.




Figure 1.1: Services (adapted from (Govers, 2017))

However, there is no clear distinction between services when looking at travel distances. All types overlap with one or more other services. Metros and local trains, for instance, can both be efficient on distances ranging from five to twenty kilometres. Above such a distance, local trains compete with intercity trains up to about sixty kilometres.

Therefore, more research is needed to further improve the synergy between public transport services. Where that collaboration is now mostly limited to the synchronization of arrival and departure times at transfer stations, a renewed look at the division of functions could yield substantial benefits for both passengers and operators. These functions translate to stopping patterns and a optimizing this parameter has potential to further minimize costs.

A stopping pattern defines where vehicles of a certain line stop to allow passengers to alight and board. The distance between those stops is an important factor in the journey of a traveller (van Nes and Bovy, 2000): more stops per unit of length on a line allow the traveller to board closer to their origin and alight closer to their destination. However, the downside of this is that the vehicle has to stop more, which results in a longer travel time (see Figure 1.2). The figure also includes line spacing, which is not an explicit optimization variable in this thesis. However, it is taken into account implicitly, as the lines can be given a frequency of 0, effectively taking it out of order and increasing the distance between lines. Frequency also has an impact on passengers, for instance through the average waiting time and the available capacity. It is therefore also taken into account in this thesis as a decision variable.




Figure 1.2: Stop spacing, line spacing and frequency (van Nes, 2015)

1.1 Problem definition

Knowing that a more homogeneous service supply could lead to more capacity in the network, the question is raised what the role of each service should be in the network. The problem definition is:

Given a public transport passenger demand, a network with multiple services and a set of possible lines, what are the optimal stopping patterns and frequencies for each line, while satisfying capacity constraints?

In other words, where does which service stop, and how many vehicles should be dispatched? Are the different services worth the negative impact of heterogeneous operations on usage of infrastructure capacity? These services should support and supplement each other, instead of just coexist.

In order to answer this main question, several sub questions need to be answered:

1. How should the set of lines, for which the stopping patterns and frequencies will be optimised, be determined?
2. What services are needed to serve passengers optimally in a network?
3. Are existing, traditional services, with a clear distinction between long distance and short distance services and according stop spacing, optimal for passengers and operators, or is another stopping pattern more desirable?

4. What frequencies are needed for the different services on the lines to satisfy passenger demand?
5. Could the current infrastructure (2017) accommodate the desired stopping patterns and the required frequencies for a future passenger demand?

1.2 Research objectives

This research develops a method to find stopping patterns and frequencies for a predetermined part of a real public transport network, which allows for different stopping patterns. An optimal solution cannot be guaranteed, due to run time limitations, but the method is designed to use the available run time as efficiently as possible. The method takes the existing network, passenger demand, passenger costs and operator costs into account. Next to that, the model is able to assess the capacity of tracks and stations, to assess whether the needed services and frequencies fit within the current capacity of the network and it is able to take network capacity into account as a constraint.

With the results of the model in this thesis, it is possible to recommend a combination of services and frequencies for a network with a given demand and identify if and where the network does not have enough capacity to accommodate this.

1.3 Research scope

The method is on the strategic level of public transport design, as it takes possible future expansions of the network into account. This means that the level of detail of the results is limited to stopping patterns and frequencies of lines.

The infrastructure is treated as given, both the current network and possible expansions. Infrastructure is defined as the links (tracks) and nodes (stations) of the network, under the assumption that any train can run on any track. This means there is no distinction between metro and train infrastructure. The same goes for the capacity of the links and nodes in the network. The overall travel demand for public transport is assumed to be fixed and not influenced by the service level of the network, as are the origins and destinations of the travellers.

These limitations are chosen to keep the model comprehensible. Adapting stopping patterns and frequencies for multiple public transport lines and multiple services already requires many variables. Increasing the number of variables will add to the complexity, making the interpretation of the model more complex as well. Extending the research would also take more time than is available for this thesis.

The model is first applied to a fictive network, to demonstrate the functions of the model and to perform numerical experiments. The results of these experiments are used to look for trends for different scenarios within the same network.

Then the model is applied to the train and metro network in the region of Amsterdam for the year 2040. In this case study, the bus and tram networks are both treated as given: passengers can use them, but the networks are not optimised. Only public transport travel demand is taken into account.

1.4 Stakeholders and actors

There are several parties involved in the problem described above, who all have different interests and therefore different objectives. How to deal with these stakeholders and actors is not part of this thesis, but they are mentioned nonetheless as the objectives are part of the problem.

- **Passengers:** the passengers in a network are the reason the network exists in the first place. People want to travel and want their journey to be as fast as possible and as comfortable as possible. They are dependent on the network to offer them a fast and comfortable connection between their origin and their destination. Stopping patterns and frequencies of public transport services have a large impact on both the speed and the comfort experienced by passengers. Services that stop nearby both the origin and destination of a passenger reduce walking time for that passenger, but every stop in between only takes extra travel time. Reducing the number of stops increases the speed of the journey, while at the same time the probability of needing a transfer in that journey increases as well, decreasing comfort. A higher frequency is always positive for passengers: it decreases their waiting time and the level of crowding in a vehicle.
- **Public transport operators:** the operators provide the services that people can use to travel from their origin to their destination. Their objective is to make money, as they are often commercial companies. As fares and subsidies, which are the operators' sources of income, are not part of this thesis, their interest in this thesis is to minimize costs.
- **Public transport authorities:** the authorities' interests lie where operators and passengers meet. They are responsible for the rules that operators have to operate within, for instance on providing a minimum level of service to all people. At the same time, most operators require subsidy, as they cannot provide the required service in a way that is profitable. This subsidy is then paid by the public transport authority. The authorities' goal perhaps describes the goal of this paper best: provide the best public transport service for the lowest price.

1.5 Reading guide

The next chapter contains the background information on the underlying problems of stopping pattern and frequency optimization. A literature overview is provided as well, in which for the research components earlier papers are discussed. A comparison is made between the existing literature and this thesis, in which the scientific relevance of this thesis is highlighted. In chapter 3, the framework for the model is laid out, discussing the components of it. In chapter 4, the model is formulated mathematically, discussing the notation, cost functions and constraints. Chapter 5 discusses the implementation of the model in the programming and static passenger assignment software packages. The model is then applied to a fictive network for numerical experiments in chapter 6. The case study for the train and metro network of Amsterdam and its results are then described in chapter 7. Finally, conclusions are drawn, limitations are listed and recommendations for future research are provided in chapter 8.

2 Background and literature review

The objective of this thesis, optimizing stopping patterns and frequencies of public transport lines, can be seen as part of the Transit Network Design Problem (TNDP). The TNDP deals with designing public transport lines to provide passengers with routes to travel from their origin to their destination in a way that is as efficient as possible. It is more complicated than traditional transport network design, as not only needs to be determined what links are needed for an optimal network, but these links also need to be grouped in routes and for all those routes frequencies need to be determined (Desaulniers and Hickman, 2007).

First, the background makes clear where this thesis is positioned within the context of public transport design. Next, a literature review is conducted to establish the state of the art on network design and give an overview of the approaches taken by different researchers to design a public transport network.

2.1 Background on Transit Network Planning

The creation of a public transport network for an area is in literature often referred to as Transit Network Planning (Ibarra-Rojas *et al.*, 2015). TNP is a framework to design a public transport network from scratch in increasing detail up to driver rostering (Figure 2.1).

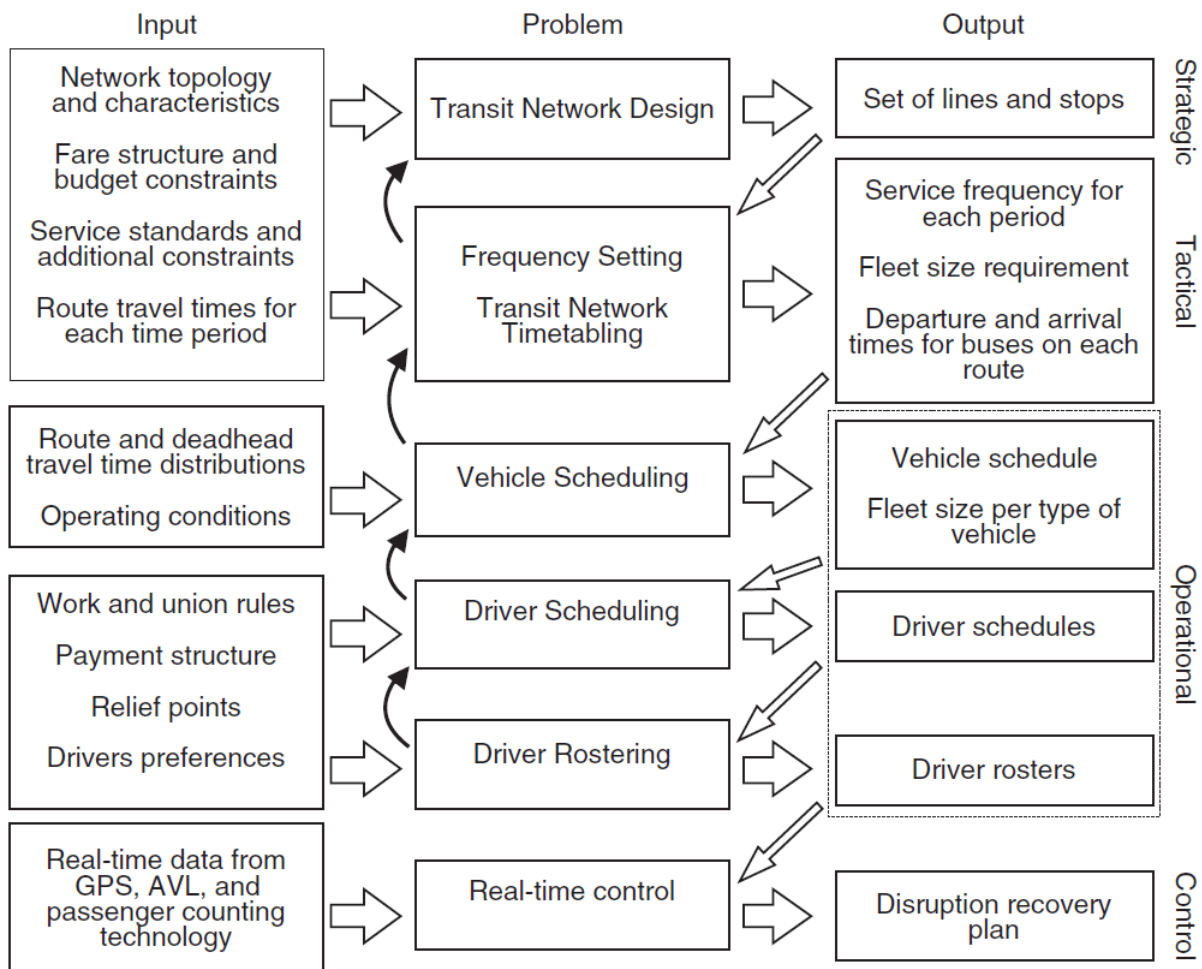


Figure 2.1: Transit planning process (Ibarra-Rojas *et al.*, 2015)

The TNP framework starts with the design of the network, so the location of lines and stops in an area. The choices that are made in this process depend on the topology of the area, the origins and destinations of passengers and the objective and constraints that are imposed. The objective is often to minimize costs, either for passengers, operators, or both. That is the reason a preliminary frequency needs to be set for each line, to determine for a line the capacity and waiting time (passenger costs) and the number of vehicles needed (operator costs). Constraints often include capacity of vehicles and infrastructure, fleet size, a budget, and passenger demand satisfaction.

On a tactical level, the frequencies of the lines are determined per time of day, as an off-peak demand does not need a peak frequency. These frequencies are then turned into a timetable, in which transfers between lines are timed to provide the best service on routes requiring more than one line.

The last steps are on an operational level and deal with the assignment of vehicles and drivers to the trips laid down in the timetable. On the operational level, mostly the costs of operators play a role, as passenger costs do not depend on the driver and negligible on the vehicle. Figure 2.1 also shows a fourth level, which is the control level. On this level, unforeseen disruptions in the network are dealt with ad hoc.

Another look at transit networks is through stop optimization, in which the number and locations of stops on a line are optimised for an objective. This differs from the TNDP in the fact that in order to optimise a stopping pattern, links and nodes should already be defined. For almost all practical issues in the Netherlands, this is indeed the case, as virtually every urban area is served by some kind of public transport. For bus networks, the routes can relatively easily be changed, as they are not bound by special infrastructure, except perhaps special bus lanes. This is not true for rail bound transport, which requires of course rail infrastructure to move. These tracks can be seen as the boundaries for the solutions space, as moving existing tracks or building new tracks is often too expensive. The train stations along the tracks are fixed as well, as moving them is far more expensive than moving a bus stop. All of this reduces the number of possible lines and stops.

2.2 Literature overview

The next section will present an overview of literature dealing with the different components of the TNDP. Topics are successively: decision variables, passenger assignment, constraints, assessment criteria, services, optimization methods, and applications. These topics are chosen because they are the important choices in model development: the decision variables determine what needs to be optimised; the passenger assignment is the way to assess a solution for the decision variables; the constraints determine the boundaries of the solution; the assessment criteria are used in the assessment of a solution; differentiating in services is a key feature of this thesis; the optimization method is a very important part of the model; the application is the practical use of the model.

In Table 2.1, an overview of the literature is provided and the place of this thesis within the context is shown for each topic described above. The sections following the table will explain this table further, including the choices made for this thesis.

Table 2.1: Thesis in scientific context

Paper	Decision variable(s)	Passenger assignment	Capacity constraints	Demand satisfaction constraints	Assessment criteria	Multiple services considered	Optimization method	Application to network
Ceder and Wilson, 1986	Route set, frequencies	None	Fleet size	100% satisfaction	Passenger and operator costs	No	No	Fictional network
van Nes, 2002	Stop spacing, line spacing, speed, frequency	None	None	Not applicable (no capacity constraints)	Total costs, social welfare	Yes, hierarchical	Analytical	Fictional networks
Fan and Machemehl, 2004	Route set, frequencies	Static	Vehicle capacity, fleet size	No, but penalty for unsatisfied demand	Passenger, operator and unsatisfied demand cost	No	Genetic Algorithm, Local Search, Simulated Annealing, Random Search, Tabu Search, Exhaustive Search	Fictional example network
Goossens <i>et al.</i> (2006)	Lines, stopping pattern	Static	Carriage capacity, number of carriages	100% satisfaction	Operational costs	Yes, no limit	Combinatorial optimization	Dutch train network
Borndörfer <i>et al.</i> (2008)	Route set, frequencies	Static	Link capacity	100% satisfaction	Passenger and operator costs	No	No	No
Lin and Ku, 2014	Stopping pattern	Board first arriving train	Train capacity	100% satisfaction	Profit company	Yes, no limit	Genetic Algorithm	Taiwan Railway Administration
Schmid, 2014	Route set, frequencies	Static	Fleet size, bus capacity per link	100% satisfaction	Ride and transfer time	No	Hybrid large neighbourhood search	Fictional network
Arbex <i>et al.</i> , 2015	Route set, frequencies	Static	Bus capacity	Only if within 2 transfers	Passenger and operator costs	No	Genetic Algorithm	Mandl's (1980) benchmark network

Paper	Decision variable(s)	Passenger assignment	Capacity constraints	Demand satisfaction constraints	Assessment criteria	Multiple services considered	Optimization method	Application to network
Yue <i>et al.</i> , 2016	Stopping pattern, timetable, stopping time	None	Station capacity	100% for each OD pair	“Profit”, decreases with stopping time and number of stops	Yes, three different types	Column Generation	Beijing – Shanghai High Speed Rail
López-Ramos <i>et al.</i> , 2016	Route set, frequency	Static	Infrastructure budget, fleet size, vehicle capacity, link capacity	100%, but not all stations need to be served	Passenger and operator costs	Yes, local and express	Lexicographic Goal Programming	Rapid transit Seville and Santiago de Chile
Gu <i>et al.</i> (2016)	Stop density, frequency, number of routes	Static	Vehicle capacity, minimum frequency	100%	Passenger and operator costs	Yes, all stop, skip-stop, express	Gradient descent method	Idealized line
This thesis	Stopping pattern, frequency	Static, zones	Node capacity, link capacity, vehicle capacity	100%, but not all stations need to be served	Passenger and operator costs	Yes	Genetic Algorithm	Fictional network & Amsterdam train and metro

2.2.1 Decision variables

The decision variables, i.e. the variables that are to be determined, do not differ very much per research, as the goal is virtually the same in each study. The goal for the TNDP is to come up with a route set and associated frequencies, which means the decision variables are route sets (sequences of links) and frequencies per route (Ceder and Wilson, 1986; van Nes, 2002; Fan and Machemehl, 2004; Schmidt, 2014; Arbex *et al.*, 2015; López-Ramos *et al.*, 2017).

For stopping pattern optimization problems, the obvious decision variable is the stopping pattern of a line (Lin and Ku, 2014). In other studies, the frequency was determined as well, to satisfy passenger demand (Goossens *et al.*, 2006). Yue *et al.* (2016) even included stopping time, as they constructed a timetable immediately.

In this thesis, stopping patterns and frequencies are both decision variables. Multiple lines will be optimised at the same time. For each line it is determined at which stations it will stop, as a choice from a predefined list, and how many times per hour the line will run. More details, such as scheduling, are not part of the optimization.

2.2.2 Passenger assignment

Most studies perform a static passenger assignment as part of their method, in order to be able to calculate the passenger costs. Few studies assign passengers to routes beforehand (Ceder and Wilson, 1986; López-Ramos *et al.*, 2017). A large part of the studies aggregates origins and destinations in the stations, inducing that all stations need to be served in order to satisfy passenger demand. Lin and Ku (2014) had all passengers board the first train to arrive after their own arrival at the station, provided it will serve their destination station as well.

This will not be the case here, as passengers will be allowed to choose from different stations to begin their public transport journey. The same applies to choosing an end station from which to walk to a destination. This allows for a more dynamic trade-off between operator costs and passenger costs, as not all stations have to be served in order to fulfil all demand.

2.2.3 Constraints

The constraints in each study can be split in two categories: capacity constraints and demand satisfaction constraints. The first category gives the limitations of the system, for instance fleet size (Ceder and Wilson, 1986; Schmidt, 2014), vehicle capacity (Fan and Machemehl, 2004; Lin and Ku, 2014), link capacity (López-Ramos *et al.*, 2017) and/or node capacity (Yue *et al.*, 2016). In this thesis, link and node capacity constraints can be enforced. Fleet size and vehicle capacity constraints are not taken into account, although crowding is. Therefore, vehicle capacity does play a role, but it does not limit passenger flows.

The second category determines whether passenger demand has to be fulfilled. Most studies impose that all passenger demand has to be satisfied, although Fan and Machemehl (2004) include a penalty for unsatisfied demand. Arbex *et al.* (2015) regard passenger demand as unsatisfied if more than two transfers are needed between the origin and the destination of the passenger.

Passenger demand satisfaction is no constraint, although unsatisfied demand will be penalized with the maximum passenger costs observed in the base case situation. This allows for demand to be disregarded if the extra costs exceed the extra benefits.

2.2.4 Assessment criteria

All studies evaluate the found solutions on operator costs, some form of passenger costs, or both. Only Yue *et al.* (2016) have a slightly different approach: a profit is defined per line, which decreases with each stop proportionally to the duration of that stop.

Operator costs depend on the number of vehicles that is needed and the length of the routes those vehicles have to travel. The latter can be expressed in either time (Mauttone and Urquhart, 2009), distance, or both (Cipriani *et al.*, 2012).

The most common indicator for passengers' costs is their travel time, split into weighted components. Those components are usually in-vehicle travel time and waiting time, often supplemented by an extra time penalty per transfer. As in most cases the origins and destinations are stations, access and egress time is not taken into account.

Studies looking into one aspect of public transport often include other costs that are studied or optimised. These extra costs are for instance in-vehicle time split into standing or seating, denied boarding (Cats *et al.*, 2016) and unsatisfied demand (Fan and Machemehl, 2004; Cipriani *et al.*, 2012).

This thesis will optimise combined passenger and operator costs, where passenger costs consist of in-vehicle travel time, waiting time, walking time, and transfer penalties, monetized using a value of time.

2.2.5 Services

Only a few studies take difference in services into account. Goossens, van Hoesel and Kroon (2006) and Lin and Ku (2014) allow for as many different services as needed, while Yue *et al.* (2016) define three different services and López-Ramos *et al.* (2016) define a local and an express service. Van Nes (2002) explicitly looked at the hierarchy between services and their relation for several decision variables.

Gu, Amini and Cassidy (2016) performed a study into stopping patterns for different transit modes, looking into skip-stop services and local-express services. For ordinary buses, bus rapid transit and rail, the services were designed and compared to an optimised all-stop alternative. It was found that when either the travel length or travel demand rises, a skip-stop service became favourable over an all-stop service for all examined modes. For very high travel lengths and travel demand, a local-express service became favourable for rail. However, the study assumed homogeneous demand over all stations. Fine-tuning the locations of for instance express service stations could provide even better results.

Although different services are considered in this thesis, they are not forced to follow a predefined pattern, such as express or local. Every service is allowed to stop at or skip any station.

2.2.6 Optimization method

The problem in this thesis, as well as other variants of the TNDP, is an NP-hard problem, as it has a very large solution space and several constraints (Nayeem *et al.*, 2014; Schmidt, 2014). NP-hard means that the time needed to solve a problem with n input variables is not limited to n^k for any k , i.e. the problem is not efficiently solvable by a computer.

To tackle this problem, heuristic methods are used. A heuristic cannot guarantee to find an optimal or perfect solution, but given enough time it will find a satisfying solution. There are several so called metaheuristics, heuristics that can be applied to a wide range of problems, and several of them will be shortly discussed here. The first five search methods (so up to exhaustive search) come from Fan and Machemehl (2004), for the others the sources are mentioned in the text.

- Local search: the local search method uses neighbourhoods. Around an initial solution, the neighbouring solutions are assessed and the best neighbour is used for the next search. If there are no better neighbours, a local optimum has been found. A global optimum cannot be guaranteed, especially in large solution spaces such as the one in this thesis. The method is therefore unsuitable.
- Simulated annealing: an improved local search method, which allows picking worse neighbours to escape local optima. A drawback is the long run time, which is a problem for this thesis.
- Random search: this algorithm simply searches randomly through the solution space until a pre-set maximum number of iterations has been reached. The complete lack of learning and lack of guarantee for finding a local optimum, let alone the global optimum, makes this method very unsuitable.
- Tabu search: this method is often used in combination with other heuristics and adds a memory, where earlier calculated solutions are stored. Moves that would go back to an earlier assessed solution are penalized or even forbidden, making the search more diverse and increasing the chance of escaping local optima. Although this is a good method, there is still a chance that a global optimum is not reached. It is therefore not the preferred method. However, the memory part is implemented, albeit in a different way (see section 5.3.7).
- Exhaustive search: this method simply compares all solutions, guaranteeing to find the best one. However, this also means that it takes the maximum computation time and it is therefore not suitable for this thesis.
- Combinatorial optimization: finding the optimum in a finite set of solutions, for instance the shortest route from A to B. it has been used by Goossens *et al.*(2006) to find the optimal combination of train routes. Combinatorial optimization requires the problem to be converted to mathematical equations only, which was not convenient for the problem in this thesis due to the required static passenger assignment.
- Hybrid large neighbourhood search: proposed by Schmidt (2014), a combination of large neighbourhood search (LNS) and linear programming (LP). The specific method used was a destroy-and-repair algorithm, where solutions were partially destroyed and recreated in the search for an optimal solution. This method is believed to be unsuitable as well, due to the lack of learning in the algorithm.
- Column generation: the method used by Yue *et al.* (2016) used the property that not all variables were expected to be relevant in the optimal solution. Column generation splits the problem in a master problem, the original problem with only a few variables, and several sub-problems, containing the other variables. This thesis takes all variables (i.e. costs) into account, so a method where some are ignored is not suitable in this case.

- Lexicographic goal programming: this type of programming, as used by López-Ramos *et al.* (2017), assumes a hierarchy in goals, where each goal is infinitely more important than the next one. As in this thesis such a hierarchy does not exist, this method is not suitable.
- Gradient descent: Gu *et al.* (2016) used that their solutions space was bounded, which made it possible to use the slopes of their solution space. The gradient descent method follows the steepest slope to find an optimal solution. Disadvantages are that it cannot cope with non-convex forms and also has trouble with integer variables. In the paper just mentioned this was solved by setting boundaries for the variables. However, in this thesis this is not preferred and the gradient descent method is therefore not chosen.
- Genetic algorithm: this method, based on biology, has been used by most researchers mentioned in Table 2.1. It is inspired by the evolution theory, where strong “parent” solutions mate to produce even stronger “children” solutions. It is an efficient technique to search for the global optimum in the solution space, as it combines the search guidance from the neighbourhood search with multiple options to escape local optima. It can also handle the integer variables needed in this thesis very well, which is why this method is chosen as the optimizing method. The genetic algorithm has had many applications in the past. It was for instance used to design a vision for the future Dutch railway network (Guis *et al.*, 2012; Keizer *et al.*, 2013) and for the optimization of the bus network of Utrecht (Van Eck, 2010). More information on the theory and implementation of the genetic algorithm can be found in section 3.4 and chapter 5.

2.2.7 Application

The model will be applied to both a fictional network and the existing public transport network of Amsterdam. The papers mentioned in Table 2.1 have only provided an application to either a fictional network or a (simplified) existing network. In this paper, the fictional network will be used to extensively show the different features of the model and the way it behaves in different scenarios.

Due to the very long run time of the model for the Amsterdam case, the model will only be run once for Amsterdam. However, it is very interesting to compare the outcome of the Amsterdam case to the different scenarios performed with the test network. The two different applications can provide a context for each other.

The downside is that the networks of both applications differ tremendously. The Amsterdam network is much bigger, more extensive, with more services and a larger difference in for instance travel distance. It will be important to keep these differences in mind when comparing both cases, as the different network might be a large factor in the way the model behaves in both networks.

2.3 Relevance

In the next sections, the scientific and practical relevance of this thesis are elaborated on.

2.3.1 Scientific contributions

The scientific contributions of this thesis are shown in Table 2.1, where its position is shown in scientific context. On eight criteria, several papers have been compared. It can be seen that none of the criteria are really new for this thesis, but that the combination of choices is:

- Decision variables: both stopping patterns and frequencies of multiple public transport lines are optimised simultaneously
- Passenger assignment: a static passenger assignment is performed, using zones to represent origins and destinations instead of stations
- Capacity constraints: node and link capacity are constraints that can be relaxed, vehicle capacity is included in a crowding factor for in-vehicle costs
- Demand satisfaction constraints: demand satisfaction is not required, a penalty is imposed on all unsatisfied demand
- Assessment criteria: the sum of passenger and operator costs is optimised
- Multiple services considered: yes, although no predefined pattern is enforced
- Optimization method: a genetic algorithm is used
- Application: the model is applied to both a fictional network and the train and metro network of the metropole region Amsterdam

The passenger assignment model of this thesis is more extensive than those used in previous studies and therefore shows a more realistic approximation of passenger behaviour. Most studies assume passenger demand to have its origin and destination fixed to a station. Demand is assumed to be fixed, similar to this thesis, which does not allow for stations to be skipped by all services if all demand should be served. This was already enforced by Fan and Machemehl (2008). The use of zones outside the stations and a fixed network connecting zones to more than one station, as in this thesis, does allow for stations to be skipped by all services, without immediately causing unsatisfied demand.

Another feature of this thesis, which differs from the mentioned papers, is the way different services are taken into account. In most papers, the services are defined up front, if they are taken into account at all. This is also caused by the mandatory stops at least one of the lines has to make at each station to serve demand, as discussed earlier. In this thesis, the services are completely free to be shaped from scratch within the allowed range of frequencies, including the possibility of line deactivation by setting the frequency to 0 runs per hour.

2.3.2 Practical contributions

The practical contribution of this thesis is the model, which can be used to analyse and optimise different public transport services in a network. The traditional hierarchy thinking is checked for its efficiency. This will be shown in the application of the model on the train and metro network of Amsterdam. The model can be deployed to use the available infrastructure as efficiently as possible and assess and optimise the added value for travellers of possible infrastructure extensions, while taking the costs for the operator into account as well. The model is

demonstrated in an application to both a fictional network and the existing public transport network of Amsterdam.

Givoni and Rietveld (2014) found that for just the train network of Amsterdam, overall welfare would decrease when even the least used train station would be closed. They concluded from this that it might be beneficial to open even more train stations. At the same time, they acknowledge that the local public transport network was not taken into account, nor was capacity or operator costs. This thesis will take this into account, providing a more complete analysis. Adding more stations is not considered in this thesis, but the model could be used to investigate the effects for both passengers and operators, by extending the network for which the model is run.

In the next chapter, a framework for the model developed in this thesis is presented. This framework will provide a general overview of the model, in which the different components of the model, such as the passenger assignment, are embedded. Other components that were discussed in this chapter, such as the constraints and the optimization method, will be discussed later.

2.4 Conclusion of the literature review

This thesis presents a new way to solve a part of the Transit Network Design Problem, by optimizing the stopping pattern and frequency of any number of lines within a public transport network. The static passenger assignment model that will be used is not new, but the degree of choice freedom exercised by the passengers within the network is. Where previous studies allowed only little choice of origin and destination station, if any at all, this thesis allows the use of a complete underlying public transport network.

At the same time, the capacity of the current infrastructure is taken into account. Potential solutions are discarded if they exceed infrastructure limitations at any point. This is discussed extensively in section 5.3.7. Both link and node capacity constraints can be enforced, although it is not mandatory. Vehicle capacity is not taken into account, although crowding is.

The fact that this thesis takes both passenger and operator costs into account, allows for a more comprehensive evaluation than a large part of the discussed papers. Both need to be in balance, even if this means that a part of the demand cannot be satisfied. Although unsatisfied demand is penalized, it is allowed, giving the model more options to find the optimal network.

An extra degree of freedom is that stopping patterns are not forced into a predetermined service, for instance express or local. Instead, each individual station can be skipped if this improves the overall costs.

The optimization method that will be used, the genetic algorithm, has proved its value several times before in a range of Transit Network Design Problems. When an exact solution cannot be found in a reasonable time, the algorithm is capable of finding a near optimal solution, making it very suitable for the problem of this thesis.

Finally, the model will be applied to both a fictional network and the existing network of Amsterdam. The numerical experiments with the fictional network allow for the assessment of multiple scenarios for the same network. This can give indications on whether different scenarios call for different public transport services and what these different services should be. Trends observed in the scenarios can be used to give insights in current situations, for instance when the number of passengers increases within a network.

This allows a good opportunity to compare theory to practice, as not only the networks can be compared, but the outcome for the Amsterdam case can be compared to the current situation as well. This gives context to the results of the model.

3 Modelling framework

A model-based optimization approach is chosen to answer the research questions. The reason for the modelling approach is that this way a simplified version of the situation (either fictional or existing) can be used to assess the impact of a solution. It also makes the optimization process automatic, allowing any number of steps per optimization round and any number of optimization iterations to be completed sequentially, without any manual actions in between. The model that will be developed will be called the Stopping Pattern and Frequency Optimization Model, or SPAFOM in short. The framework for it will be set up in the next section, the sections after that will discuss the several components of the framework. The last section shows the information flows to, from, and within the model.

3.1 Framework

The optimizable decision variables are the stopping pattern and frequency of a selection of public transport lines within a network, here called the pool of lines. A solution is defined as a combination of a stopping pattern and a frequency for each line in the pool of lines. After the solution has been checked for feasibility, the costs for the operator can be calculated directly.

A static passenger assignment model is used to run a simulation of the network, including the stopping pattern and frequency for the pool of lines. From this simulation, the passenger flows over the network can be obtained and used to calculate the passenger costs.

The operator and passenger costs will then be evaluated and improved. This is where the actual optimization takes place. With this new solution, the cycle starts again. This cycle is executed until a predetermined stop factor is reached.

This is in short the framework of the model. A visual representation of the framework is shown in Figure 3.1. In the next sections, the components of the framework will be explained conceptually one by one. For a more detailed explanation on the implementation of the conceptual model, the reader is referred to the next chapters.

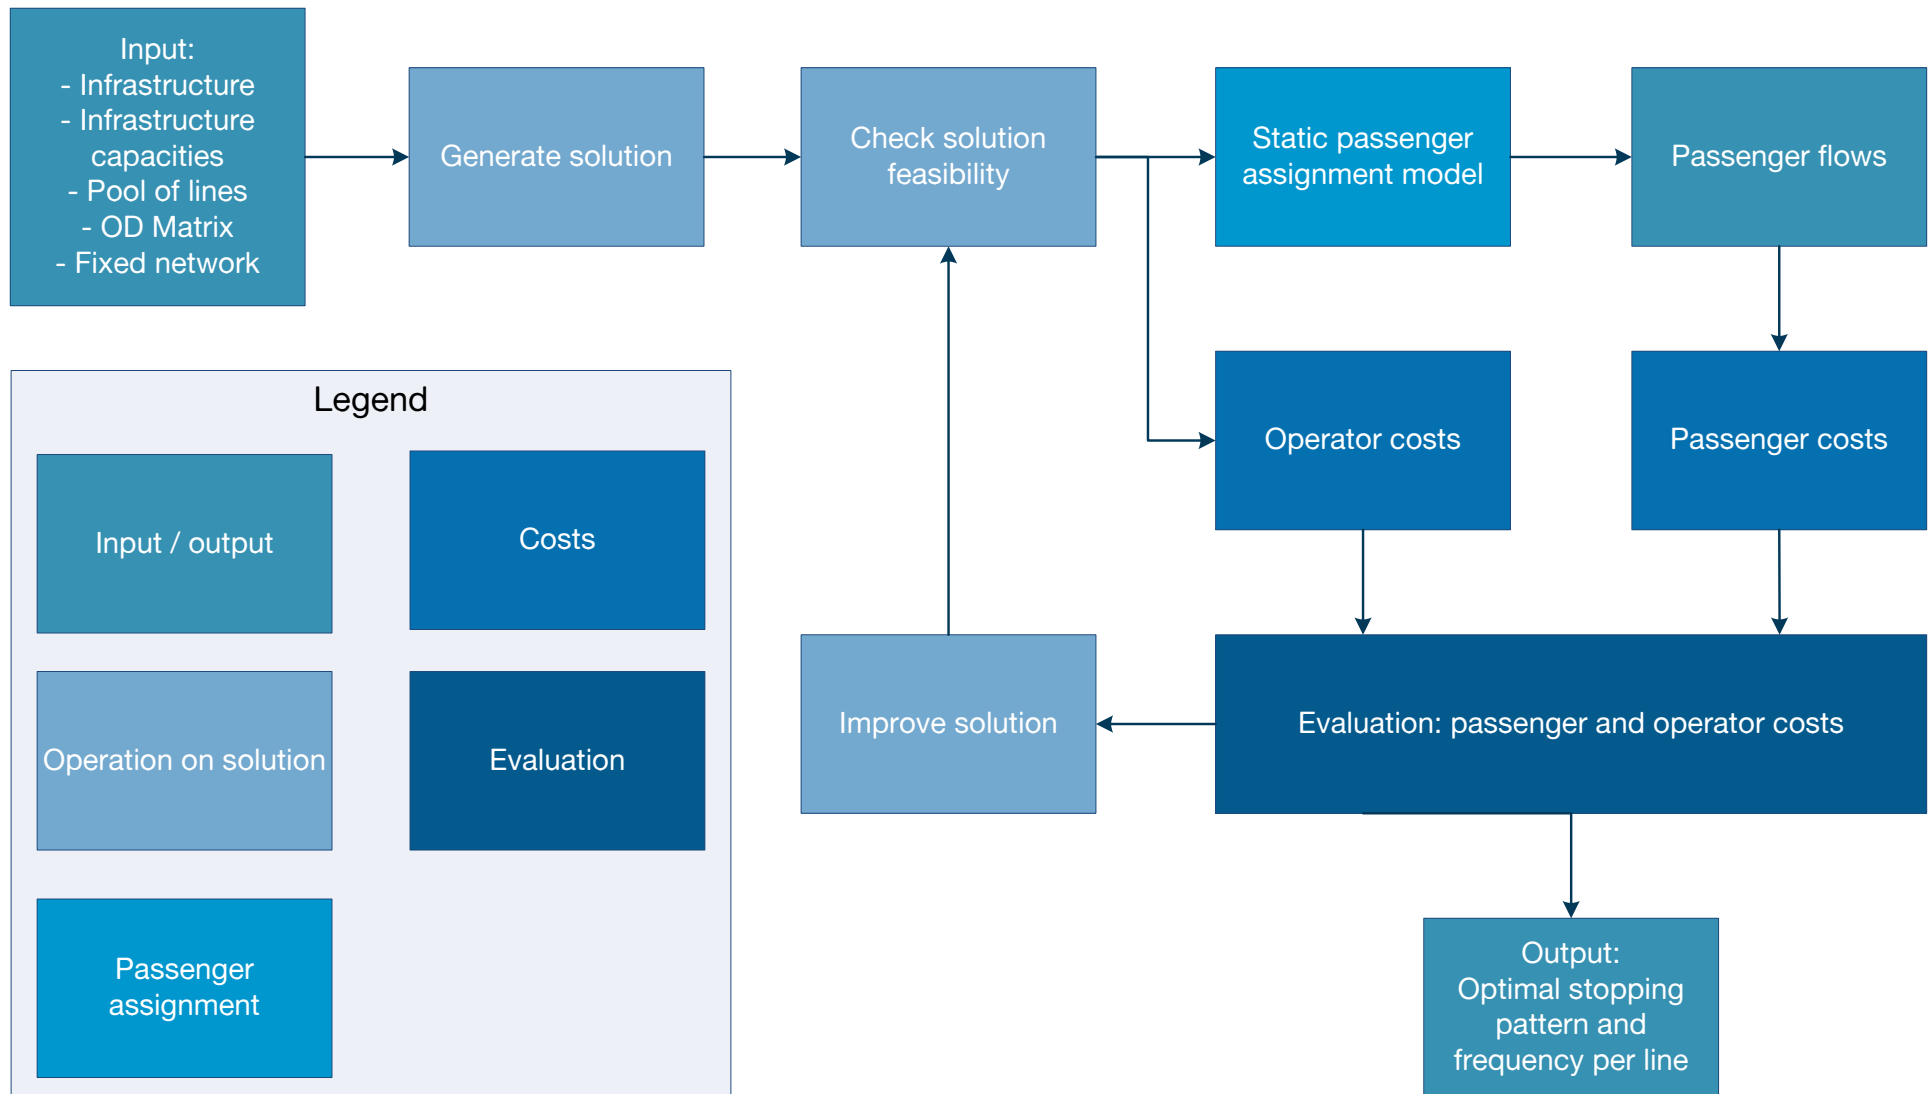


Figure 3.1: Stopping Pattern and Frequency Optimization Model framework.

3.2 Components of SPAFOM

The next sections will provide more detailed information on each component of SPAFOM.

3.2.1 Input

The model will require a network and a pool of lines to be optimised. The network is assumed to be fixed and represents the area within which the optimization will take place. It needs to cover at least the complete pool of lines, the relevant origins and destinations and the infrastructure between origins, destinations and public transport access points. Other types of infrastructure, for instance transit lines that are not part of the optimization problem, can be put in as well and will be used as underlying network. The infrastructure will be represented as links and nodes, as the model is on the strategic level and more details are not necessary. The capacity of the links and nodes is treated as given as well.

Another important input is the pool of lines, which will be constructed beforehand and are case specific. These are the lines that will be used to optimise the network. Each line will be represented as a sequence of stops, which are either served or skipped, and a frequency. The OD-matrix of passengers is an input as well. This will be an OD-matrix on zone level, so travellers are not fixed to a station.

The last input is a part of the underlying network, to accommodate travelling from zones to stations if they cannot be accessed directly. The relevant part of this network has to be determined beforehand and depends on for instance the importance of a line for the accessibility of a zone and the extent to which the line is an alternative for the examined network. This part of the network is treated as fixed, but the costs of passengers and operators using the network are taken into account.

3.2.2 Generate solutions

The first solutions, which are a combination of a stopping pattern and a frequency for each line, can either be generated randomly or be designed. The latter case is to make sure that the number of iterations needed to determine a sufficiently optimised solution is as low as possible, reducing the needed run time of the model. The initial solution can meet this quality by using for instance the current schedule of the lines, as it can be assumed that this is already quite good. An additional advantage is that the improvements that will be made by the model are the same improvements that should be made in real life. A random solution has the advantage of being unbiased, allowing for a more diverse solution exploration.

3.2.3 Check solution feasibility

The generated solution has to meet a couple of constraints, in order to be feasible. These constraints include infrastructure capacity for both lines and links, as well as symmetry in the lines. This means that a line has the same frequency back and forth and stops at the same stations both ways. It can then be checked what making the lines “logical” costs in terms of benefits for travellers and/or operators. In this thesis, however, only “logical” solutions are assessed, so the time needed to run the model will only be used for solutions that can be used in practice. Section 5.3.3 will discuss this elaborately.

3.2.4 Assignment model

The solution will then be assessed. To do this, a passenger assignment will be performed. The model will contain the input components mentioned earlier: available infrastructure, capacity of links and nodes, OD matrix, relevant underlying network, and a set of lines with stopping patterns and frequencies.

A journey in the model consists of the aspects as shown in Figure 3.2. It starts at an origin, from which a passenger can walk to any station the network allows him to. At the station, the passenger has to wait for the public transport vehicle to arrive and can then board that vehicle. The passenger then travels in that vehicle, until he reaches the station where he should alight. At this station, the passenger can transfer to another mode. The passenger then has to wait again for the vehicle to arrive and can then board that vehicle. He travels in that vehicle, until he reaches the right station. From that station, he can walk to his destination.




Figure 3.2: Schematic representation of a journey of a passenger

All the components impose time or another sort of discomfort on that passenger. As it is assumed to be fixed that a passenger travels from an origin to a destination, the passenger will look for the least inconvenient way to do this. This is what the passenger assignment model is for: finding the best route for each passenger and assigning that passenger to that route.

3.2.5 Passenger costs

The assignment model will have as output the passenger flows over the network for the solution of stopping patterns and frequencies. These data will be fed back to the optimization model, where the total passenger costs will be calculated. These costs consist of perceived travel time, waiting time, access and egress time, transfer penalties, and unsatisfied demand costs, all expressed in monetary values. How these costs are calculated exactly is explained in chapter 4.

Fares are not taken into account in the optimization model. The reason for this is that every euro a passenger spends on travelling will go to the public transport operator. As passenger and operator costs are treated equally important, the negative impact on passenger costs is balanced by the extra revenue for the operator. There is thus no influence on the total costs in the objective function, which aims to minimize both.

3.2.6 Operator costs

The costs for the operator are computed as well. Relevant costs depend on the number of trains needed, which depends in turn on the frequency and cycle time of a line. Operators want these costs to be as low as possible. This not only increases their profit, but also makes them more competitive in the open tenders for concessions. The ultimate goal for authorities is a self-sustaining public transport service, where the costs are covered by the revenues from the fares. Dutch Railways (NS) and RET in Rotterdam have shown this is possible (*RET Jaaroverzicht 2015 - Jaarverslag 2015, 2016, Financiën in het kort | NS 2016, 2017*). However, it is not the goal of this thesis to research how a subsidy-free public transport service can be established.

3.2.7 Evaluation

The new solution will be compared to the previous solution. Both passenger and operator costs will then be evaluated. The actors want their costs to be as low as possible and the public transport authorities want to minimize both costs as well. Therefore, the sum of both passenger and operator costs will be minimized. If the stopping factor has been reached, the solution will be provided as output in the form of an optimal stopping pattern and frequency per line. This solution will then be evaluated on case specific factors that were not in the model, such as operating costs compared to the current or a forecasted situation.

The model can also be used to assess whether infrastructure expansions are needed and if so, where they are needed most. Two approaches can be used to assess the need for extra infrastructure. The first method is enforcing the demand satisfaction constraint. This could cause passengers not reaching their destination. By imposing a penalty in terms of total costs for the system, the unsatisfied demand is then taken into account. The second method is relaxing the infrastructure capacity constraint. This means that every link and node in the network can handle an infinite amount of vehicles. Comparing the results of the model to the actual capacity will show where extra capacity is needed and how much.

3.2.8 Improve solution

The solution will then be improved using a genetic algorithm. As discussed in section 2.2.6, this method suits the problem best. A detailed description of this algorithm will be discussed in section 5.1.

3.3 SPAFOM information flow chart

In Figure 3.3, the information flow chart for SPAFOM is shown. This flow chart shows the framework discussed above and includes the input and output of the different steps. It also shows where decisions need to be made and what the consequences of these decisions are.

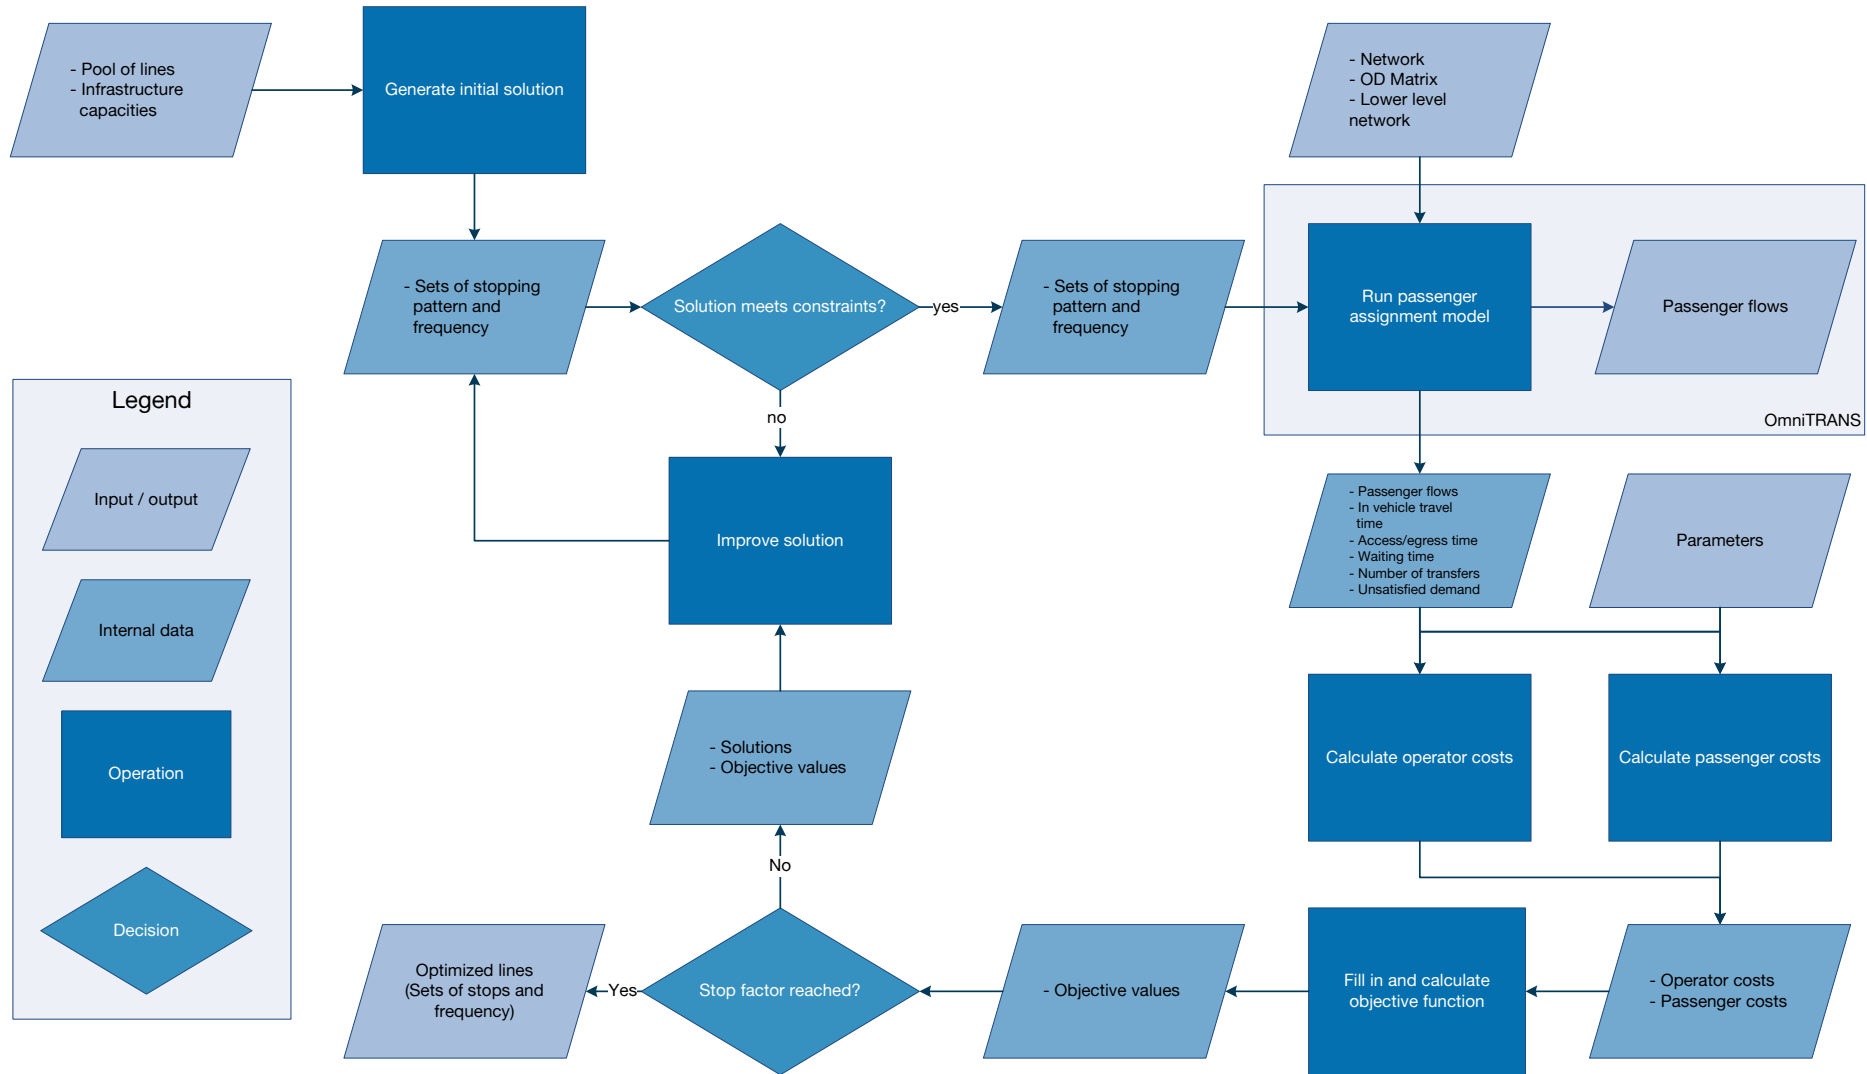


Figure 3.3: Stopping Pattern and Frequency Optimization Model information flow chart

3.4 Conclusion of the framework

This chapter has described the framework of the model and its components: the input, the generation of the first solutions, the feasibility check of the solutions, the assignment model, passenger and operator costs, the evaluation of those costs, and the improvement of the solutions. The information flows within the framework have been shown as well.

In the next chapter, the model will be formulated mathematically.

4 Model formulation

In the following chapter, the model described before will be described mathematically. The input, passenger costs and operator costs mentioned in the previous chapter will be given a mathematical form, as well as the constraints discussed in chapter 2. This mathematical formulation is needed to perform the application in the next chapter. The first section provides an overview of the symbols used in this chapter.

4.1 Notations

The tables below summarize the symbols and notations used throughout this chapter with a short explanation. Table 4.1 shows the notation of general symbols, Table 4.2 the notation of decision variables, Table 4.3 the notation of passenger cost parameters and Table 4.4 the notation of operator cost parameters.

Table 4.1: Symbols

Symbol	Units	Meaning
N	[-]	Set of nodes (OD)
N^S	[-]	Set of stations ($N^S \subseteq N$)
N^Z	[-]	Set of centroids ($N^Z \subseteq N$)
A	[-]	Set of links
A^T	[-]	Transit links ($A^T \subseteq A$)
A^W	[-]	Walking links ($A^W \subseteq A$)
L	[-]	Pool of lines
F	[-]	Set of frequencies
$t_{a,l}^{iv}$	[minutes]	In – vehicle travel time on link a ($a \in A^T$) for line l
$t_{z,i}^{walk}$	[minutes]	Walking time from centroid z to station i
$t_{i,l}^{dwell}$	[minutes]	Dwell time of line l at station i
χ_s	[-]	Number of tracks at station s
θ_s	[-]	Number of platforms at station s
κ_a	[-]	Number of tracks on link a ($a \in A^T$)
γ_l	[passengers]	Capacity of a vehicle on line l
$\lambda_{x,y}$	[passengers]	Demand between x and y ($x \in N^Z, y \in N^Z$)
$q_{l,a}$	[passengers]	Passenger occupancy of line l on link a
$\tau_{l,s}^{stop}$	[-]	Binary value, line l stops at station s (1) or not (0)
$w_{l,f}$	[-]	Binary value, line l has frequency f or not
ϕ_l^{max}	[-]	Maximum load factor of a vehicle on line l
ϕ_l	[minutes]	Minimum headway for line l
μ	[€]	Costs for unserved demand

Table 4.2: Decision variables

Decision variable	Units	Meaning
$f_l(w_{l,1}, \dots, w_{l,f})$	[veh / hr]	Frequency of line l
$s_l(\tau_{l,1}, \dots, \tau_{l,n})$	[yes or no]	Stopping pattern of line s

Table 4.3: Passenger cost parameters

Passenger costs	
β_l^{iv}	Costs per in vehicle travel minute
β_l^{cr}	Costs for crowding
β_l^{wait}	Costs per waiting minute
β_l^{walk}	Costs per walking minute
β^{tr}	Transfer penalty
β_{unsat}	Unsatisfied demand penalty ($Costs_{pax}^{max}$)

Table 4.4: Operator cost parameters

Operator Costs	
β_l^{disp}	Costs per dispatched train
β_l^{op}	Costs per train operating hour

4.2 Cost functions

The objective function consists of costs for (1) passengers and (2) operators. The combination of both should be minimized, in order to fulfil the objective. Each of the costs depends on the input of the model and a number of variables that follow from the solution or the assignment model, which will be explained in the following sections.

4.2.1 Input

The input of the model is a network, consisting of a set of nodes N and a set of links A . Nodes can be either stations (N^S) or centroids (N^Z), links can be either transit links (A^T) or walking links (A^W). Furthermore, a pool of lines L , containing all public transport lines, and a set of possible frequencies F are defined.

4.2.2 Passenger costs

The passenger costs are calculated using data from the static passenger assignment model. They consist of travel time and a penalty for transfers. The travel time can be split into access and egress time, waiting time, and in-vehicle time.

4.2.2.1 Access time

Access time is the time a traveller needs for the first part of the journey, which is walking (other access or egress modes are not taken into account) from his origin centroid z to the preferred boarding station i . This centroid represents a number of addresses aggregated for an area. The total time spent walking between a centroid and a station depends on the number of passengers walking between them ($q_{z,i}$) and the time a passenger needs to walk the distance ($t_{z,i}^{walk}$) for every combination of centroid and stop. The resulting time is multiplied by a factor β^{walk} , which represents the discomfort that people perceive for walking relative to travelling on-board a vehicle. The above is formulated as Equation 1.

$$\sum_{z \in N^Z} \sum_{i \in N^S} \beta^{walk} \cdot t_{z,i}^{walk} \cdot q_{z,i} \quad \text{Equation 1}$$

4.2.2.2 Waiting time

For the calculation of the waiting time, a uniform arrival time of passengers is assumed, both at their first station and for any transfer. This means that every passenger waits on average half the headway time of a line for each line he boards. As the model works with frequencies, the average waiting time is calculated using the frequency of each line. The headway of a line in minutes is calculated as sixty minutes divided by the frequency. The average waiting time is half that, so the waiting time can be expressed as $\frac{30}{f_l}$. Multiplied by the total number of passengers using this line (q_l), the total passenger waiting time is obtained.

Passengers can time their arrival at a station to a certain degree. Under the assumption that people will never wait more than half an hour for a service, a maximum waiting time of 30 minutes is assumed. This limits the maximum average waiting time to 15 minutes.

People prefer sitting on-board a vehicle over waiting, therefore an extra factor β_l^{wait} is applied. Note that the factor may differ per line, as for instance an average bus stop is less comfortable to wait at than an average train or (underground) metro station. A summation over every line gives Equation 2.

$$\sum_{l \in L} \beta_l^{wait} \cdot \min\left(\frac{30}{f_l}, 15\right) \cdot q_l \quad \text{Equation 2}$$

4.2.2.3 In-vehicle time

The total passenger-in-vehicle time is calculated for each link and each line as the travel time for line l over transit link a ($t_{a,l}^{iv}$), multiplied by the number of passengers on-board of that line for that link ($q_{a,l}$). A factor β_l^{iv} is applied to provide the possibility to differentiate in terms of comfort level between different lines. If the comfort level of each line is assumed to be the same, so there is no preference for a mode, this factor is no longer needed. This is due to the fact that all other time factors are scaled relative to in-vehicle time, so $\beta^{iv} = 1$. With this factor included, in-vehicle travel time is calculated as in Equation 3. The second part (Equation 4) is a penalty for discomfort because of crowding and depends on the volume over capacity ratio (V/C-ratio). If the V/C ratio of a line on a certain link is 0.5 or less, so no more than half the vehicle's capacity is used, no discomfort penalty is used. If the vehicle's occupancy is twice its comfortable capacity, the penalty is twice the in-vehicle time. Between the V/C-ratios of 0.5 and 2, a linear function is assumed. This is a simplification of the crowding factor used in Cats and Hartl (2013), so it could be implemented more easily. The effects on the results are negligible, as the original and simplified functions do not differ much.

$$\sum_{l \in L} \sum_{a \in A^T} \beta_l^{iv} \cdot (t_{a,l}^{iv} + t_{a,l}^{perceived}) \cdot q_{a,l} \quad \text{Equation 3}$$

$$t_{a,l}^{perceived} = \begin{cases} 0 & \text{if } \frac{q_{a,l}}{\gamma_l} \leq 0.5 \\ \left(\frac{q_{a,l}}{\gamma_l} - 0.5\right) \cdot \frac{2}{1.5} \cdot q_{a,l} \cdot t_{a,l}^{iv} & \text{if } 0.5 < \frac{q_{a,l}}{\gamma_l} < 2 \\ 2 \cdot q_{a,l} \cdot t_{a,l}^{iv} & \text{if } \frac{q_{a,l}}{\gamma_l} \geq 2 \end{cases} \quad \text{Equation 4}$$

4.2.2.4 Egress time

Egress time is basically the same as access time, only the direction is different. It is the last part of a passenger's journey: the walk from the stop he alighted to his destination centroid. The equation is therefore the same as the one for the access time, with a change of directions, as can be seen in Equation 5.

$$\sum_{j \in N^S} \sum_{z \in N^Z} \beta^{walk} \cdot t_{j,z}^{walk} \cdot q_{j,z} \quad \text{Equation 5}$$

4.2.2.5 Transfer penalty

A separate penalty is applied for each transfer within the network. This penalty is expressed in extra perceived travel time for every transferring passenger (q_l^{tr}) per transfer and represents the traveller's discomfort of the transfer. This extra time is represented by β_l^{tr} . Note that only the line a passenger transfers to influences the magnitude of the transfer penalty, as this determines the circumstances of the out-of-vehicle time during a transfer for the largest part. Transferring to or from a walking trip is not considered as a transfer. The transfer penalty is calculated as in Equation 6. The extra waiting time involved in a transfer is incorporated in Equation 2.

$$\sum_{l \in L} \beta_l^{tr} \cdot q_l^{tr} \quad \text{Equation 6}$$

4.2.2.6 Unsatisfied demand penalty

Any demand left unsatisfied by the network, so the demand not able to reach its destination, is penalized as well. This unsatisfied demand ($\lambda_{x,y}^{unsat}$) is penalized with the maximum passenger costs $Costs_{pax}^{max}$ in the original network, so the network before the optimization, and is denoted as β_{unsat} . This leads to Equation 7. The use of a penalty for unsatisfied demand obviates a constraint on servicing each station. It gives the model the flexibility to deny a (small) share of the passengers to benefit the overall costs. The relatively high costs make sure that this decision cannot be taken lightly.

$$\sum_{x \in N^Z} \sum_{y \in N^Z} \beta_{max} \cdot \lambda_{x,y}^{unsat} \quad \text{Equation 7}$$

4.2.3 Operator costs

The operator costs can be calculated directly from the solution of stopping patterns and frequencies. They can be specified a priori, as they do not depend on the output of the transit assignment model. A distinction is made between fixed costs per vehicle and variable costs per vehicle per hour or per kilometre.

4.2.3.1 Fixed costs

The fixed costs are the costs associated with using one additional vehicle, for instance a train or bus. This depends on the frequency and cycle time of a line. The cycle time depends on the time a vehicle needs to traverse all links of its line ($t_{a,l}^{iv}$), the total time lost stopping at stations ($\tau^{stop} \cdot t_{l,s}^{st}$), and the layover time at both ends of the line ($t_l^{layover}$). The cycle time is divided by sixty to convert to hours and multiplied by the frequency of the line, which results in the total vehicles needed for a line. This number is rounded up to an integer, as only whole vehicles can be used. Multiplying with the fixed costs per vehicle gives the total fixed costs (see Equation 8).

$$\sum_{l \in L} \sum_{a \in A^T} \sum_{s \in N^S} \beta_l^{fix} \cdot \left\lceil \frac{(t_{a,l}^{iv} + \tau^{stop} t_{l,s}^{st} + t_l^{layover})}{60} \cdot f_l \right\rceil \quad \text{Equation 8}$$

4.2.3.2 Variable costs

The variable costs are calculated much the same as the fixed costs. The only difference is that, instead of number of vehicles, the amount of time (Equation 9) or distance (Equation 10), depending on the type of variable costs, is multiplied by the costs per time or distance unit of operation per vehicle.

$$\sum_{l \in L} \sum_{a \in A^T} \sum_{s \in N^S} \beta_l^{var,time} \cdot (t_{a,l}^{iv} + \tau^{stop} t_{l,s}^{st} + t_l^{layover}) \cdot f_l \quad \text{Equation 9}$$

$$\sum_{l \in L} \sum_{a \in A^T} \sum_{s \in N^S} \beta_l^{var,dist} \cdot t_{a,l}^{iv} \cdot v_l \cdot f_l \quad \text{Equation 10}$$

4.2.4 Objective function

Summing up Equation 1 to Equation 10 and minimising the total equation leads to the combined objective function as follows in Equation 11:

$$\begin{aligned}
\min \bigg\{ & \sum_{z \in N^Z} \sum_{i \in N^S} \beta^{walk} \cdot t_{z,i}^{walk} \cdot q_{z,i} + \sum_{l \in L} \beta_l^{wait} \cdot \frac{30}{f_l} \cdot q_l \\
& + \sum_{l \in L} \sum_{a \in A^T} \beta_l^{iv} \cdot t_{a,l}^{iv} \cdot q_{a,l} + \sum_{l \in L} \beta_l^{tr} \cdot q_l^{tr} \\
& + \sum_{j \in N^S} \sum_{z \in N^Z} \beta^{walk} \cdot t_{j,z}^{walk} \cdot q_{j,z} \\
& + \left(\sum_{x \in N^Z} \sum_{y \in N^Z} \lambda_{x,y}^{unsat} + \sum_{l \in L} \max(0, f_l \cdot \gamma_l - q_l) \right) \cdot \beta_{max} \\
& + \sum_{l \in L} \sum_{a \in A^T} \sum_{s \in N^S} \beta_l^{fix} \cdot \left\lceil \frac{(t_{a,l}^{iv} + \tau^{stop} t_{l,s}^{st} + t_l^{layover})}{60} \cdot f_l \right\rceil \\
& + \sum_{l \in L} \sum_{a \in A^T} \sum_{s \in N^S} \beta_l^{var,time} \cdot (t_{a,l}^{iv} + \tau^{stop} t_{l,s}^{st} + t_l^{layover}) \cdot f_l \\
& + \sum_{l \in L} \sum_{a \in A^T} \sum_{s \in N^S} \beta_l^{var,dist} \cdot t_{a,l}^{iv} \cdot v_l \cdot f_l \bigg\}
\end{aligned}$$

Equation 11

The objective function has to meet a couple of constraints, which will be explained in the following section.

4.3 Constraints

The solution has to meet a couple of constraints in order to be feasible. These constraints include line constraints, such as line symmetry, and capacity constraints. These constraints will later be used to limit the solution space.

4.3.1 Stopping pattern and frequency constraints

Each line has to meet some constraints concerning its stopping pattern and frequency, as a line can only run one way. The way back is served by a different line. Both lines need to be a mirrored image of the other, so the same stopping pattern and frequency for each pair of lines back and forth. If a passenger travels from one station to the next, he can always take the same line back and that line has the same frequency as the first line. The reason for this is twofold: first, it leads to a lot of uncertainty for passengers if they cannot take the same line back to their origin as they took to travel to their destination. Second, some vehicles can be used twice or more within an hour if their cycle time is half an hour or less. However, this requires that sufficient vehicles are returned to their starting station, so both the stopping pattern and frequency of a line need to be the same back and forth.

The constraints are enforced by Equation 12 and Equation 13.

$$s_{l,i,j} = s_{l,j,i} \quad \forall l \in L, i \in N^S, j \in N^S \quad \text{Equation 12}$$

$$f_{l,i,j} = f_{l,j,i} \quad \forall l \in L, i \in N^S, j \in N^S \quad \text{Equation 13}$$

Another constraint for the frequency of each line is that there can only be one frequency: lines cannot run both 2 and 4 times per hour. This is enforced by Equation 14, which assures that only one binary value concerning the choice of frequency is 1.

$$\sum_{f \in F} w_{l,f} = 1 \quad \forall l \in L \quad \text{Equation 14}$$

4.3.2 Node capacity

The node capacity, which is the capacity of a station, depends on the number of tracks at the station, both with and without platforms. Obviously, a track is only suitable for stopping if it has a platform where passengers can board and alight. The capacity of a node is therefore twofold: all vehicles need to be able to pass a node, similar to the capacity of the link, and there needs to be enough “platform time” for all stopping trains.

Equation 15 deals with the first: the total time vehicles need to pass the station depends on the frequency of lines and how many of those stop. If a line stops at a station ($\tau_{l,n}^{stop} = 1$), the dwell time t_l^{dwell} , assumed to be the same for every station for a line, and minimum headways between two consecutive services (ϕ_l) of the line are summed and multiplied by the frequency of the line. This is the total time a line stopping at a station needs at that station. For lines that do not stop at the station, the needed time is only the minimum headway between services multiplied by the frequency. The time of both the stopping and passing vehicles cannot exceed the available time, similar to link capacity.

The availability of platforms is checked in Equation 16. This equation is largely similar to the previous one: the time needed by stopping trains is checked against the availability of tracks with a platform.

There are a couple of assumptions that are important to mention. First, it is assumed that all vehicles enter the station on the same side and exit on the other side; none of them changes direction. This is only true for stations that are not a terminal for a line. Assuming the stations and lines are symmetric in both directions (the latter is a constraint), the assumption of uniformity in direction holds for these stations.

Second, it is assumed that a platform track can only accommodate one stopping vehicle at the time. In reality, this is sometimes not the case. However, as this assumption always results in an underestimation of the capacity of a station, it is not deemed to be important enough to specify in detail, as such details are on the operational planning level and too detailed compared to other components of the model.

$$\sum_{l \in L} \tau_{l,s}^{stop} \cdot (t_l^{dwell} + \phi_l) \cdot f_l + (1 - \tau_{l,s}^{stop}) \cdot \phi_l \cdot f_l \leq 60 \cdot \chi_s \quad \forall s \in N^S \quad \text{Equation 15}$$

$$\sum_{l \in L} \tau_{l,s}^{stop} \cdot (t_l^{dwell} + \phi_l) \cdot f_l \leq 60 \cdot \theta_s \quad \forall s \in N^S \quad \text{Equation 16}$$

4.3.3 Link capacity

The capacity of a link depends on the number of tracks on the link (κ_a), the minimum headway of a line (ϕ_l) and the frequencies (f_l). This constraint is expressed as in Equation 17, where the total time needed for all services to pass a point on a link cannot exceed the available time of sixty minutes per track.

$$\sum_{l \in L} \phi_l \cdot f_l \leq 60 \cdot \kappa_a \quad \forall a \in A^T \quad \text{Equation 17}$$

4.3.4 Demand satisfaction

Altering stopping patterns may cause certain centroids to become disconnected. The demand from and to this centroid will then not be served, leading to unsatisfied demand. It will not be enforced that all demand needs to be served. Instead, each passenger that cannot reach its destination will get the highest generalised travel cost that was observed in the original network (μ), which will be added to the passenger costs. See section 4.2.2.6 for this.

4.4 Conclusion of the model formulation

The previous sections have provided mathematical formulations of the input of the model and the calculation of all passenger and operator costs. These formulations have been combined into the objective function. The constraints this objective function is subject to have been formulated as well. In the next chapter, these formulations are put into practice in the implementation.

5 Model implementation

With the framework set up in chapter 3 and theoretical basis laid out in the chapter 4, it is now time to put the theory into practice. The model is implemented using the programming language Ruby (*Ruby Programming Language*, 2017) and the passenger assignment software OmniTRANS (*OmniTRANS - DAT.Mobility*, 2017). Ruby version 1.8.7 is integrated in OmniTRANS and therefore chosen over other languages.

Ruby can be used to directly access the databases of OmniTRANS, which hold all information of the network in the passenger assignment model, such as links, nodes, passenger travel demand, and public transport lines. Special commands are available to read and write the databases. The steps in the framework that was introduced earlier are used to structure the next sections. However, short introductions to genetic algorithms and OmniTRANS databases are provided first, as this knowledge helps understand how the implementation takes place.

After these introductions, the input of the model is discussed. The cycle of the genetic algorithm is then explained: generating solutions, checking the feasibility of each solution, implementing the solution in OmniTRANS, calculating passenger and operator costs, evaluating these costs, and improving the solution. The cycle is then completed, as a feasibility check is then performed. Finally, the output of the model is discussed.

5.1 Introduction to genetic algorithms

The optimization technique used in the implementation is a genetic one. This is an evolutionary algorithm, based on the evolution theory in biology (Mitchell, 1996). There are several steps, as can be seen in Figure 5.1.

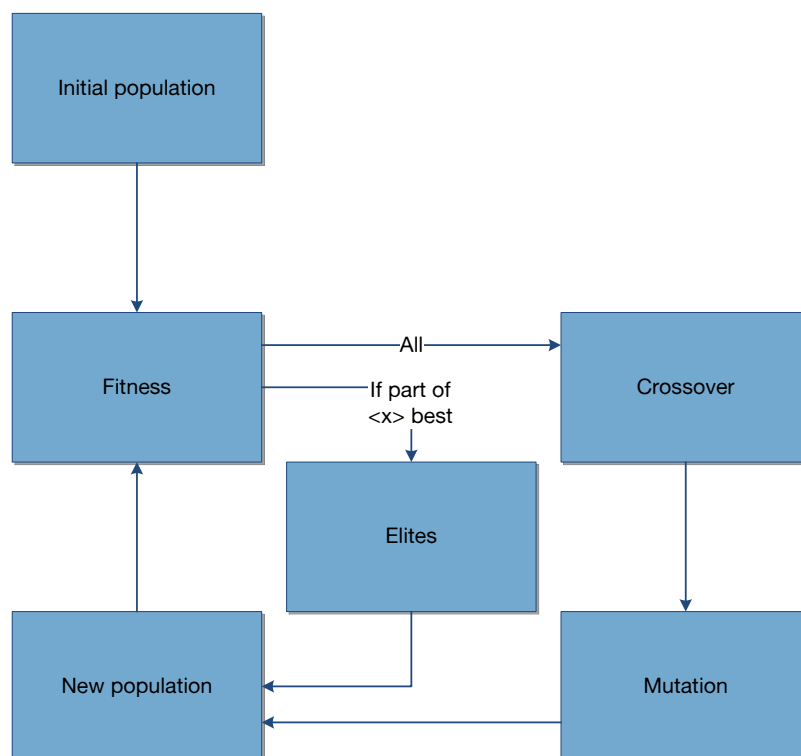


Figure 5.1: Genetic algorithm

One starts with an initial population, which is a collection of possible solutions for the decision variables. The decision variables of this model are stopping patterns and frequencies. These need to be represented in such a way that they can be optimised by the genetic algorithm. A binary approach is chosen for this, meaning that the solutions are strings of 0's and 1's. A 0 or 1 is called a gene. The genes together form a chromosome, which is one solution for all lines that need to be optimised. Multiple solutions together form a generation. Multiple solutions together form a generation.



Figure 5.2: Generations, chromosomes and genes

The binary approach is an easy way to represent stopping patterns, as will be shown hereafter. For the frequency of a line, this approach is not intuitive. However, the main advantage is that each solution can be represented as one binary string. Therefore, this approach was chosen, together with a repair mechanism that will be explained in section 5.3.3.

The first part of the solution represents the stopping patterns. Each binary value represents a station on a line and whether a vehicle on that line stops (1) or does not stop (0) at that station. An example is given in Figure 5.3, where a line with five stops is represented by five binary values. In this case, the line would stop at stations 1, 3 and 4, while going through stop 2. The last link is not traversed and station 5 is never reached.




Figure 5.3: Solution representation

The second part of a solution represents the frequency per line. For each line there is a binary string with the same number of binary values as there are possible frequencies. Only one of those values is 1, as a line can only have one frequency. An example is given in Figure 5.4, where the result is a frequency of 4 runs per hour. Each line has its own frequency representation in the solution, so each line can have a different frequency.

0	1	2	3	4	5	6	8	10	12	15	20	30	60
---	---	---	---	---	---	---	---	----	----	----	----	----	----

..., 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, ...

Figure 5.4: Frequency solution representation

Each solution is then tested for its fitness. This fitness is the result of the objective function described in section 4.2 by Equation 11. A predetermined amount of solutions with the best fitness scores are preserved for the next generation, which is called elitism. This is the first way information can be passed on to the next generation and it is represented by the short loop in Figure 5.1. Another way is through crossover: based on the fitness of all solutions, including the elite solutions, a crossover is performed between two parents. The better the fitness, the higher the chance a solution has of “reproducing”. This is the long loop in Figure 5.1.

When two parents are chosen, two children are produced with a random combination of characteristics of both parents. Each child then has a chance of a random change, similar to mutation. The children, which are either mutated or not, are then added to the new population with the elites until the population has the original size. This population is the next generation.

The fitness of this population is then determined and from this point the cycle starts again. The cycle can be ended either by defining a maximum number of generations or by checking the change between the fitness scores of the best solutions of the last couple of consecutive generations and determining a convergence factor. For more theoretical information on genetic algorithms and evolutionary computation, the reader is referred to for instance Mitchell (1996) and Garis (2004).

5.2 Introduction to OmniTRANS databases

The databases of OmniTRANS are divided into tables. Each table contains information for one object type, e.g. a stop, transit line, link, node, etc. Most information is grouped per PMTURI, which stands for Purpose, Mode, Time, User, Result and Iteration. The first four are used to specify input (e.g. a traveller has Purpose work, goes by train at 10 AM and is a commuter) and the latter two are used to specify output of the model (e.g. a result for an All or Nothing assignment, the first iteration). The tables are grouped accordingly. Some information only depends on the object type, other information differs per Mode and Time, and again other information differs per Purpose, Mode, Time and User group. For this thesis, it is sufficient to know that information is stored in such tables.

5.3 SPAFOM implementation

In the remaining sections of this chapter, the steps of the model are explained. These steps are visualized in Figure 5.5. The sections follow the model flow chart, discussing the input of the model (5.3.1), the generation of the first solutions (5.3.2), the feasibility check (5.3.3), the implementation of the solutions in OmniTRANS (5.3.4), the different costs (5.3.5 and 5.3.6), the evaluation of those costs (5.3.7), the way the solutions are improved (5.3.8), and the output of the model (5.3.9).

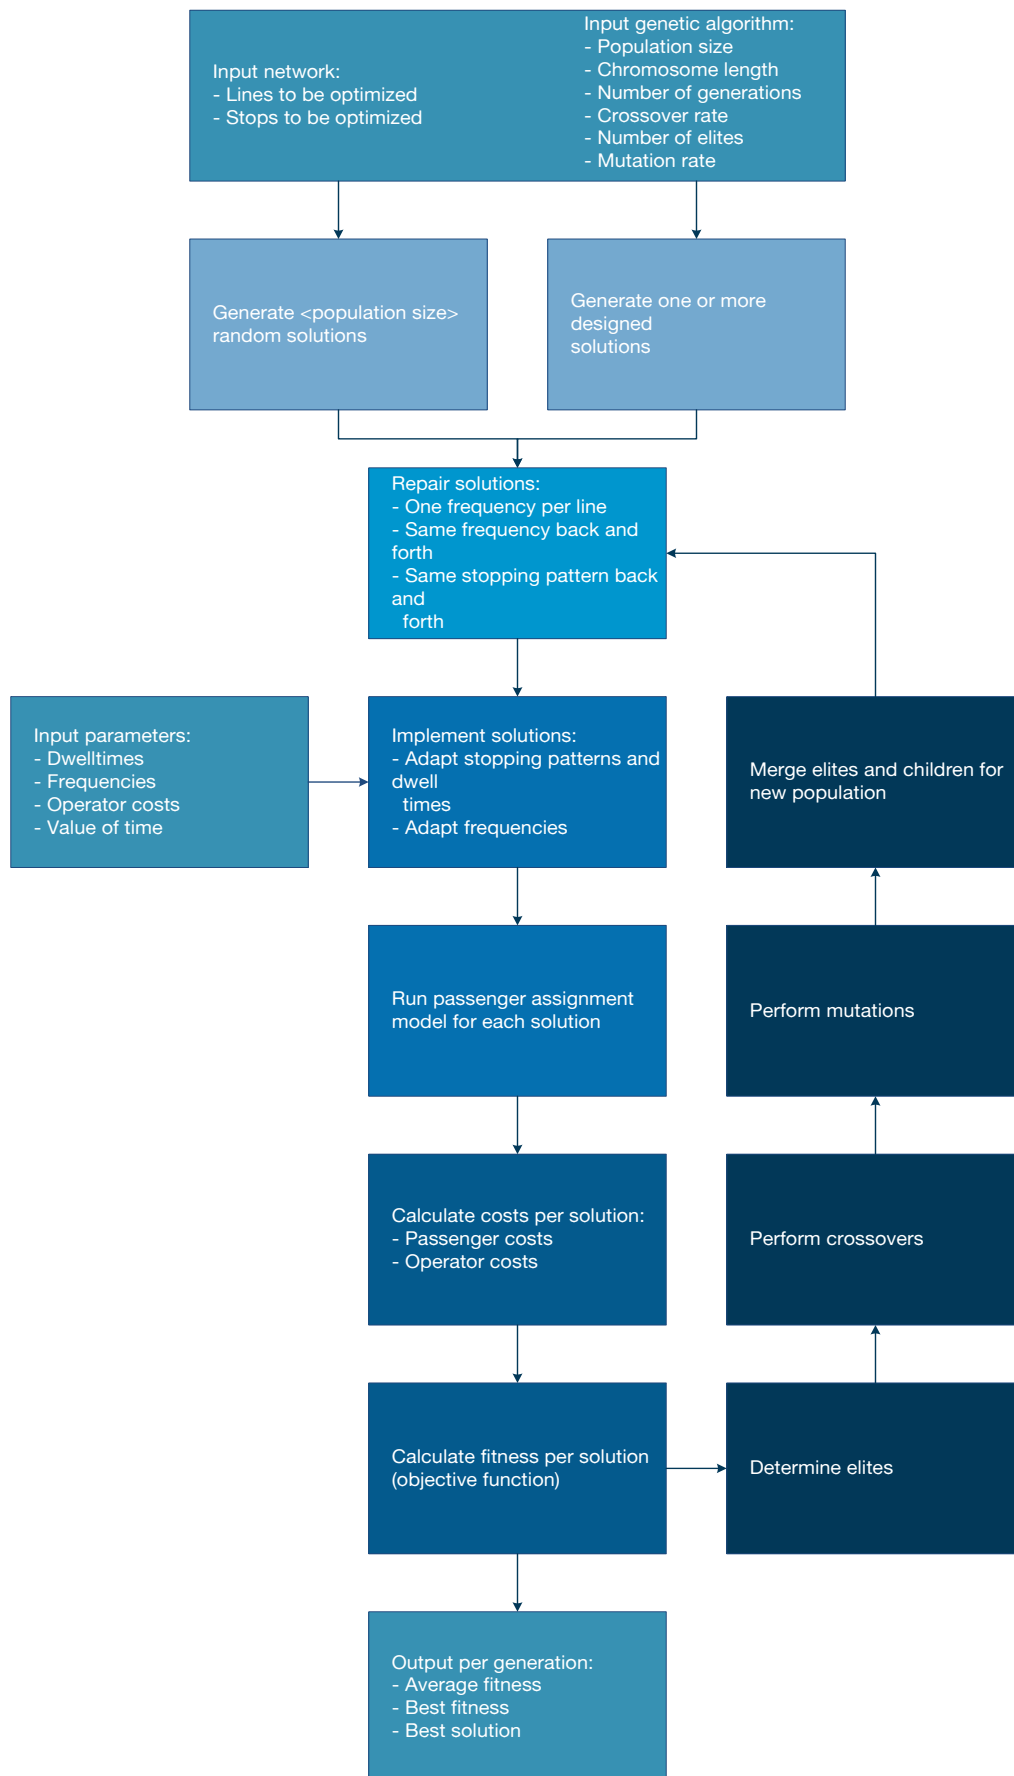


Figure 5.5: Model implementation flow chart

5.3.1 Input

The input of the model is separated into three categories: input considering the network to be optimised, input for the genetic algorithm, and other parameters. These parameter categories, and the parameters themselves, are discussed below.

- Input needed to construct the network:
 - Numbers of the lines that need to be optimised. These numbers can be derived from the OmniTRANS model and need to be put in manually.
 - The numbers of the stations that are part of the optimization have to be provided as well. This means that one part of a line can be treated as given network, while optimizing the other part. The only requirement is that if a stop has to be optimised in one direction, it has to be optimised in the other direction too.
- Input needed for the genetic algorithm
 - Population size: the number of solutions considered per generation.
 - Chromosome length: the number of decision variables per solution. This is calculated automatically from the input of the network. The length is determined by the number of stations that are part of the optimization, as each station is a gene of the chromosome, and the number of possible frequencies. Example: if a line has five stations, which are all part of the optimization, and the line can have a frequency of 0, 1, 2 or 4 times per hour, its contribution to the chromosome length is $5 + 4 = 9$ genes. More generally, the chromosome length is the number of stations plus the number of lines * number of possible frequencies.
 - Number of generations: the amount of cycles that should be run.
 - Crossover rate: the chance that a crossover between two solutions takes place. If there is no crossover, parent solutions will become children solutions and will thus be present in the population of the next generation. In SPAFOM, this rate is 1, as the best parents already become children through elitism.
 - Number of elites: if set higher than zero, the best solutions of a generation will be transferred to the next generation. This makes sure that the best solution of a generation is never worse than the best solution in the previous generation.
 - Mutation rate: the chance that a decision variable in a solution is changed. Each variable in a decision has the same chance of changing.
- Other parameters:
 - Dwell times: the process of braking, halting, and accelerating is simplified to one value, which represents the time lost. This dwell time value could differ per mode, but is the same for each stop on a line for a certain mode, independent of the number of people accessing and egressing a vehicle.
 - Frequencies: another input is a list of allowed frequencies for the network. This way, it can be made sure that a frequency is an integer (so no half runs per hour) and that clockwise logical runs are obtained.
 - Passenger costs parameters: the parameters to calculate passenger costs, such as the parameters for walking time, in-vehicle time, etc. (see the previous chapter), can be adjusted manually. These parameters are automatically adapted in the passenger assignment model as well.

- Operator costs: the costs for the operator are not part of OmniTRANS, so they have to be provided per mode. The operator costs are divided in fixed and variable costs. The first one is calculated per vehicle, the second one per vehicle-hour or vehicle-kilometre.
- Value of time: the value of time is used to convert time to costs for passengers.
- Vehicle capacity: the capacities of all different vehicles need to be given in order to calculate the crowding factor per link and per line.

5.3.2 Generate solutions

The first generation of solutions can either be designed, chosen randomly, or a combination of both. The first option means that each initial solution is constructed by taking certain design rules into consideration. A solution could for instance represent the current situation in a network, or have a local/express service pattern. It depends per case what design is implemented. This option can also be used to evaluate service designs, as it will become clear whether there is a better solution at the end of the model run. Either the initial solution has survived all generations as elite or a better solution has taken its place, in which case the costs can be compared.

In the random solution generation process, a solution is filled with 0's and 1's at random, until it has the length it needs to have. This method will probably yield worse solutions at first, but the diversity of the solutions might make up for that in the long run.

To combine the best of both ways, a part of the initial population of solutions will be designed beforehand, while the other part will be generated randomly. This will result in a shorter run time and will still yield good results fast. The number of elites will always be set higher than the number of designed solutions, so random solutions do not all extinct in the first generation and diversity is kept for feature generations.

5.3.3 Check solution feasibility

There are some constraints on a solution that need to be enforced in order for a solution to be feasible. The generation of solutions is checked for feasibility on the following criteria:

1. Each line has only one frequency. A line cannot run both 4 and 6 times per hour.
2. Each line has the same frequency in both directions. As a line back and forth is represented as two lines in OmniTRANS, all pairs of lines should have the same frequency to represent one bidirectional line.
3. Each line has the same stopping pattern in both directions. A passenger that takes one line from a certain station should be able to come back at that station with the same line in the opposite direction.
4. Node capacity is not exceeded
 - a. Platform capacity is not exceeded
 - b. Total station capacity is not exceeded
5. Link capacity is not exceeded

If a solution does not comply with any of the criteria, it is repaired (criteria 1 to 3) or discarded (criteria 4 and 5):

1. There are two ways a solution cannot comply with this criterion. Both require a different repair method:
 - a. No frequency is chosen. All binary values for choosing a frequency for a line are 0. In this case, a frequency is chosen randomly
 - b. More than one frequency is chosen. In this case, all but one of the chosen frequencies are set to 0.
2. If after step 1 the frequency for a line is not the same as its mirror, the lowest frequency is chosen. This way, a solution has a higher probability to comply with capacity constraints in 4 and 5 than when the highest frequency would have been chosen.
3. If a stopping pattern is not the same in both directions, the mirror line is adapted to match the direction coming first in the solution string.
4. Node capacity consists of two components, as discussed in section 4.3.2: platform capacity and through capacity. These components are checked simultaneously, as they influence each other. The difference is shown in Figure 5.6: platform capacity only includes the tracks at a node next to a platform that can be used by passengers to access and egress a vehicle; the total node capacity also includes tracks that pass a station, as if it were a link.




Figure 5.6: Node and link capacity

- a. Platform capacity: under the assumption that each track with a platform can harbour one vehicle at a time and the time a vehicle needs per stop is the sum of its minimum headway and its dwell time, the summed headways and dwell times of all stopping vehicles cannot exceed sixty minutes per track at a platform.
- b. Total node capacity: once the solution fits for all stations with regard to platform capacity, the passing transit lines are taken into account as well. The total time needed for stopping and passing lines at a station cannot exceed the available capacity, which includes both platform and through tracks.

5. The link capacity is assessed per direction. As lines are forced to be identical in both directions, only one direction of a link has to be checked. To assess link capacity, the headways of all passing lines in that direction are summed. The sum for a link cannot exceed its capacity of 60 minutes per track.

These reparations make sure that a solution is feasible before it is assessed. Therefore, no calculation time is wasted assessing infeasible solutions. Solutions that violate one or more of the capacity constraints after they have been repaired for the first three criteria are discarded. Random start solutions and children are only added to the population if they comply with all constraints. New random (first generation) or offspring solutions are generated until the new population has been filled. A test with 10.000 random solutions showed that almost half of them, 4.300, comply with the capacity constraints of the test network. Therefore, the run time of the model is not affected disproportionately by this operation¹.

The reparation process is mostly random: the frequency is selected randomly and the stopping pattern is selected randomly. However, there is a bias towards lower frequencies, to increase the chance of satisfying the capacity constraints.

5.3.4 Implement solution in OmniTRANS

The solution is implemented in OmniTRANS by adjusting the tables in the databases. Per station and line the stop type is read from the solution (stop or no stop) and the dwell time is adjusted accordingly (no dwell time or dwell time as defined in the input). Per line the frequency is read from the solution as well and written to the right table. When the solution is implemented, a static, deterministic All-or-Nothing assignment is performed. This is the default assignment method in OmniTRANS for public transport. There is one other option, which is volume averaging, but this is an iterative method used to take crowding into account. As it is iterative, the run time of the assignment model, accountable for most of the total model run time, will grow with the number of iterations. The run time is already an important limitation in the application of the model (see chapter 6), especially in the case study. Crowding is therefore not taken into account in the route choice, although it is accounted for in the passenger costs.

5.3.5 Passenger costs

The costs for passengers are divided into several components, as discussed in the previous chapter. These costs are calculated using the various tables OmniTRANS provides in its database.

- Walking time: the total walking time is calculated by first calculating the travel time on each walking link. This is multiplied by the number of passengers using the link and the sum of all links is the total time spent walking by all passengers. The method used here not only includes walking as a way of accessing and egressing the public transport network, but also between stations, if it occurs.
- Waiting time: as the transit line runs are frequency based, the total waiting time is calculated under the assumption that passengers arrive uniformly at all stations, no matter which mode they use to get at the station. Therefore, the average waiting time is half the

¹ The lower the percentage of unfitting solutions is, the longer the algorithm has to try before a fitting solution is found, which could in some cases lead to extremely high running times.

headway of the line they are waiting for. However, there is a limit of fifteen minutes, under the assumption that people will avoid a longer waiting time when the average waiting time is more than fifteen minutes.

- In-vehicle time: the time spent by passengers in a vehicle is calculated separately for each mode, to allow for a different weight of for instance comfort. The OmniTRANS database provides an overview per line with in-vehicle time already calculated, so it is simply summed per mode. The discomfort penalty for crowding is calculated using the occupancy of each line for each link and the capacity and frequency of that line, which is the same for all links the line traverses.
- Transfer penalties: per station, an overview is available with all transfers. All passengers that board a vehicle after just alighting another one are counted as transferring passengers. This ensures passengers are not counted twice, but only as they board. The penalty is the same for each vehicle combination. The extra waiting time due to the transfer is taken into account in the calculation of the waiting time; this part only calculates the penalties on top of that.
- Unsatisfied demand penalty: this is the penalty for passengers who cannot reach their destination. Such a passenger is assigned the largest travel costs of the original network. The travel cost skim matrix is accessed to search for where these costs are assigned and the demand between the OD-pairs where this is the case is multiplied by the penalty. The sum of all these cases is the total cost for unsatisfied demand. As mentioned before in the previous chapter, the unsatisfied demand penalty replaces the need for service constraints on stations.

Each travel time component has its own weight, to represent the preferences of passenger on how they like to spend their travel time. For example, passengers rather sit in a vehicle than that they have to walk. Each time component is thus multiplied by its weight, as well as the value of time (VOT). The VOT represents the monetary value a passenger gives to an amount of travel time. This way, travel time is converted to travel costs, so it can be summed with the operator costs.

5.3.6 Operator costs

The operator costs are separated into fixed costs per vehicle and variable costs per vehicle per hour. The run time per mode is obtained by summing the run time of each line served by that mode per link and adding the dwell time per station. This number is then multiplied by the costs per hour for that mode.

The run distance per line is obtained by multiplying the run time on a link by the speed on that link and the frequency of the line. The total distance per mode is then multiplied by the costs per metre for that mode.

The total fixed costs depend on the number of vehicles that is needed to run all services. For each line, the service time (which was calculated for the variable costs) is divided by sixty to go from hours to vehicles and rounded off upwards to whole vehicles. This number of vehicles is then per mode multiplied by the fixed costs per vehicle for that mode.

The totals of both variable and fixed costs are summed and form together the operator costs.

5.3.7 Evaluation

The passenger and operator costs are summed to get the value of the objective function, which is the fitness of each solution. The average and best fitness value is saved, together with the best solution. The elite solutions that should be kept for the next generation are determined by sorting the solutions on fitness value and selecting the best.

The solutions and their fitness values are saved after each generation, so if a solution is part of a next generation, there is no need to run the static assignment again. Instead, the fitness value of the solution is simply retrieved, thereby reducing the computation time significantly. This idea is adapted from the tabu search method, which uses memory to increase the diversity of the search. In this case, it is used to improve run time.

5.3.8 Improve solution

In this section, the actual genetic optimization takes place. The input for this is the first generation of solutions with their fitness values. A predefined amount of solutions is directly put in the next generation, the so called elite solutions. These solutions provided the best fitness values and are therefore preserved. All solutions, including the elite solutions, could then be chosen for a crossover. In Figure 5.7 this is shown graphically for a population of six solutions, of which the two best solutions are chosen as elites.

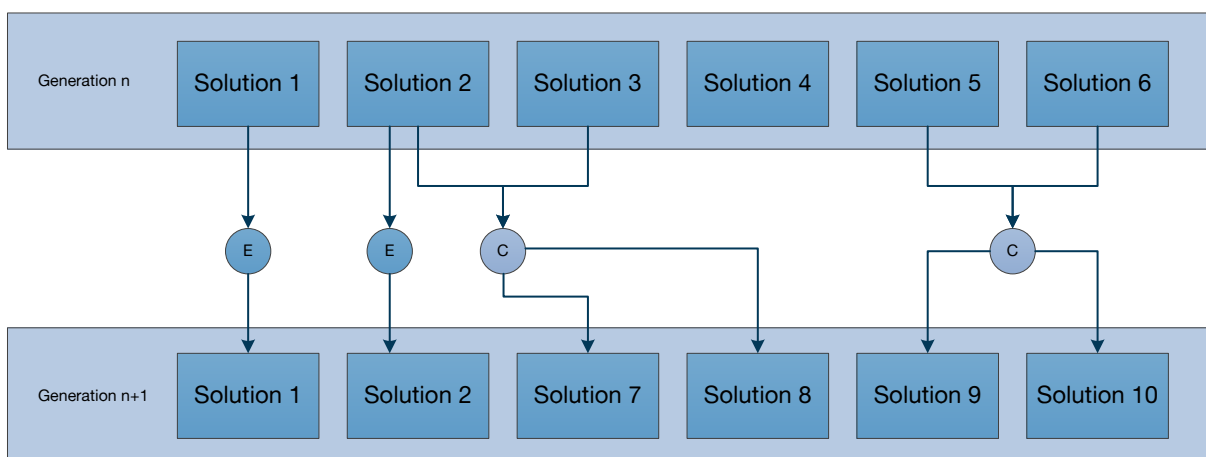


Figure 5.7: Composition of a new generation with elites (E) and crossover (C)

In the example, solutions 1 and 2 are the best solutions, so they are chosen as elite solutions and transferred directly to the next generation (A). This leaves four vacant places in the next generation, which need to be filled by children. This is done with two crossovers (B), the process of which will be explained in the next section. Solution 4 is nor an elite solution, nor a parent, and is therefore eliminated completely.

5.3.8.1 Crossover

During the crossover, two parents are selected based on their fitness values. This selection is based on a roulette wheel selection process. In this process, all fitness results are inverted (as this is a minimization problem, so higher values are actually worse solutions) and added together. A random number between zero and the total inverted fitness is drawn and the solution in which's range this random number falls is chosen as the parent. This way, it is ensured that the solution with the best fitness value has the highest chance of reproducing. An example is shown in Figure 5.8, which includes the following steps:

1. The inverted fitness values of all solutions are set in a row
2. A random point is chosen, where a line is drawn
3. The solution at the place of the line is chosen as the parent.




Figure 5.8: Example roulette wheel selection

The parents are always different solutions, although it is in theory possible that a solution occurs twice in a generation. In that case, the solutions are seen as separate individuals and could both be chosen as parent.

A uniform crossover is performed between the two parents. This results in two children, a blue child and a red child. Each gene of the blue child (so every 0 and 1) is the gene of one of the parents at that place in the chromosome. The red child automatically gets the gene of the other parent at that place. An example is shown in Figure 5.9, where the blue genes are chosen for child 1 and child 2 automatically gets the gene of the other parent.

Parent 1	1, 0, 1, 1, 1, 0, 0, 1
Parent 2	1, 1, 0, 0, 1, 1, 0, 0
Child 1	1, 0, 0, 1, 1, 0, 0, 1
Child 2	1, 1, 1, 0, 1, 1, 0, 0

Figure 5.9: Example uniform crossover

The nature of the problem makes other forms of crossover, for instance single or double point crossover, less suitable than the aforementioned method. In these cases the solution strings of both parents are namely cut in two (single point) or three pieces, after which pieces are exchanged between the parents to form children. This would mean that stopping patterns of all lines would

remain the same, except for the line where the cut was made. Only the mutation mechanism (see the next section) would then be able to perform this vital part of the optimization process.

5.3.8.2 Mutation

Each child then also has a chance of being mutated. Each gene has a very low chance of changing from 0 to 1 or the other way around. Mutation expands the search area covered by the solutions, which leads to a more diverse and complete search. However, it also interrupts the improvement process. A mutation can make a large difference, but only afterwards it becomes clear whether this difference is positive or negative. With enough generations, a “bad” mutation will be compensated and a “good” mutation will be improved, so mutation is a very valuable aspect of the genetic process.

5.3.8.3 New generation

The children, whether they are mutated or not, form a new generation, together with the elite solutions from the previous generation. From this point, the cycle starts over, with a reparation, implementation, etc., until the termination factor is reached.

5.3.9 Output

The output of the model is a list of the average fitness, best fitness and best solution per generation. Other data is possible as well and depends on the scenario that is being used. The solution can be run again to visually display it in OmniTRANS, for instance compared to the base case. The output that is being saved per generation is:

- Generation number
- Average fitness
- Best fitness
- Total passenger costs
- Total operator costs
- Walking costs
- In-vehicle costs
- Transfer costs
- Number of transfers
- Waiting costs
- Costs for unsatisfied demand
- Total fixed operator costs
- Total variable operator costs
- Best solution
- Run time for all generations combined

The output provides the possibility to extensively analyse the best solution and how it came to be the best of all assessed solutions. It shows the considerations made between all different costs in the model, what costs were dominant and where savings were made. The results of the applications in the next chapter will make this clearer, as then examples can be given.

5.3.10 Model verification and validation

The model was verified in parts, as each block of code was tested before combining the blocks. The test network is very suitable for this, as the results of each block were easy to verify by calculating the correct results using for instance Excel. Moreover, the calculations for the test network are much faster than the calculations for the existing network of Amsterdam, due to the size difference of the network. Each cost component is for instance a block that can be calculated independent of each other. The individual results were later matched with the outcome of the combined blocks, so the complete model, which outputs all components separately. As each block was independent, all results of the complete model matched the results of the individual block.

The model could not be validated in the available time. Most of the validation was done through earlier research, for instance the assumptions on the valuation of waiting time, in-vehicle time, access and egress time, and the transfer penalty. The assumptions on the structure of the system have not been validated explicitly either. However, as existing commercial software was used, together with an existing model, the assumptions of the system are directly obtained from these sources and assumed to be accurate. The route choice process of passengers on the basis of travel costs, for instance, is very intuitive and is a commonly used way to do this. Although such a way can never exactly capture the way an individual passenger chooses his route, the aggregated results are representative for all passengers.

5.4 Conclusion of the model implementation

In the previous sections, a detailed explanation on the working of SPAFOM has been provided, as well as an explanation on the genetic algorithm and how it was used in SPAFOM. The input of the model has been described, the way initial solutions will be generated and checked for feasibility, how the solutions are implemented in OmniTRANS, how the results of the assignment in OmniTRANS are used to calculate passenger and operator costs, how these costs are evaluated, how the solutions are then improved using the genetic algorithm and reassessed, and finally what the output of the model is.

In the next chapter, SPAFOM will be applied to a test network and the case of Amsterdam.

6 SPAFOM numerical experiments

SPAFOM will first be demonstrated in a fictional, simple network. In this demonstration, the capabilities and limitations of the model will be shown using several scenarios. The fictional network is suitable for this, as the results are predictable and easier to interpret than in an extensive existing network. It also provides the opportunity to run multiple scenarios, due to the limited run time.

First, the generic specifications of the model will be discussed in section 6.1. These specifications are used in the numerical experiments with the fictional network, which are performed in section 6.2. These experiments consist of six scenarios in the same network. A sensitivity analysis on the unsatisfied demand penalty is performed as well. Section 6.3 discusses the limitations of the experiments and conclusions are drawn in section 6.4.

6.1 Model specifications

Several types of data are needed as input for the model. These types of data have been discussed in section 5.3.1 and are discerned in input for the genetic algorithm (section 6.1.1) and other input (section 6.1.2).

6.1.1 Genetic algorithm

The input for the genetic algorithm has been described before: it consists of the population size, number of generations, chromosome length, crossover rate, number of elites and mutation rate. The values for these parameters are derived using trial and error and rules of thumb, which is discussed below.

- Population size: 26 solutions per generation are judged to provide a wide enough search area, while at the same time limiting the run time. More solutions per generation did not yield better results, while fewer solutions were judged to have too much risk of sticking at a local optimum.
- Number of generations: at 300 generations, the fitness of the best solution hadn't changed significantly for a large number of generations in several test runs. It is therefore judged that this amount of generations is enough to find a solution near enough to optimal. The only way to know for sure is by running the model indefinitely, effectively calculating the results of all possible solutions. This method is infeasible due to time constraints. Even a small network with three lines and three stations per line has $2^{3 \cdot 3} \cdot 14^3 = 1.4 \text{ million}$ possible solutions. In the very optimistic case that one calculation only takes one second, it would take over 16 days to run all solutions. The numerical experiment in this thesis has $2e^{25}$ possible solutions, so running them all is not an option.
- Chromosome length: only the chromosome length is fixed beforehand, as it has to represent exactly all decision variables: 92 possible stops + 20 lines * 14 possible frequencies = 372 genes.
- Crossover rate: the use of elites makes it unnecessary to save even more "old" solutions by making them children. Therefore, a crossover rate of 1 is chosen, so all crossovers will take place and lead to new solutions.

- Number of elites: 10 elites, which is a bit less than half the generation, works well. Again, trying different numbers yielded this result. For the first generation, this means that half of the elites could be a designed solution, while the other half consists of the best random solutions. It is assumed that this provides a basis broad enough to ensure diverse populations in feature generations.
- Mutation rate: this one is set to 0.01. This means that one in a hundred genes will change on average, so almost four in each solution. However, as this process takes place before the reparation process, there is a high chance that the mutation is made undone afterwards, especially if it concerns a change in frequency.

6.1.2 Other parameters

The other parameters are set as follows:

- Dwell time²: the dwell time is set to 2 minutes for trains. This is the time a train “loses” by stopping at a station and includes decelerating and accelerating, opposed to skipping a station. This number assumes a so called short stop of 30 seconds at a station. OmniTRANS assumes acceleration and deceleration to/from the maximum speed is instantaneous, which is the reason acceleration and deceleration is included in the dwell time. For metro the dwell time is set to 1 minute. This is due to the shorter stopping time and faster acceleration and deceleration.
- Headway: for train, the minimum headways are prescribed by ProRail and depend on the situation at a station (ProRail, 2017). In the worst case, the headway between two consecutive trains should be four minutes. Therefore, this value is used. For metro, headways of 90 seconds are technically possible. Two minutes is often used to allow for planning margins (LeighFisher, 2014), so that is the number that will be used in this model.
- Frequencies: The used frequencies in this case are: 0, 1, 2, 3, 4, 5, 6, 8, 10, 12, 15, 20, 30 and 60 runs per hour. This way, a schedule where vehicles depart on whole minutes is always possible. The exception is a frequency of eight vehicles per hour, but as this is a valid value in practice, it is still included.
- Passenger cost parameters: these parameters are adapted from Cats *et al.* (2016):
 - Walking time: 2.0 times the actual walking time.
 - Waiting time: 2.0 times the actual waiting time
 - In-vehicle time: 1.0 – 2.0 times the actual in-vehicle time, depending on the level of crowding (see Figure 6.1). Until a vehicle is halfway filled on average, no crowding is assumed to be experienced by travellers, hence the parameter is 1.0. From then on, the crowding experience linearly increases to 2.0, until the vehicle is filled with twice the amount of passengers it can comfortably carry. Passengers that travel with a vehicle that is filled to twice its comfortable capacity experience their travel time to be twice as long as it is in reality.

² The values for the dwell times were derived using the expert judgement and experience of Royal HaskoningDHV employees. The actual dwell time differs per station and depends for instance on the number of accessing and egressing passengers, but this is not taken into account in this thesis.

This is regardless whether the amount of passengers can actually be reached, because in reality people can be denied boarding due to physical capacity or legislation. Unfortunately, the static passenger assignment is not able to take denied boarding into account. This would take one or more iterations per solution, which is determined not to be worth the extra calculation time. No distinction is made between metro and train in terms of comfort, so the parameters are applied independent of mode.

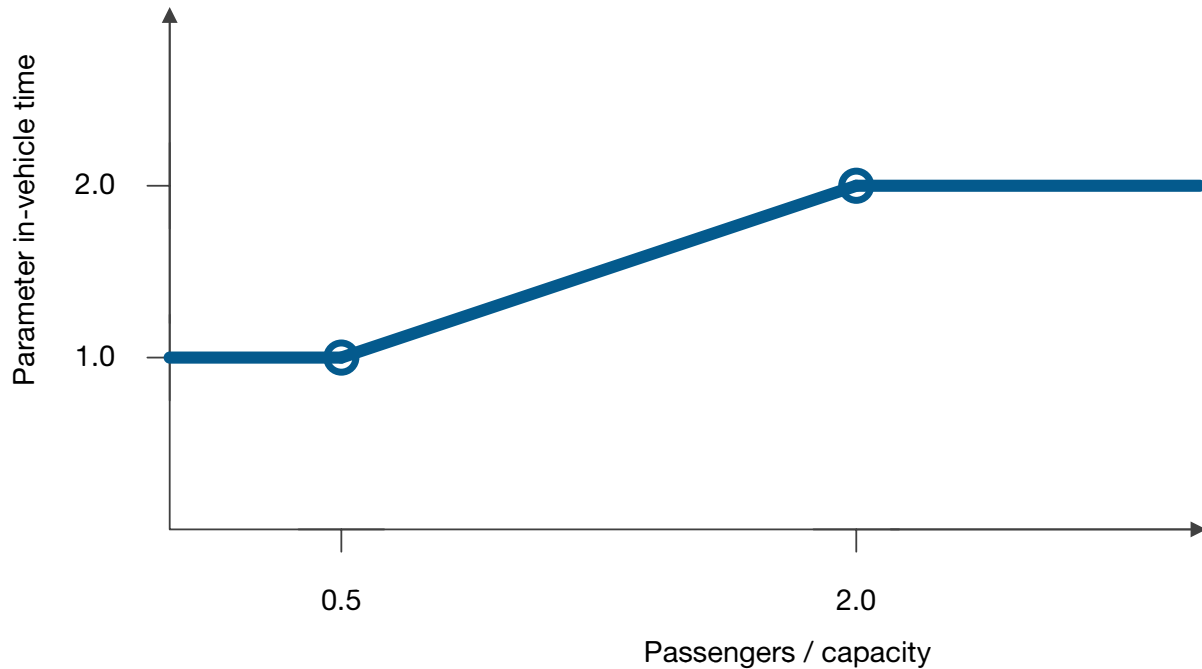


Figure 6.1: Parameter in-vehicle time

- Transfer penalty: 5.0 minutes per public transport to public transport transfer. This is the same in Verkeerskundig Noordvleugelmodel (VENOM) and has been found by Ortúzar and Willumsen (2011) as well.
- Value of time: The value of time, which represents the monetary value a passenger awards to travel time, is adopted from Kouwenhoven *et al.* (2014), which is to the author's best knowledge the latest study on value of time for The Netherlands. The average value for all surface transport modes was taken, as this number represents the overall VOT the best.
- Unsatisfied demand costs: these costs represent the penalty for passengers who are not able to reach their destination. It is derived from the original network, where all lines stop at all stations and run once per hour, as the maximum passenger costs in that network.
- Operator costs:
 - Fixed costs: fixed costs depend on the number of extra vehicles that are needed in comparison to what is available. They are not part of the test network, as it cannot be known how many vehicles of the total needed are already available. The large purchasing costs of such a vehicle (a couple of millions of euros for one train) will probably influence the outcome of the model in such a way that not running a line is always the cheapest option. In a real case, such as the case study later on, the number of vehicles needed can be compared to the number (expected to be)

available after the model has run, to get an idea of the costs to purchase, lease or renovate the extra vehicles needed, if any.

- Variable costs for train per hour: the variable costs per hour comprise of the costs for personnel on the train, which is assumed to be the maximum of 12 VIRM type carriages (see Figure 6.2) as two train sets of six carriages (two times VIRM-VI). It is assumed that each train has one driver and two conductors. The data for the costs come from a study by AT Osborne (AT Osborne, 2017).



Figure 6.2: NS VIRM train

- Variable costs for train per kilometre: the variable costs per kilometre consist of energy costs and maintenance costs. The first is calculated using key figures from CROW (CROW, 2015) for two coupled VIRM-VI trainsets. For maintenance, no key figures were available for trains, only for metro. Therefore, the maintenance costs are assumed to be proportional to the vehicle's weight. Two coupled VIRM-VI trains are four times as heavy as an M5 metro, so the costs are assumed to quadruple (*Metro | GVB, 2017; NS, 2017*). The same factor can also be found for energy usage.
 - Variable costs for metro per hour: the same costs are used for metro as for train, this time assuming one driver per train and one conductor per four vehicles, as not all metros have a conductor. Instead, teams of security employees are deployed to randomly patrol on metros. Experts from Royal HaskoningDHV agree with the assumption of one conductor per four vehicles.
 - Variable costs for metro per kilometre: for these costs the key figures from CROW were used as well, which indicated the energy and maintenance costs per kilometre.
- Vehicle capacity:
 - Train: for the capacity of a train, a VIRM with twelve carriages is used. This is the maximum amount of trainsets that is used in the Netherlands. With a capacity of 110 people per coach, the capacity of a VIRM in this setup is 1320 people (NS, 2017). This number takes comfort rules, for instance on standing places per m², into account.

- Metro: for the metro the newest type in Amsterdam, the Metropolis M5 metro, is used. This type has a capacity of 960 places, mostly standing places (*Metro* | GVB, 2017).

Table 6.1: Overview of model input for numerical experiments

Data type	Value	Source/method
Network	Original network	-
Genetic algorithm		
Population size	26	Best practice
Number of generations	300	Best practice
Crossover rate	1.0	Elites make not doing crossover unnecessary
Number of elites	10	Best practice
Mutation rate	0.01	Best practice
Other		
Dwell time train	2 minutes	Practice / expert judgement
Dwell time metro	1 minute	Practice / expert judgement
Headway train	4 minutes	ProRail (2017)
Headway metro	2 minutes	LeighFisher (2014)
Capacity train	1320 passengers	12 cars, NS (2017)
Capacity metro	960 passengers	6 cars, <i>Metro</i> GVB (2017)
Frequencies	0, 1, 2, 3, 4, 5, 6, 8, 10, 12, 15, 20, 30, 60 runs per hour	Possibilities for cyclic schedules
Parameter walking time	2.0	Cats <i>et al.</i> (2016)
Parameter waiting time	2.0	
Parameter in-vehicle time	1.0-20	
Parameter transfer	5.0	VENOM, Ortúzar and Willumsen (2011)
Value of time	€8.75/hour	Kouwenhoven <i>et al.</i> (2014)
Unsatisfied demand penalty	80.0 minutes	Original network
Operator costs train		
Fixed costs	0.0	Not taken into account
Variable costs per hour	€129.63	AT Osborne
Variable costs per km	€12,00	CROW
Operator costs metro		
Fixed costs	0.0	Not taken into account
Variable costs per hour	€59,82	AT Osborne
Variable costs per km	€3,00	CROW

Each station on a diagonal transit line has a total of 4 platform tracks, while all other stations have 2 platform tracks. There are 2 through tracks available at each station and each link has 2 tracks. As all lines are symmetrical back and forth, the tracks at all stations and links are automatically divided equally over the directions.

The stations are served by a total of 10 lines, all bidirectional. Each of the lines is represented by a different colour. Generally, there are two types: the diagonal lines with a distance of 15 kilometres between each station and the angular lines with a distance of 5 kilometres between each station. All lines stop at each station they pass and have a frequency of one run per hour in the original configuration.

The network is almost completely symmetrical. Only the second walking links are asymmetrical. This is done to be able to look at the behaviour of SPAFOM at zones with only one connection and zones with two connections in different situations.

The zonal data are almost symmetrical, just like the network. These data comprise of the number of residents and jobs. The zones can be distinguished in three groups:

- Centroid 1, connected to the middle station: this zone acts as the centre of the network and, in the base case, has most of the jobs (80.000). It also has the most residents of all zones (20.000).
- The corner centroids (zones 2 to 5): these have a relatively high number of residents (10.000) and some jobs as well (5.000).
- The other centroids (zones 6 to 20): these zones are smaller, with a relatively small amount of residents and jobs (both 1.000 for all zones).

The production of each zone is 0.9 times its number of residents and 0.1 times its number of jobs. The attraction of each zone is the opposite: 0.9 times its number of jobs and 0.1 times its number of residents. The trips are balanced on the production total, but as the number of residents and jobs is the same this is actually not necessary. These numbers were chosen to make sure that the number of travellers is easy to manipulate in scenarios, if needed.

The passenger flows are shown in Figure 6.4, where each centroid is split into origin (O) and destination (D). Each coloured line between an origin and a destination represents the passenger flow between from origin to destination. The width of the line depends on the amount of passengers; the colour depends on the origin. The slice of outer ring per zone depends on the total amount of passengers starting or concluding their journey at the zone. As was designed, centroid 1 is the largest in both ingoing (D1) and outgoing (O1) flows. The corner zones are large as well, while the other zones are relatively small.




Figure 6.4: Passenger flows in base case scenario

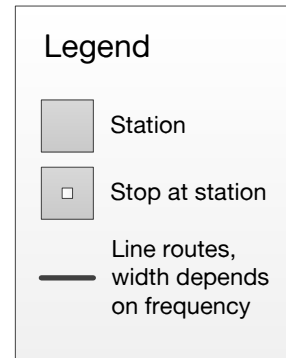

The exact number of travellers on the network is not important, as it is just a fictional network and any number can be used. However, the magnitude of the number of travellers is, as passenger costs are calculated per traveller. It can have an effect on the number of vehicles, especially when infrastructure capacity is taken into account. A comparison between the scenarios for costs per passenger will provide more insight in this.

6.2.2 Scenarios


The model will be tested in several scenarios, to test one or more aspects of the model or its expected outcome: the base case, with all parameters as described above; the high demand case, where the amount of residents and jobs in all zones is ten times higher than in the base case; the asymmetric demand case, where a corner zone has the highest demand; a high operator costs case, where the operator costs are assumed to be ten times higher than in the base case; a passenger cost only case, where operator costs are neglected and a passenger optimum is sought; a case the same as the base case, but without the infra limitations.

As mentioned before, several starting solutions are designed to give SPAFOM a head start and reduce the run time needed to reach a near optimum. The five designed initial solutions for all scenarios, shown in Figure 6.5, are as follows:


1. All lines stop at all stations they pass and have a frequency of once per hour. These are the same services with which the maximum passenger costs were calculated for the unsatisfied demand parameter.
2. The diagonal lines do not stop in the centre and have a frequency of four times per hour; the other lines stop at all stations and have a frequency of six times per hour. This represents real life services in the form of an intercity (diagonal) and metro (other).
3. All lines skip each second station they pass; diagonal lines run twice per hour and other lines four times per hour. This kind of service has fast services between stops that are served, but also leave some unsatisfied demand.
4. Diagonal lines skip centre and have a frequency of four times per hour; other lines do not stop at corners and have a frequency of twice per hour. This is a solution that is designed to connect the corners with each other and the rest of the zones with each other. It represents two separate, unconnected services and is used to produce variety in future generations.
5. No services on diagonal lines, other lines stop at all stations and run five times per hour. This network has many direct connections that run frequently, but at the cost of a longer travel time due to relatively many stops.




1. All lines stop at all stations and run once per hour.




2. Diagonal lines skip the centre station and run four times per hour.
Other lines stop at all stations and run five times per hour.



3. All lines skip each second station. Diagonal lines run twice per hour, other lines run four times per hour.



4. Diagonal lines skip centre station, run four times per hour.
Other lines skip corner stations and run twice per hour.



5. No diagonal lines.
Other lines stop at each station and run five times per hour.

Figure 6.5: Designed solutions for the test network scenarios

The scenarios that will be tested are the following:

- The base case scenario: as described above. It is used as a reference case for the other scenarios.
- The high demand scenario: the same as the base case scenario, except that zonal data are multiplied by a factor of 10. This leads to ten times as many travellers on the same network.
- The asymmetric demand scenario: instead of a centred demand, like in the base scenario, in this scenario the demand is centred on one of the corners of the network: centroid 2. Again, the centroids can be divided in three groups, based on their size:
 - Centroid 2: the largest centroid, with 20.000 residents and 30.000 jobs.
 - Centroids 1 and 3 to 5: relatively large centroids, with 5.000 residents and 2.500 jobs.
 - Centroids 6 to 20: relatively small centroids, still with 1.000 residents and 1.000 jobs.

This leads to a different movement pattern of passengers, as can be seen in Figure 6.6. This figure shows the passenger flows between all centroids, split out in origins (O) and destinations (D). Centroid 2 is clearly the largest one, as it both produces and attracts the largest amount of trips.




Figure 6.6: Movement pattern between centroids

- The high operator costs scenario: in this scenario, the operator costs are multiplied by 100, making them more important compared to passenger costs.
- The passenger costs only scenario: the opposite of the previous scenario, as operator costs are assumed to be zero. This leads to a passenger optimum.
- Unlimited infrastructure capacity scenario: the base case scenario is executed again, but this time without the limitations on infrastructure. In other words, all links and nodes have an infinite capacity. Vehicles' capacities are not unlimited, so crowding still leads to extra in-vehicle costs.

6.2.3 Results

The results of each scenario consist of the evolution of the average fitness per generation, the fitness of the best solution per generation and the different cost components of that best solution. The final best solution is discussed and compared to the initial best solution.

6.2.3.1 The base case scenario

The evolution of the fitness of the best solution for the base case scenario is shown in Figure 6.7.




Figure 6.7: Evolution base case scenario

The second designed solution, where the diagonal lines do not stop in the centre station, scored the best of all starting solutions at the beginning of the evolution. In this solution, the platform capacity of the central station is used completely. Figure 6.7 shows that the largest steps in the evolution process are caused by reductions in waiting time, which can be attributed to higher frequencies. This is made possible by a change in stopping patterns of some lines, leading to those lines not serving the central station anymore and making room for other lines to increase service frequencies.

Another important development in the process takes place around the 140th generation. Here, some demand is not served anymore, leading to unsatisfied demand costs. This increase in unsatisfied demand costs for a small amount of passengers is smaller than the decrease in waiting

costs for other passengers. The in-vehicle costs, transfer costs and waiting costs of the unserved passengers are avoided as well, so the net result of not serving a small part of the demand is positive as a whole. This also becomes clear from Figure 6.8. The resulting network is shown in Figure 6.9.




Figure 6.8: Cost comparison start and final solution base case scenario

While almost all costs have been decreased, it is clear that the waiting costs have decreased the most. The walking costs, which make up a considerable part of the total costs, stay more or less constant during the whole process. This was due to the starting solution that was chosen, which allowed all passengers to enter the public transport network at their nearest station.

The number of transfer, and with that the transfer costs, have stayed constant as well. As these costs were a relatively small part of the total costs, it was not profitable in terms of total costs to provide as many direct trips as possible. Instead, the available resources (operator costs and capacity) have been concentrated on the lines with the most passengers.

The slight reduction in in-vehicle costs is a result of the higher frequencies as well, as the result is an increase of the capacity of a line and a reduction of crowding costs. As mentioned before, another cause of this decrease is that a part of the passengers cannot travel to their destination anymore, leading to a decrease in passenger trips.

The variable operator costs have increased slightly, due to the higher frequencies of the line. Even though some lines have disappeared, overall more vehicles are used in the network, which corresponds to the increasing frequencies and the reduced waiting costs. The relatively small contribution of operator costs to the total costs makes it worth the extra operator costs to operate more vehicles. The resulting network and passenger flows are shown in Figure 6.9.

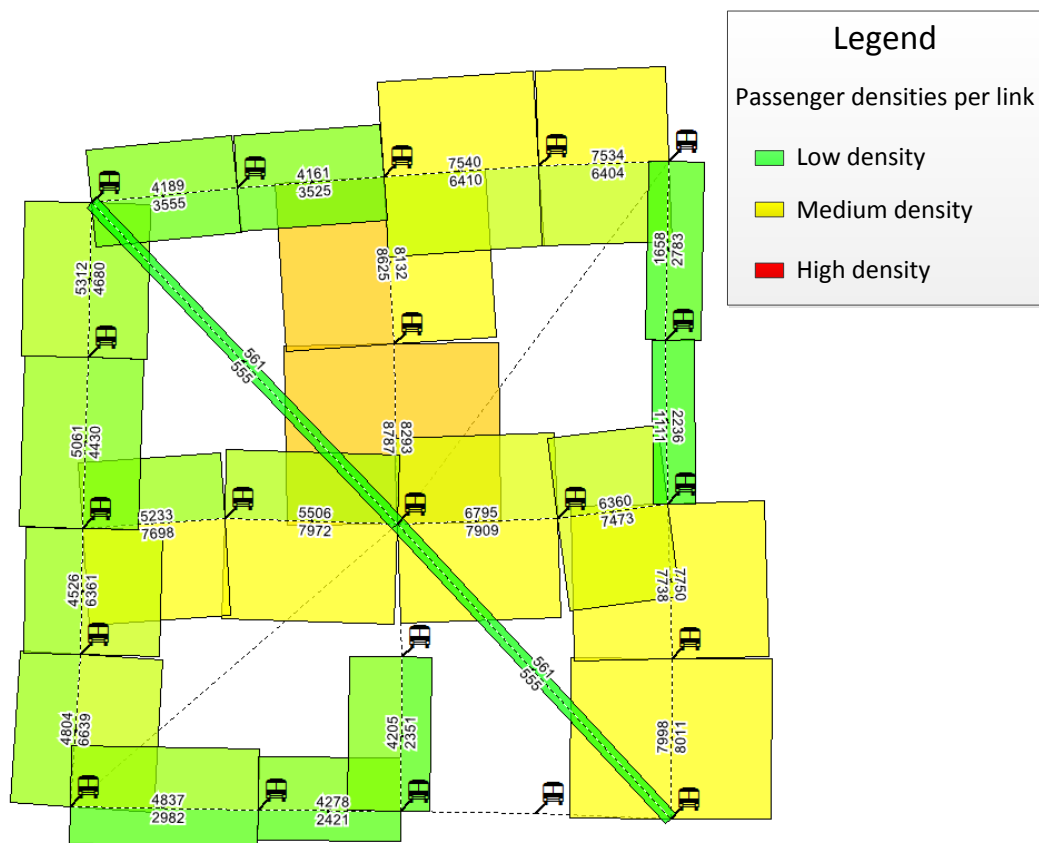



Figure 6.9: Final solution base case scenario and passenger flows per link

Some features of the starting solution can still be seen, but all but one line have changed in one way or another. Only the line from the upper left corner via the middle left station to the centre has not changed at all. Other lines have been shortened and/or are increased in frequency. One of the diagonal lines has disappeared completely, and so has one of the angular lines. This latter event has caused the unsatisfied demand, as one of the stations in the lower right quadrant is not served anymore. The zone connected to this station has no connections to another station and is thus disconnected from all other zones.

It is clear that the central station is the bottleneck in the network, as almost all changes take place in this station. Three lines do not stop in the middle anymore (four if the disappeared angular line is included), to make room for an increase in frequency of other lines. This has a large impact on the waiting time, as can be seen in Figure 6.8. Other than that, not many changes are seen in the stopping patterns. If a station has enough capacity, a line will stop there. The extra time this takes for passengers who do not alight or enter the vehicle at that station is apparently worth the extra convenience for those passengers that do use the station. The remaining diagonal line does not get a stop in the central station, which means that it is beneficial from a cost perspective to provide many zones a direct and frequent but slower connection, instead of providing one zone a faster and direct connection.

This also becomes clear from the passenger flows, as passengers tend to use the more frequent lines more often³. This becomes especially clear from the upper right quadrant, where the more frequent line serves almost three times the amount of passengers as the other line, despite the symmetrical OD-matrix. Besides, the more frequent line stops at the central station, where the other line does not.

The remaining diagonal line is not used much, which explains why the other diagonal line was removed. The unserved station in the lower right quadrant also stands out in the figure, due to the other removed line. This removed line was probably the reason one of the diagonal lines remained in the solution, to serve the lower right station.

Overall, it can be concluded that the network performed better when lines have high frequencies, even if this leads to other lines being not operated anymore or not to their full extent. Maximisation of vehicles in the network leads to a minimization of waiting time, which appears to be the cost factor that can be reduced the most.

³ NB: Passengers are aggregated if a link contains multiple transit lines. However, for route choice the line frequencies are not. See also section 6.3.

6.2.3.2 The high demand scenario

The evolution of the high demand scenario is shown in Figure 6.10.




Figure 6.10: Evolution high demand scenario

The costs for the solutions in this scenario are more than ten times higher than in the base case scenario, because the amount of passengers is ten times higher. Another part of the cost increase is the higher level of crowding, as the capacity of the infrastructure has not changed and neither has the amount of vehicles, nor their capacity.

Again, the second designed solution comes out best in the beginning and again most improvement is seen for the waiting costs. Unlike the base case scenario, there is no unsatisfied demand in any of the best solutions per generation. It looks like there are no variable operator costs as well, but this is caused by the tenfold increase of the passenger costs, while the operator costs have stayed the same.

Otherwise, there are no large differences with the base case scenario. The resources are concentrated on the angular lines. Frequency is more important than speed and, as a result, waiting costs are reduced the most, even though the in-vehicle costs are larger due to crowding. The variable operator costs, which are put on a secondary axis in Figure 6.11 to be visible, have not risen very much. This indicates that with a small effort an operator can significantly reduce in-vehicle and waiting costs for passengers. Walking costs have not changed at all, which means that every passenger can still use his nearest station to use a transit line, as was the case in the first solution. Transfer costs have risen slightly, which is caused by the lines that terminate before reaching the central station. Therefore, passengers to the central station and further have to make an extra transfer. The higher frequencies of the lines compensate this, reducing both in-vehicle and waiting time.




Figure 6.11: Cost comparison first and final solution high demand scenario (operator costs on separate axis)

In the final solution, the same pattern as in the base case scenario can be seen: the capacity of the central station is the most important factor (see Figure 6.12). In order to still decrease waiting time by increasing frequencies on some lines, frequencies of other lines need to decrease or lines do not stop in the centre anymore. The most striking feature of the final solution is that the diagonal lines do not run anymore. Instead, from all but the upper right corner station angular lines run with a high frequency towards the central station.

Again, the passenger flow distribution is as expected. If a passenger can choose between two lines, the chance he will choose one line depends on the frequencies of both lines. The same patterns are seen as in the base case scenario.

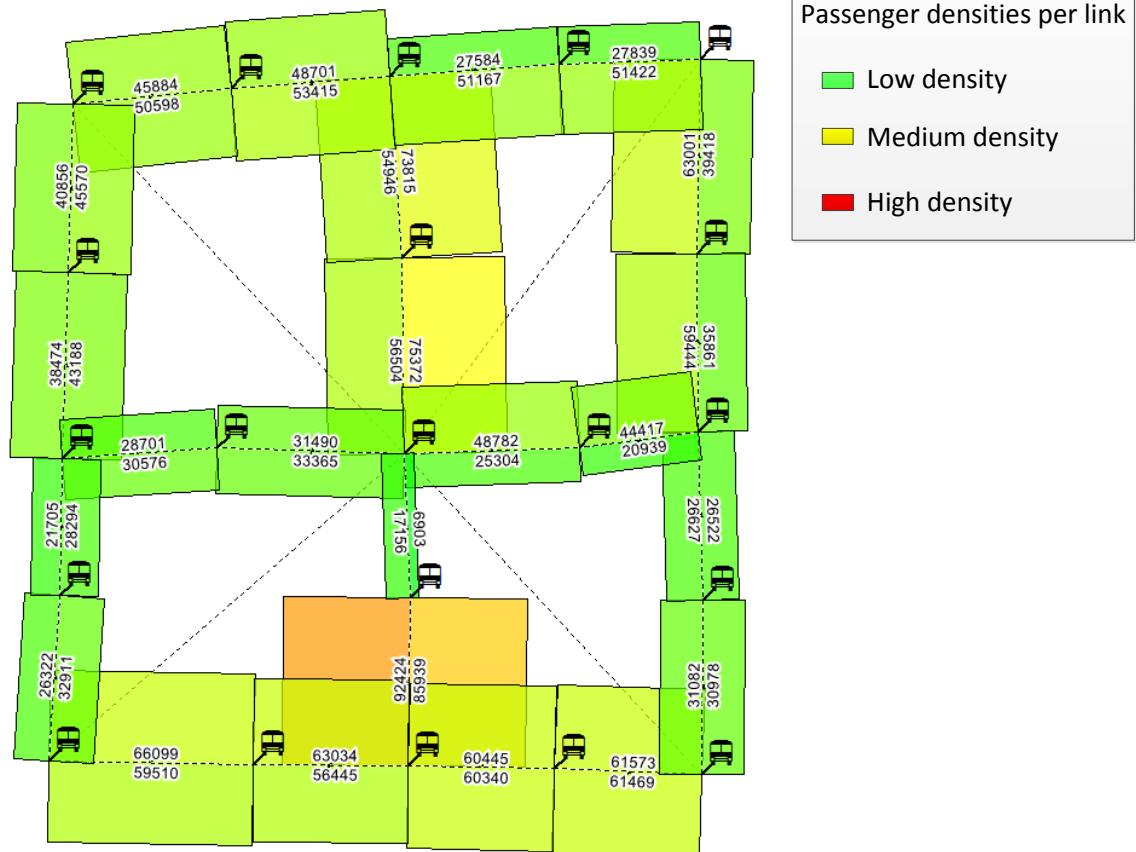
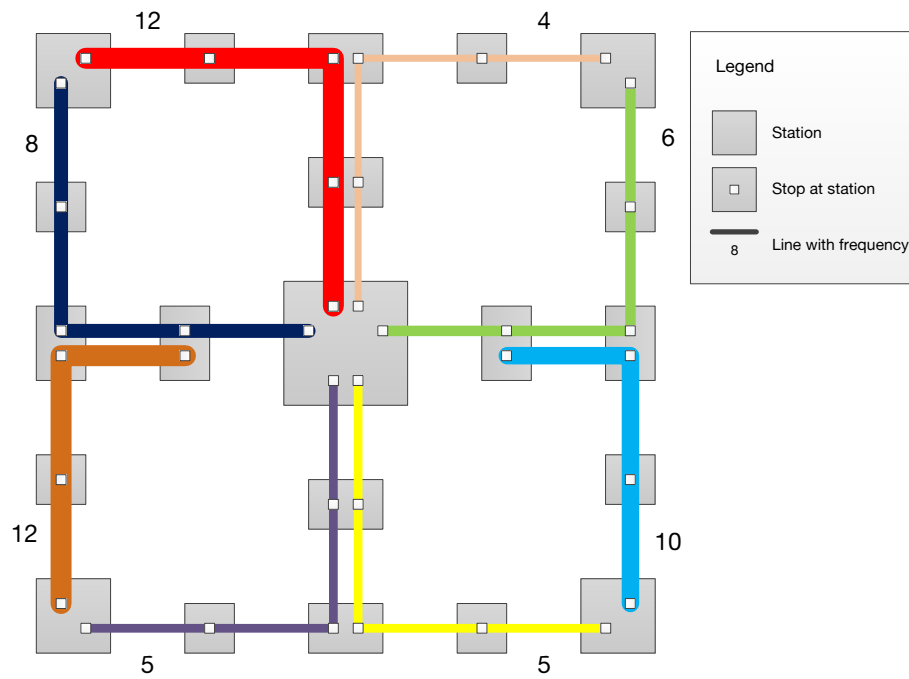


Figure 6.12: Final solution high demand scenario

6.2.3.3 The asymmetric demand scenario

Figure 6.13 shows the evolution of the different costs for the asymmetric demand scenario.




Figure 6.13: Evolution asymmetric demand scenario

Just as in the base case scenario, designed solution 2 scores the best in the beginning. From there on, the total costs decrease proportionally to both the in-vehicle and waiting costs. All other costs stay about constant and there is no unsatisfied demand during the process (see also Figure 6.14).




Figure 6.14: Cost comparison first and final solution asymmetric scenario

Where in the previous scenarios the waiting costs were clearly the largest factor of the cost reduction, in this scenario the in-vehicle costs are reduced just as much. From Figure 6.15, it becomes clear that this is not because of the faster routes, as none of the lines skip a stop that was served in the first solution (although there are some lines that terminate earlier than before). Instead, an increase in frequency for almost all lines has reduced the level of crowding on those lines.

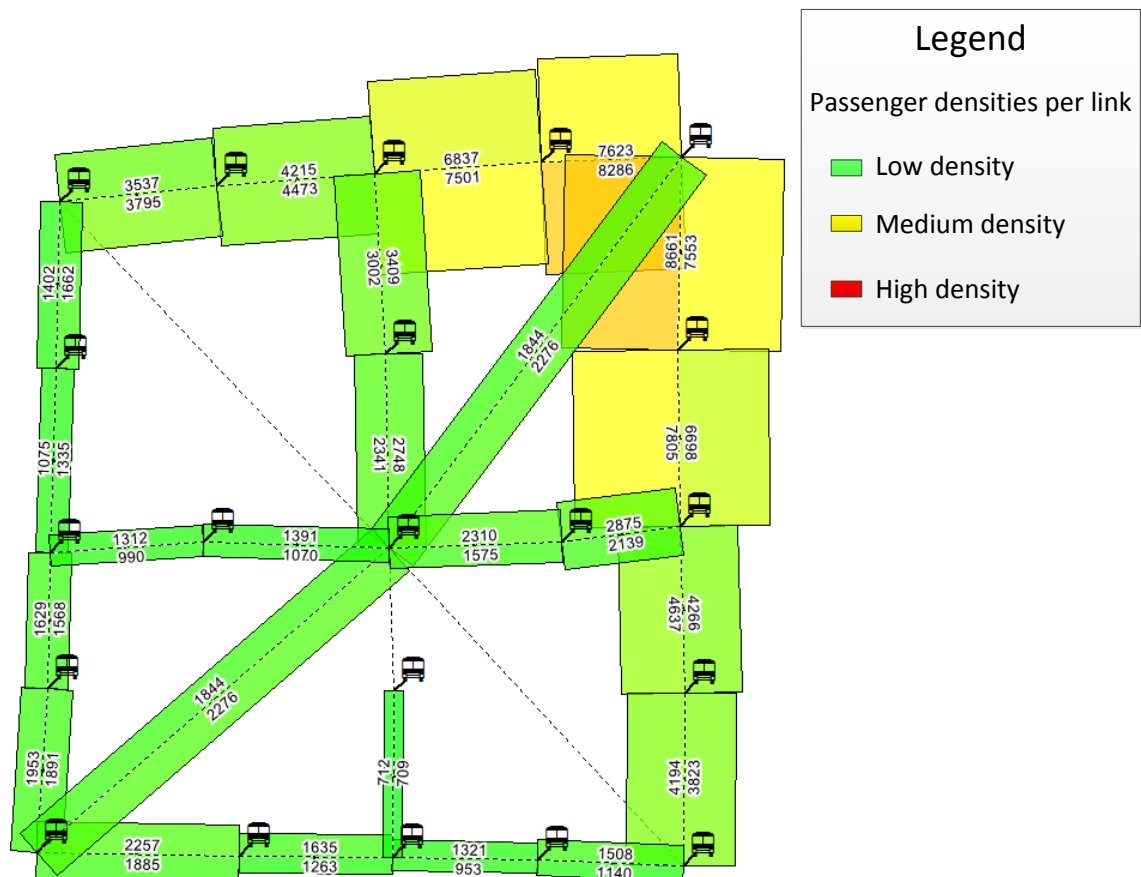



Figure 6.15: Final solution asymmetric scenario and passenger flows per link

Furthermore, the lines closer to the largest zone in the upper right corner tend to an increased frequency compared to other lines. Where in other scenarios, each angular line has approximately the same importance, in this scenario there is a clear hierarchy. In this hierarchy, more important lines get more resources. Another trend is that lines, which are less important to access the largest zone, adapt their stopping pattern to allow for an increased frequency of the other lines within the capacity boundaries. Practically, this means that lines do not stop in the central station anymore, which allows for a significant increase in frequency.

Because of the high production and attraction of the upper right corner, the number of routes a passenger can take to his destination is limited compared to the other scenarios. This shows in the passenger flow distribution, as lines closer to the upper right corner carry more passengers than lines further away, regardless of the frequency.

6.2.3.4 The high operator costs scenario

The evolution of the costs in the high operator costs scenario is shown in Figure 6.16.




Figure 6.16: Evolution high operator costs scenario

The impact of the extra variable operator costs is immediately clear, as the relative size of the operator costs is much bigger. Designed solution 5, where there are no diagonal lines, all other lines stop at all stations and run five times per hour, is the best solution of the first generation. The amount of improvement over 300 generations is smaller than in other scenarios as well, but still some improvement is made.

Figure 6.17: Costs comparison first and final solution high operator costs scenario shows that most improvement is made in in-vehicle costs. Reduced waiting costs make up the other part of the improvement, as walking and transfer costs are constant. Variable operator costs have increased slightly.




Figure 6.17: Costs comparison first and final solution high operator costs scenario

The stopping pattern of the final solution (see Figure 6.17) looks very much like the first solution, as was expected from the small overall improvement in Figure 6.16. Some frequencies have changed and two lines do not stop in the central station anymore. The reason these lines terminate before the central station, is that the large central zone can be reached via the new terminal station as well.

Again, the preference for lines with higher frequencies becomes clear from the passenger flow distribution. As there are no large differences in frequencies between the lines, the passenger flow distribution is reasonably uniform over the network.

Overall, the results of the model were to be expected. In-vehicle and waiting costs are still the passenger costs that can be decreased most, at the expense of extra operator costs. The increased operator costs in this scenario limited the improvement of passenger costs, which resulted in a somewhat minimalistic version of the first solution. Still, some frequencies were increased and most stations with more than one stop are still served by more than one line, despite the extra operator costs. This shows that passenger costs can outweigh operator costs, even when the latter are very high, and that resources can best be concentrated on the lines serving the most passengers.

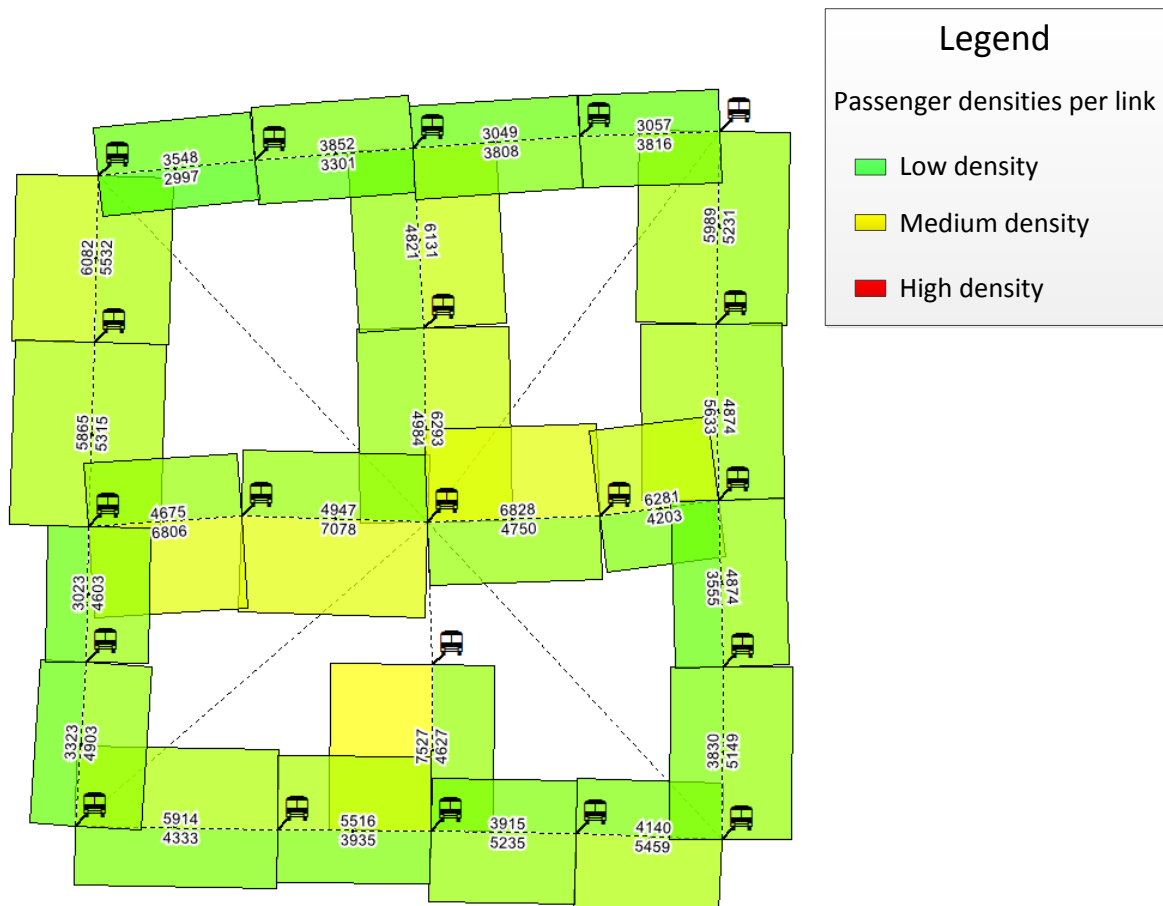
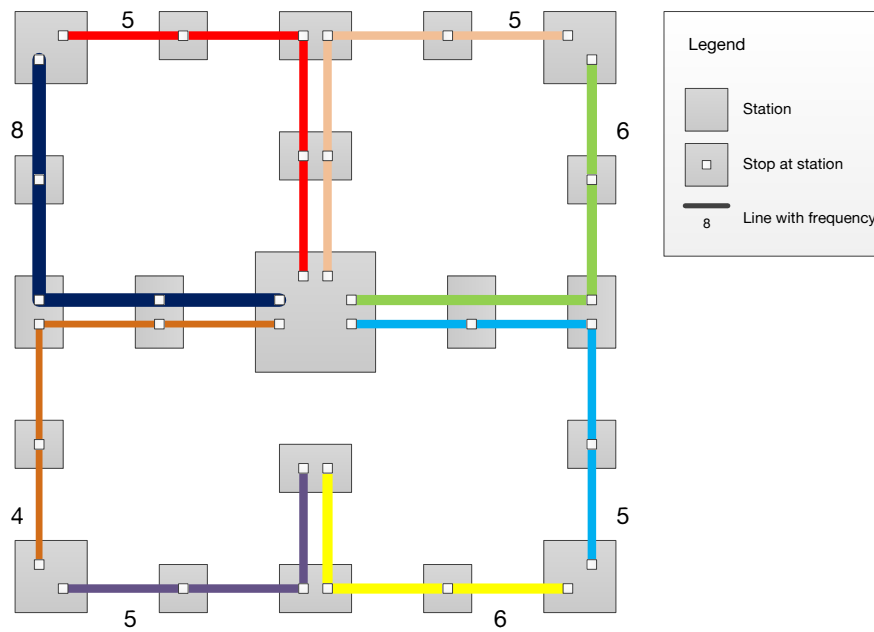


Figure 6.18: Final solution high operator costs scenario and passenger flows per link

6.2.3.5 The passenger costs only scenario

In Figure 6.19, the results for the passenger costs only scenario are shown.




Figure 6.19: Evolution passenger costs only scenario

Just as in most previous scenarios, the best solution of the first generation is the second designed solution. This remains the best solution for a short period, after which several optimizations take place. Around halfway, the final solution is already found. From Figure 6.20 it becomes clear again that in-vehicle and especially waiting costs can be reduced most, while walking and transfer costs stay constant and there is no unsatisfied demand.




Figure 6.20: Cost comparison passenger costs only scenario

The final solution is shown in Figure 6.21. The solution does not differ much from the first solution. All changes can be related to fitting as many vehicles in the network as possible, by changing stopping patterns to make room in the central station. As the influence of the operator costs was already very limited in earlier scenarios and the infrastructure capacity was guiding for the solutions, there is not much difference with earlier scenarios.

What catches the eye is the presence of both diagonal lines. As they do not stop in the central station, they are not subject to the capacity constraints there: they simply use the through tracks. As operator costs are ignored in this scenario, there is no reason to eliminate the lines, as they are free of charge. Still, the diagonal lines have a low frequency compared to the other lines, which means they do not do much in terms of passenger costs reduction.

The diagonal lines do not serve many passengers, compared to the other lines. This shows that relatively few passengers profit from them. The same goes for a part of the network that is often not served in other scenarios, which is the link from the central station downwards. This is caused by a combination of two factors: first, the central zone is connected to that link, so there are many passengers who depend on that link. Second, passengers travelling from the lower right station to the central zone have an extra transfer if they use the lower right line, as their line does not reach the central station. It is therefore easier for them to use the other line, even though the line's frequency is a bit lower.





Figure 6.21: Final solution passenger costs only scenario and passenger flows per link

6.2.3.6 The unlimited infrastructure capacity scenario

The evolution of the scenario without infrastructure limitations is shown in Figure 6.22.




Figure 6.22: Evolution unlimited infrastructure capacity scenario

Also in this scenario, the second designed solution is the best of the first generation. Not hindered by capacity constraints, the total costs are reduced significantly within the first hundred generations. This continues up to almost the end of the model run, where a decrease in especially waiting costs is achieved: only a third of the original waiting costs have remained (see Figure 6.23).

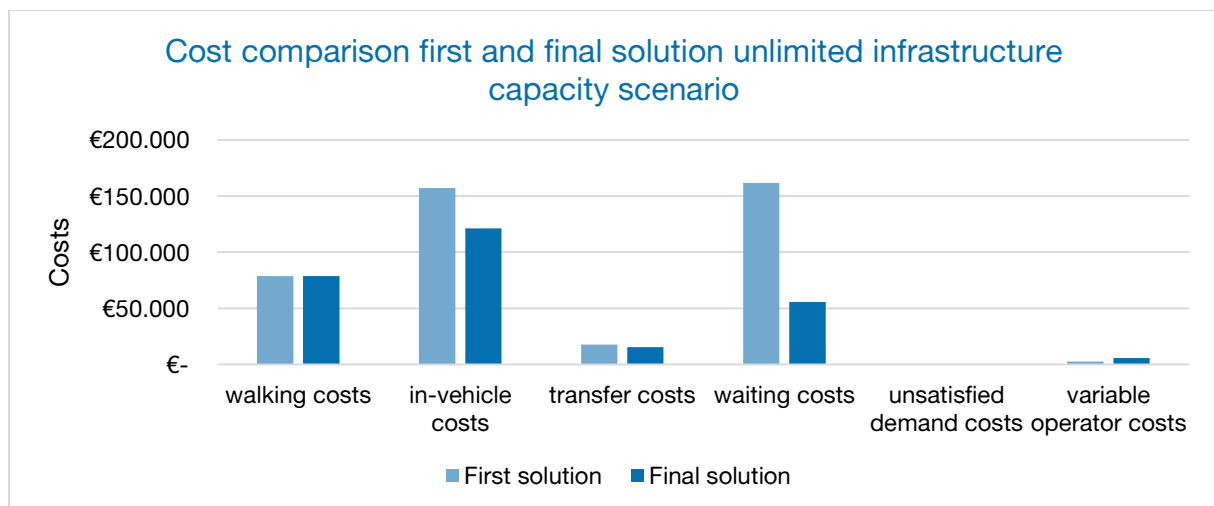


Figure 6.23: Costs comparison first and final solution unlimited infrastructure capacity scenario

Moreover, the in-vehicle costs have decreased significantly as well. Although the operator costs have more than doubled, the effect on total costs has been very large. Figure 6.24 shows that this is mainly due to a very large increase of frequencies for almost all lines. Despite the lack of capacity constraints, the diagonal lines are still removed from the network.

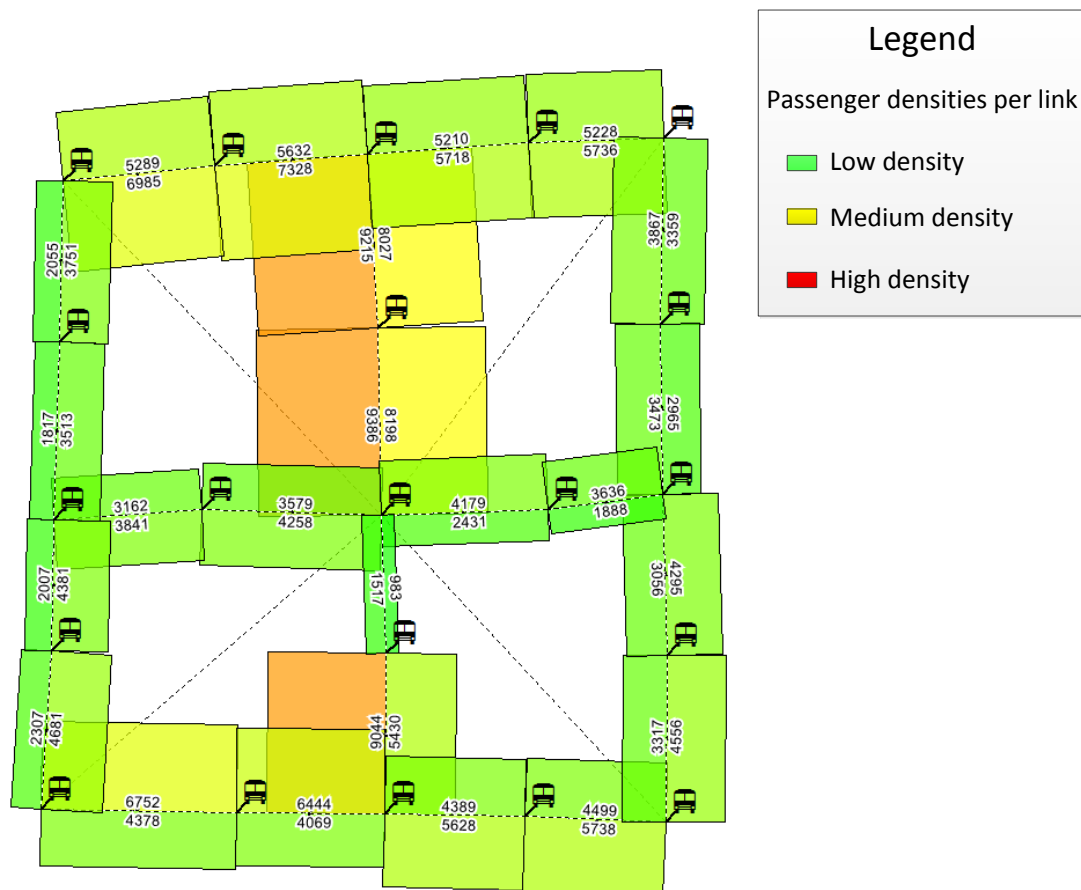
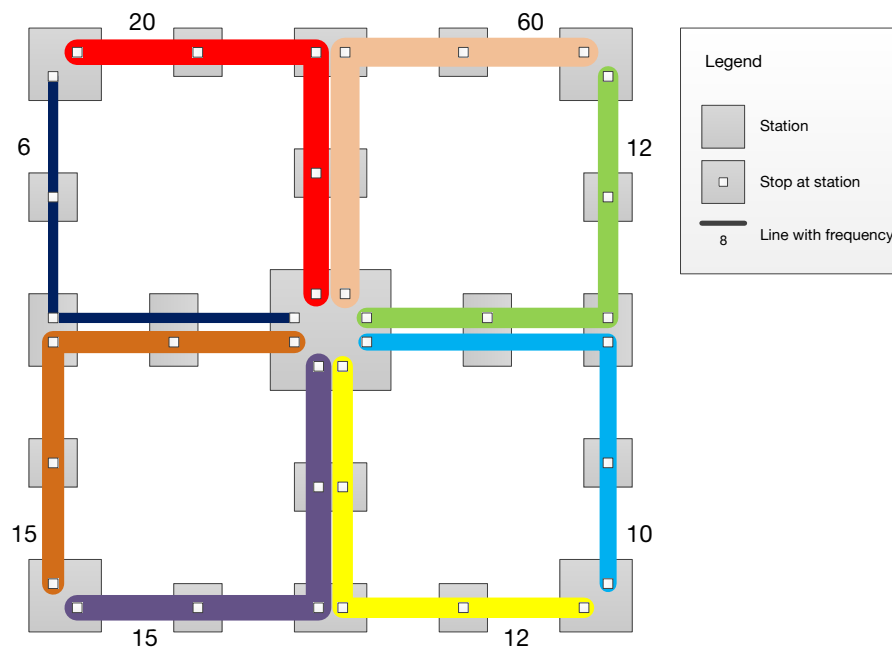


Figure 6.24: Final solution unlimited infrastructure capacity scenario and passenger flow per link

It is clear that the lines with the highest frequencies attract the most passengers, as was seen in the previous scenarios. All in all, the passenger flow distribution is as expected. The trend of the evolution in Figure 6.22 gives reason to suspect that, given extra calculation time, the end result would be that all angular lines run with a maximum frequency. This is already the case for one line. The effect of this on the waiting costs is so large, that the limited effect on operator costs is more than compensated. This corresponds with the results of earlier scenarios, where the waiting costs were the dominant factor for reducing costs as well.

6.2.3.7 Scenario comparison

In the previous sections, already some attention has been paid to the similarities and differences between each scenario. In this section, the results of the scenarios will be compared per cost component. This comparison only concerns the final solutions for each scenario.

First, the total costs, split in passenger and operator costs, are compared (see Figure 6.25). To make a fair comparison, the costs per passenger are used, as the scenarios do not all have the same amount of passengers.




Figure 6.25: Comparison of passenger and operator costs per passenger per scenario

This comparison shows that scenario with unlimited infrastructure capacity clearly has the least passenger costs per passenger and the scenario with high operator costs the most. This corresponds with earlier findings that frequency is an important part of the optimization and that allowing more (unlimited capacity) or fewer (high costs) has a large impact on the final solution.

The other three scenarios score about the same. As the result of the scenario with only passenger costs does not differ from the base case, it can be concluded that including operator costs in the way it was done in this thesis does not have a significant impact. The high operator costs scenario shows that including a larger operator costs per unit of time and/or distance does have a significant impact.

Next, a comparison of costs per passenger is shown in Figure 6.26 for all cost components.




Figure 6.26: Comparison of costs per passenger

The walking costs are the same for each scenario, only the base case deviates slightly. This is caused by the passengers whose demand is not satisfied, as these passengers cannot reach their destination and therefore do not walk. In each scenario, every travelling passenger can do so via his closest station.

The in-vehicle costs are fairly constant over all scenarios, although there are two outliers: the in-vehicle costs per passenger in the unlimited infrastructure capacity scenario are relatively low, while those costs for the high demand scenario are very large. The reason for both is the amount of vehicles compared to the number of passengers. In the unlimited infrastructure capacity scenario, the number of vehicles is relatively high, while in the high demand scenario, by design, there are more passengers for the same amount of (allowed) vehicles.

The transfer costs per passenger are more or less the same for each scenario, with the exception of the asymmetric demand scenario. This is caused by the change in origin and destination from passengers in that scenario. As the largest zone is no longer in the middle, but in a corner, fewer lines can be used to reach the largest zones. Therefore, more passengers have to transfer to travel to or from that zone.

The other differences between zones are caused by the way the solution dealt with the capacity constraint of especially the central station. This lack of capacity resulted in all scenarios in a reduction of lines reaching the central station, but how many lines were shortened differed per scenario. The unlimited infrastructure scenario did not have to do this, hence in this scenario fewer transfers were needed.

The results of the comparison of the waiting costs per scenario are more or less the same as for comparison of the in-vehicle costs. Keeping in mind that the final solution for the base case scenario included unsatisfied demand, only the unlimited infrastructure capacity scenario and high operator costs scenario significantly differ from the rest. Again, this is caused by the amount of vehicles in the network, as higher frequencies lead to less waiting time and thus less waiting costs.

The comparison of operator costs per passenger (Figure 6.27) is straightforward. The high operator costs scenario has by far the highest operator costs per passenger, over €2,50. The graph is cut at €0,10 to show the other scenarios' costs as well. The unlimited infrastructure capacity scenario resulted in the second largest operator costs, as this scenario allowed for more vehicles. There is not much difference between the asymmetric demand and base case scenarios; the asymmetric scenario's costs are slightly higher because of the extra diagonal line. The costs for the high demand scenario are low, because of the extra passengers that travel with the same amount of vehicles, and the scenario with only passenger costs has no operator costs by design.




Figure 6.27: Comparison of operator costs per passenger

The conclusion of the comparison is that the differences in costs between the scenarios are mostly caused by the difference in vehicles and passengers. A higher amount of vehicles in a network is beneficial for passengers, as both in-vehicle (crowding) and waiting costs are reduced when the frequency of a line is increased. However, a higher demand for travelling leads to higher costs, if there are not enough vehicles available.

Furthermore, the orientation of the travellers does not have much influence on either passenger or operator costs. An asymmetric demand leads to more transfers, if the largest zone(s) with respect to production or attraction of passengers is served by fewer lines. This directly leads to more waiting costs, as every new transfer leads to extra waiting time as well. However, the concentration of passengers on lines means that frequency increases can concentrate on those lines as well if capacity is limited, as was the case.

An enormous increase of operator costs per unit of distance and/or time has little influence on the costs for passengers. Although the total passenger costs were slightly higher than in the base case, roughly the same solution was found. Neglecting operator costs influences the final solution even less, as the relative share of operator costs in the total costs was already very limited.

The best option to reduce total costs turned out to be increasing infrastructure capacity. This allows for higher frequencies and longer lines throughout the network, reducing in-vehicle and waiting costs. These factors were found to be the most influential in all scenarios, and were decreased by concentrating resources on the most important lines.

6.2.4 Sensitivity analysis

Not all parameters could be determined with absolute certainty. A sensitivity analysis is therefore performed, using the base case scenario to test the influence of the height of the unsatisfied demand penalty. As discussed before, the maximum passenger costs were taken from the original network to represent the cost of unsatisfied demand, but there is little scientific evidence to back that up. All other parameters have been discussed in literature extensively and the values of those have therefore been established with enough certainty.

The analysis is performed by running the base case scenario with both a 25% decrease and a 25% increase in unsatisfied demand costs per unconnected traveller. The results of the former are shown in Figure 6.28.




Figure 6.28: Results sensitivity analysis -25% unsatisfied demand costs

Surprisingly, there is no unsatisfied demand in the final solution of this scenario. It was expected that there would be unsatisfied demand, as it was a part of the base case scenario final solution and the penalty for it has been reduced. Instead, it follows the trend of the other scenarios, where in-vehicle and waiting costs are reduced the most, while other passenger costs remain more or less constant and variable operator costs increase slightly (see Figure 6.29).




Figure 6.29: Cost comparison first and final solution sensitivity analysis -25% unsatisfied demand costs

In Figure 6.30, it can be seen that the solution does not really differ from all other solutions. Just like the costs, the network shows the same trends as in the other scenarios: lines are modified in such a way that frequencies can increase. However, it is remarkable that in this case both diagonal lines are part of the final solution. This was only the case in the scenario without operator costs.




Figure 6.30: Final solution sensitivity analysis -25%

The results of the sensitivity analysis with a +25% increase for unsatisfied demand costs (see Figure 6.31) show no surprises. As expected, waiting and in-vehicle costs are reduced the most and there is no unsatisfied demand (see also Figure 6.32).




Figure 6.31: Results sensitivity analysis +25% unsatisfied demand costs




Figure 6.32: Cost comparison first and final solution sensitivity analysis +25% unsatisfied demand costs

The network resulting from the final solution (see Figure 6.33) does not differ from any of the other solutions very much as well. One of the diagonal lines has disappeared and some lines do not stop in the middle anymore, so other lines' frequencies can increase. This reduces waiting time in-vehicle costs.




Figure 6.33: Final solution sensitivity analysis +25% unsatisfied demand costs

The results from the sensitivity analysis show that the increase or decrease of the unsatisfied demand penalty does not have a significant impact on the outcome, if the value of the penalty stays within a $\pm 25\%$ boundary around the maximum observed passenger costs for a passenger within the original network. Although the chance that a solution with unsatisfied demand as a result is chosen as elite or parent is higher respectively lower, the results are comparable to the results of the base case scenario.

6.3 Limitations of the numerical experiments

There are a few limitations that need to be addressed before conclusions can be drawn from the results of the scenarios. The first limitation is the way waiting time had to be calculated. Because the model cannot produce timetables, the waiting time of passengers had to be estimated using the average waiting time under uniform arrivals. For higher frequencies, this probably approaches the truth as people will always keep some buffer between their planned arrival time at a station and the planned departure time of the line. However, for low frequencies the uniform arrival time method could lead to unrealistically high waiting times, which passengers can avoid by timing their arrival. On the other hand, people arriving by another public transport line may not have the option to time their arrival. Since the model aggregates all passengers, this cannot be derived from the results. An agent based model would be needed to calculate individual waiting times, which is beyond the scope of this thesis.

Another limitation is the exclusion of fixed operator costs. As the network is fictional, any number of vehicles could be available to run the network and all extra needed vehicles need to be bought. The extra costs of buying an extra vehicle would influence the results too much and the choice for the number of available vehicles is too arbitrary to include fixed costs in the model. Still, the operator costs are underestimated since they are incomplete.

The network is a limitation in itself as well, as it is not known whether other design choices would lead to other results. Scenarios were used to see the effects of variations, but the network always stayed the same. The Amsterdam case study could be used to see the effects of a different network, but the differences between both is so large that it is not possible pinpoint what difference in network resulted in what deviation in the results.

6.4 Conclusions of the numerical experiments

From the results of the scenarios, a few conclusions can be drawn with regard to optimal stopping patterns and frequencies. Concerning the stopping patterns, the results show that it is beneficial to serve as many stations as possible to reduce walking time. The saved in-vehicle costs for passengers that do not need to be at the skipped station do not outweigh the extra walking costs or even unsatisfied demand costs of the passengers that do use that station. An exception can be when a line can disappear altogether and most stations are served by another line as well.

The trend in all solutions is that frequencies of lines tend to be maximised within the available infrastructure capacity constraints. Lines that could stop at a bottleneck in the infrastructure, which is in this case the central station, regularly terminate before reaching the bottleneck so other lines' frequency can increase. Despite the extra operator costs, the overall effect of increasing frequencies was overall positive. This was partially a result of the way waiting costs and operator costs were calculated, as discussed in section 6.3.

The diagonal lines did not contribute much in reducing passenger costs and were often eliminated to save the operator costs. The high frequencies of the other lines made it easy to transfer between the lines, as the extra waiting time was relatively low. Therefore, it was more beneficial for passengers to take a detour that included a transfer than wait for a direct line. The amount of passengers using the diagonal was not very high altogether, which also explains why they often disappear in the final solution.

Passengers have a clear preference for lines with high frequencies, as this reduces their waiting time. This aspect of the total costs was found to be reduced the most in all results, undoubtedly because of the higher frequencies that were observed on almost all lines. This leads to the conclusion that it is better to have fewer lines with high frequencies and many, leading to more transfers and longer in-vehicle time, than many lines with low frequencies. The next chapter, in which SPAFOM is applied to the network of Amsterdam, has to reveal whether this is also the case in a much larger and more complex network.

7 Case study: Amsterdam

In this chapter, the Stopping Pattern and Frequency Optimization Model (SPAFOM) is applied to the train and metro network of Amsterdam. The morning peak for the year 2040 is chosen for the passenger demand. The form of the network is how it is expected to be in 2040, based on the network of 2017.

There are several reasons the network of Amsterdam is suitable for the model. First, the metro lines run for a large part in parallel to the train lines. Despite the technical differences between train and metro infrastructure, there are chances to optimise both services in an integral way. By choosing for 2040, any results that would require adaptations of infrastructure would be feasible, as there is enough time to build the adaptations. Second, there are already some ideas to integrate the train and metro network, or at least use the capacity of one for the other. A concrete plan is to extend the newly built North-South metro line from Zuid to Schiphol Airport or even further, using the existing train tracks and tunnel (Financieel Dagblad, 2016). Another idea has been to run regional stop trains through the metro tunnel between Amsterdam Centraal and Amstel.

One of the goals of the case study is to investigate whether an integrated approach to the Amsterdam public transport network could solve some of the issues that play a role there. The network could for instance take the form of Karlsruhe, where regional train lines share infrastructure with local tram lines (see Figure 7.1). In the case of Amsterdam, it would be the metro and train that share the same infrastructure.



Figure 7.1: Local (S-Bahn) train in Karlsruhe at shared tram and train station

Furthermore, Amsterdam Centraal suffers from a lack of capacity on multiple fronts. The platforms are too narrow for the large amount of passengers, especially during rush hours, but to solve this some of the tracks in the station need to make way for a widening of the platforms. This would reduce the capacity of the station for trains, as the disappearing tracks were used to let through going trains overtake turning trains.

The Amsterdam network is not only much bigger than the fictional test network described in chapter 6, it also provides an underlying network of buses as an alternative to train and metro within Amsterdam. An adapted version of the public transport model of the Amsterdam Metropolitan region, the Verkeerskundig Noordvleugelmodel (VENOM), is used as input. This input is described in section 7.1. Section 7.2 describes the optimization process, and the results are discussed in section 7.3. The feasibility of the results is discussed in section 7.4. The limitations of the model for Amsterdam, and therefore of the results, are discussed in section 7.5. Finally, conclusions are drawn in section 7.6.

7.1 Model input

The input of SPAFOM is for the largest part adapted from the regional traffic model of the metropolitan region of Amsterdam, VENOM. The network consists of the public transport network as it is modelled in the VENOM2016 public transport model, scenario 2040 High, with some adaptations. The biggest difference is that all bus lines that do not reach the study area have been eliminated, as well as all links outside the Amsterdam metropolitan area that do not carry a relevant transit line. The number of centroids to be considered in the assignment has been halved as well, all to decrease the run time of the model. This network will be called the original network throughout this chapter, to distinguish it from any network that is a result of SPAFOM.

Before the removal of unnecessary links, nodes, and transit lines, the run time of just the passenger assignment model was almost an hour, which means that each solution would take at least an hour to be calculated. After the removal, the run time was only four minutes, due to both the reduced amount of origin-destination pairs that needed to be routed through the network and the reduced number of routes available.

All centroids roughly outside of the catchment area of buses to and from Amsterdam have been aggregated in the nearest train station, as the travellers to or from these centroids have no alternative but the train to reach Amsterdam by public transport. Their origin and destination are fixed, so the movement of one of these or both will not create differences between different solutions in terms of passenger costs. They will always enter or leave the Amsterdam region by the same route, so the part of their journey outside the Amsterdam region is the same for each solution, whether the centroids are aggregated or not.

Not all train stations were aggregation candidates: only the train stations with a direct line to the study area and train stations of large cities have been selected. This way, all passengers enter and exit the study area using the right transit line, which is all that matters for the model. The journey outside the study area is irrelevant, because it is the same for each solution due to the lack of alternatives to the train.

The transit lines that are taken into account in the model are shown in Figure 7.2. All train, metro and bus lines in the study area are shown. Trams are not taken into account, as they were not modelled in VENOM for 2040 and could not be added in the available time. Due to the opening of the North-South metro line, the tram routes will change significantly, but this was not yet modelled.

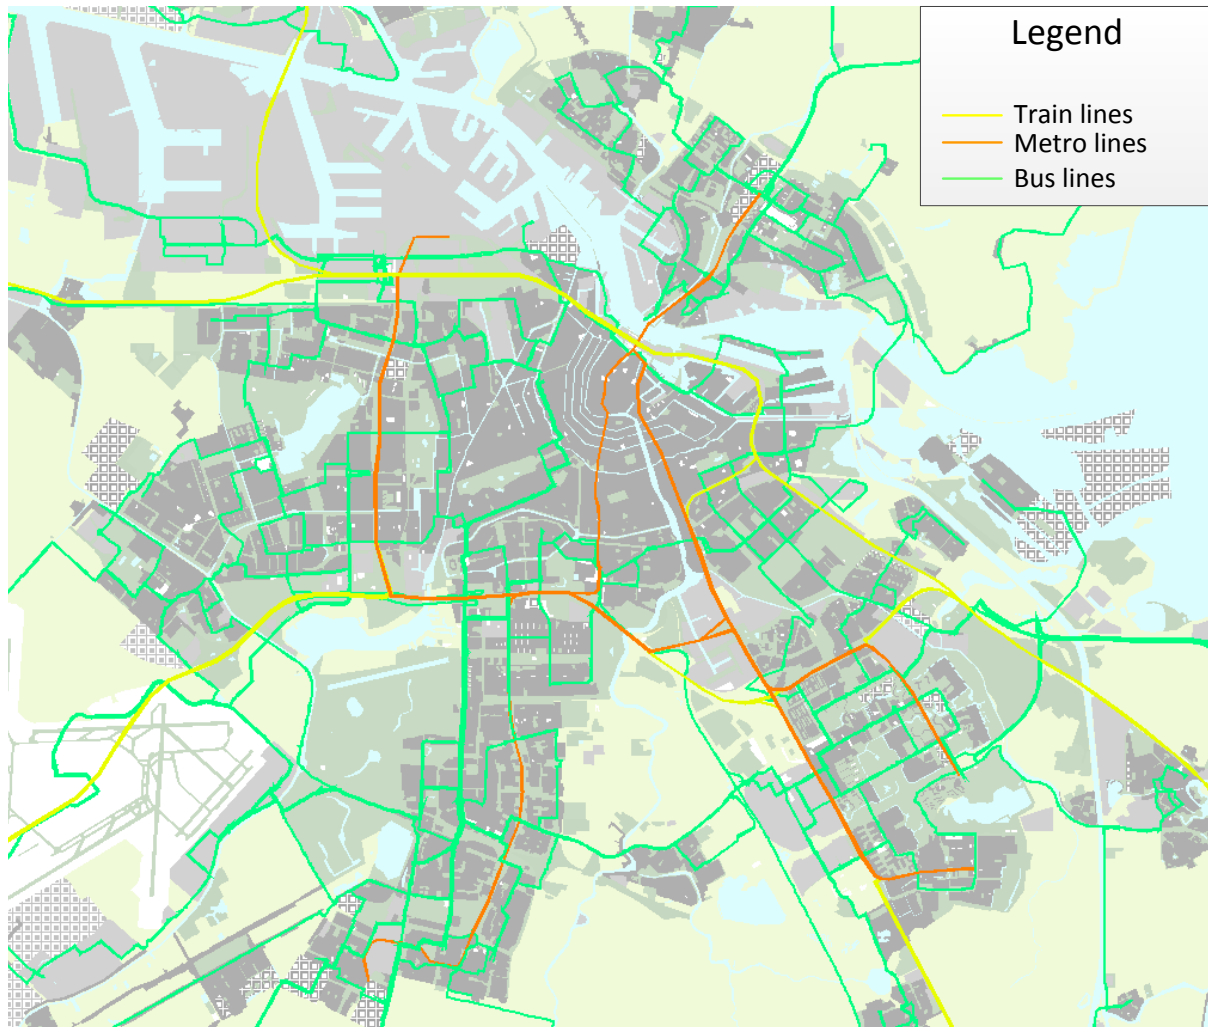


Figure 7.2: All transit lines in study area

The facts and figures of the Amsterdam case study are shown in Table 7.1.

Table 7.1: Facts and figures Amsterdam case study

Number of stations	453
Number of lines	108 (bidirectional)
Number of stations subject to design	48
Number of lines subject to design	21
Number of decision variables	193
Number of links	16.440
Number of origins / destinations	1.860
Passenger demand	481.875 passenger trips

The rail infrastructure and the stations that are part of the optimization are shown in Figure 7.3. All depicted stations are part of the optimization, except those in italic. The latter stations are the next large nodes in the network, where lines from multiple directions come together. All metro and train infrastructure currently expected to be used in 2040 within the region of Amsterdam is part of the optimization network. This includes the new North-South metro line, but excludes the Amstelveen line from station Amsterdam Zuid towards Amstelveen, as the latter will be reconstructed into a tram line (*Amstelveenlijn | Waarom vernieuwen we de Amstelveenlijn?*, 2015). The North-South metro line is extended towards Hoofddorp, a possibility that has previously been marked as useful, but not necessary (Gemeente Haarlemmermeer, 2005). Whether this extension is realised using dedicated infrastructure or the existing train tunnel, which is still a societal discussion, is not part of this thesis and not relevant for the model, as infrastructure capacity is not taken into account.

For the origin and destination relations, the high scenario for 2040 developed in the future exploration prosperity and living environment (Toekomstverkenning Welvaart en Leefomgeving, WLO) is chosen as origin and destination (OD) matrix. This is a scenario developed by the Dutch planning bureau for the living environment (PLO) and the central planning bureau (CPB) in 2015, based on fast economic growth and an accompanying large growth in mobility. Studies in 2016 by Provinces and the central government, gathered by Vervoerregio Amsterdam (Transport region Amsterdam), conclude that this high scenario corresponds to the regular prognoses and that it is thus actually quite low to be “high”.

The public transport lines have been adapted from VENOM as well. A map with all lines subject to the optimization is shown in Figure 7.4. Most of them have been copied as they were. If the same line was in the model twice or more, all but one version were erased and their frequencies were aggregated. Stop train services via Amsterdam Central station and Amsterdam Amstel station via Muiderpoort have been cut at Amsterdam Centraal, if needed, and routed through the metro tunnel between Centraal and Amstel stations. The reason for this is that this way the effect of having regional services towards the centre can be studied. Intercity services via this route were not adapted and follow their current route via Amsterdam Muiderpoort station. The bus lines, other than those not within the study area, were not adapted either.

Figure 7.5 and Figure 7.6 show the number of passengers each zone within the study area produces and attracts respectively in the observed hour. The origin-destination matrix does not change during the optimization process, as the same matrix is input for all solutions. In other words, it is assumed that people do not change their travel behaviour when transit lines change, other than picking another route if that is beneficial for them.

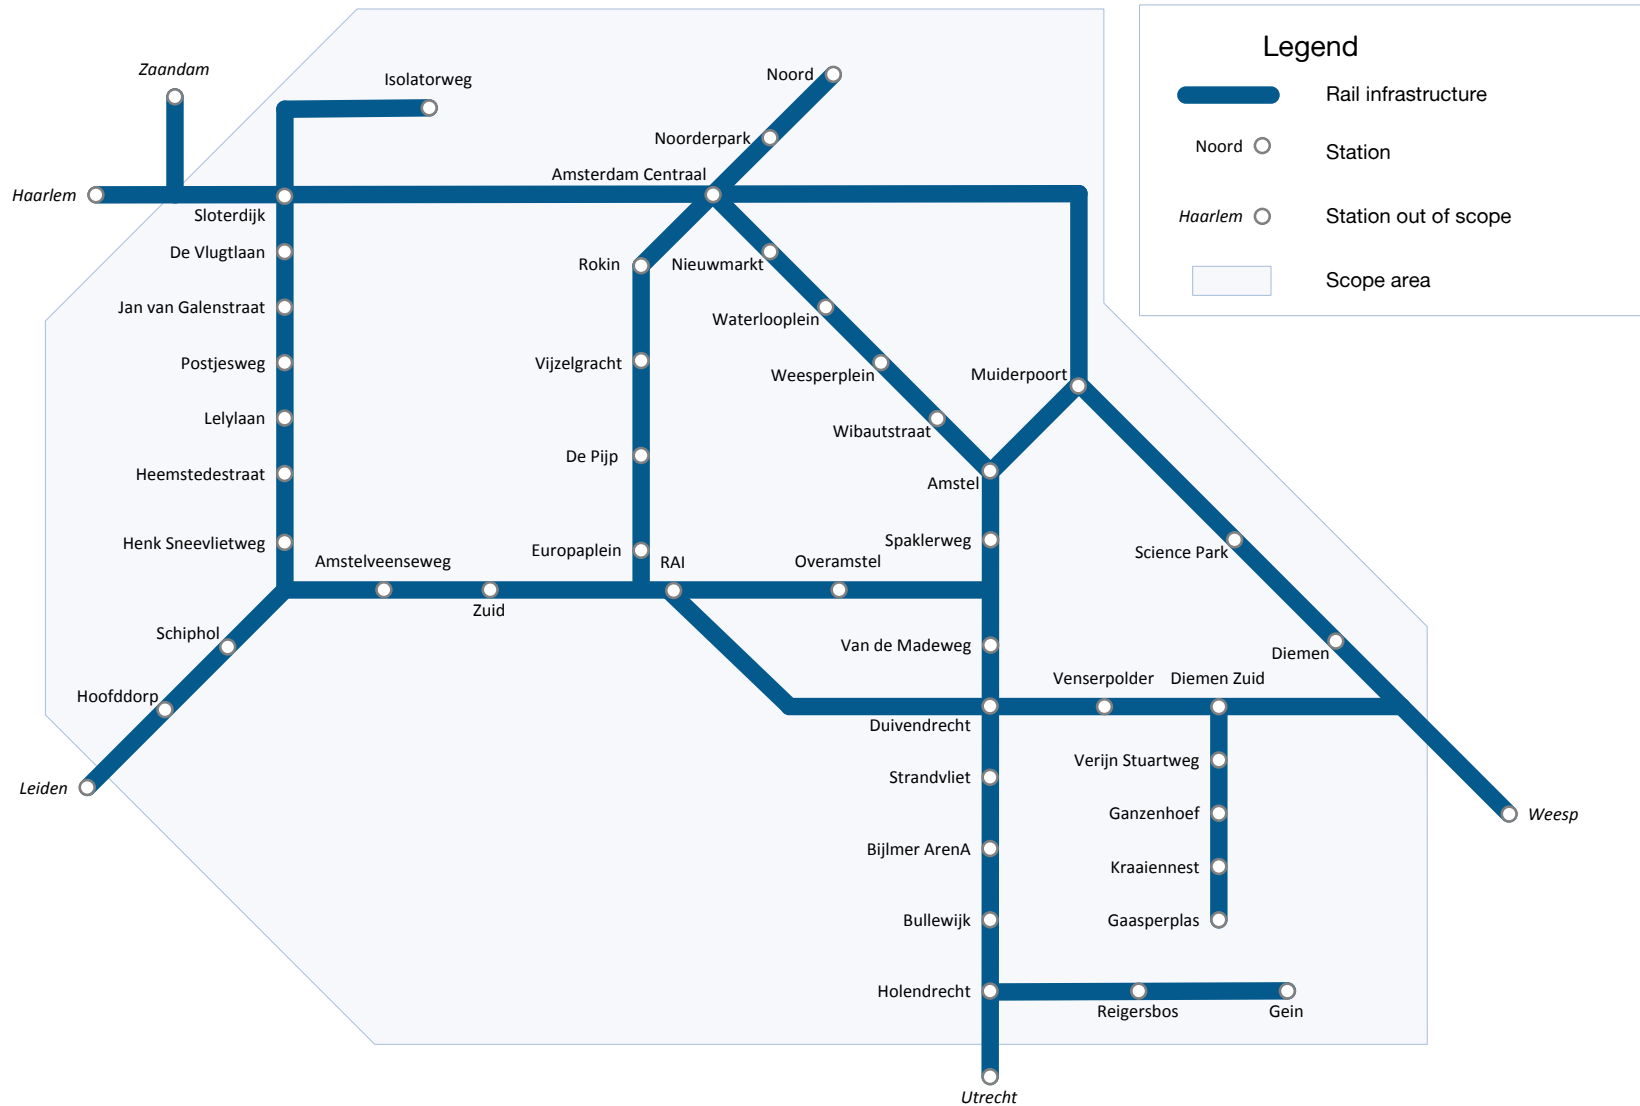


Figure 7.3: Available links and nodes for Amsterdam case study

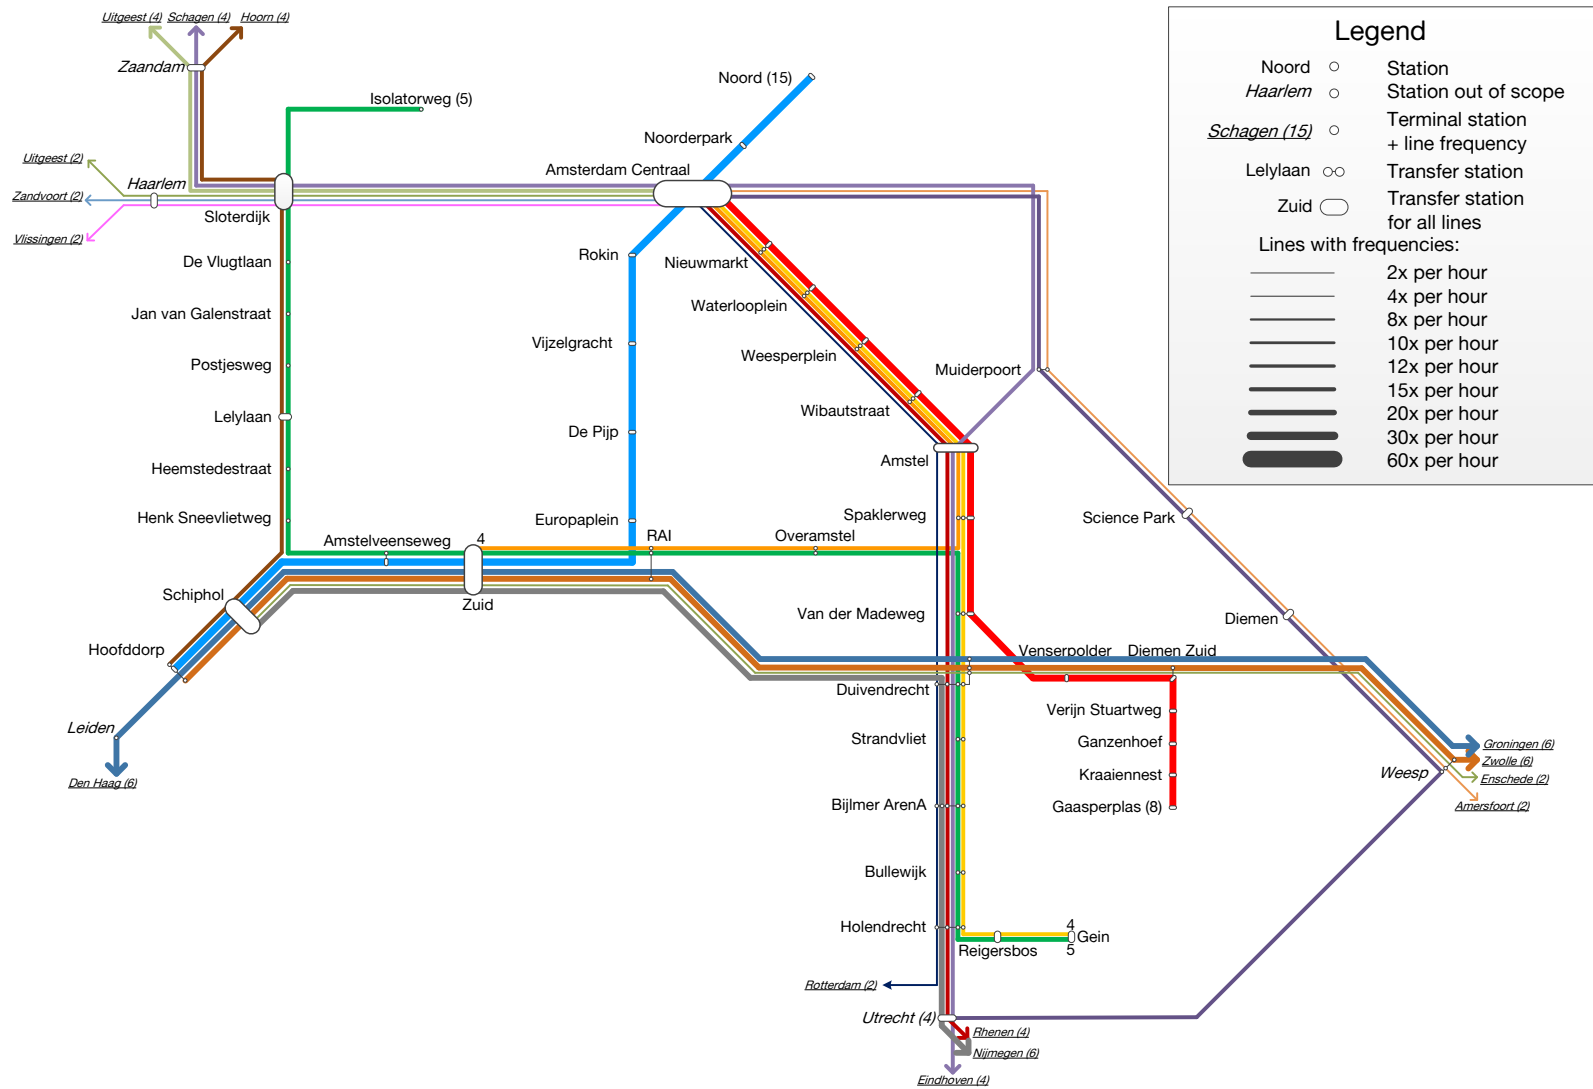


Figure 7.4: Current stopping pattern and frequencies of lines to be optimized




Figure 7.5: Number of passengers produced per zone within study area for morning peak hour, including lines to optimize




Figure 7.6: Number of passengers attracted per zone within study area for morning peak hour, including lines to optimize

The passengers come from all zones within the study area, with only a few zones producing significantly more than the others. These zones are Zeeburgereiland, a part of Amsterdam that is being developed as this thesis is written (*Zeeburgereiland: nieuwe stadswijk aan het water - Gemeente Amsterdam*, 2017) and Schiphol Airport. Furthermore, there are several zones within and near the centre of Amsterdam that produce a lot of passengers as well.

The map with the number of attracted passengers looks a bit different. It can be seen that the destinations of travellers are more concentrated than their origins. These areas are concentrated around the centre of Amsterdam and some train and metro stations. The areas that attract many passengers are characterised by one common attribute: there is a lot of employment. This makes sense, as in a regular morning peak a large majority of passengers travels to work or education.

7.2 Optimization process

The optimization process uses the genetic algorithm, which was described in chapter 5. It starts with a mix of designed and random solutions, as the numerical experiments showed that this sped up the process. The possibility of bias towards designed solutions is countered by having more elites than designed solutions, so some random solutions are saved for later generations to ensure diversity throughout the process. The designed solutions consist of:

- A solution in which the stopping patterns and frequencies are the same as in the 2040 High scenario of VENOM.
- A solution in which each line stops at each station it passes and has a frequency of one run per hour.
- A solution without metro, but with a doubled frequency for all train lines. The stopping pattern of all trains is the same as in the previous solution.
- A solution without trains that were marked in VENOM as stop trains. All other lines run as in the first solution.
- A solution in which the stop trains' stopping patterns outside the study area is as in the first solution and within the study area the trains only stop at the first and last station that also has an intercity service stopping there.

The random set of solutions is repaired, so each solution consists of pairs of symmetrical lines regarding both stopping pattern and frequency. Infrastructure capacity is not a constraint in this case, as it proved to be very limiting in the numerical experiments. For the Amsterdam case, the goal is to find out in what direction the Amsterdam public transport network should develop towards 2040. Between now and then, there is plenty of opportunity to adapt the infrastructure if needed, making the current capacity of the network less relevant. Therefore, capacity is not a hard constraint, but it will be reflected on later.

All of the solutions, together the first generation, are then evaluated using the static passenger assignment model. The model is adapted to represent each of the solutions, after which the model is run. Each time a solution is run by the model, the output is collected, which consists of passenger costs (walking costs, waiting costs, in-vehicle costs, transfer costs and unsatisfied demand costs) and variable operator costs, using the input from Table 7.2.

Table 7.2: Input Amsterdam case study

Data type	Value	Source/method
Network	Original network	-
Genetic algorithm		
Population size	32	Rule of thumb ⁴
Number of generations	As many as possible	-
Crossover rate	1,0	Elites make not doing crossover unnecessary
Number of elites	16	Best practice
Mutation rate	0,01	Best practice
Other		
Dwell time train	2 minutes	Practice / expert judgement
Dwell time metro	1 minute	Practice / expert judgement
Headway train	4 minutes	ProRail (2017)
Headway metro	2 minutes	LeighFisher (2014)
Dwell time other	0,5 minutes	Practice / expert judgement
Capacity train	1.320 passengers	12 cars, NS (2017)
Capacity metro	960 passengers	6 cars, <i>Metro</i> GVB (2017)
Capacity bus	91 passengers	Citaro standard, (<i>Bus</i> GVB, 2017)
Frequencies	0, 1, 2, 3, 4, 5, 6, 8, 10, 12, 15, 20, 30, 60 runs per hour	Possibilities for cyclic schedules
Parameter walking time	1.0	VENOM
Parameter waiting time	0,5	VENOM
Parameter in-vehicle time	1,0-2,0	(Cats <i>et al.</i> , 2016)
Parameter transfer	5,0 minutes	VENOM, Ortúzar and Willumsen (2011)
Value of time	€8,75/hour	Kouwenhoven, de Jong and Koster (2014)
Unsatisfied demand penalty	80,0 minutes	Original network
Operator costs train		
Fixed costs	0.0	Not taken into account
Variable costs per hour	€129,63	AT Osborne
Variable costs per km	€12,00	CROW
Operator costs bus		
Fixed costs	0,0	Not taken into account
Variable costs per hour	€49,85	AT Osborne
Variable costs per km	€0,65	CROW

⁴ $\frac{(\text{degrees of freedom})^2}{1000}$, rounded up to 32 for convenience

Fixed operator costs are not part of the evaluation, because they could not be reasonably calculated and the inaccurate result would influence the evaluation too much. As the costs of purchasing a vehicle lies in the same order of magnitude as the total passenger costs, the model would probably converge to a solution in which the available fleet is deployed as efficiently as possible. This is very interesting, but not the goal of this thesis. Therefore, an unlimited fleet size is assumed. This assumption will be reflected on in the results.

The sum of all costs is the score for each solution. All costs per solution are saved, so each solution does not need to run more than once. The 16 best solutions, based on their score, are then selected for the next generation of results. The other 16 solutions are formed by crossing solutions from the previous generation based on their score. A higher score, so in this case lower total costs, means a solution has a higher chance of being crossed with another solution to form a new solution.

Each new solution is then run by the model, as described earlier, for as long as the limited time to write this thesis allowed. The algorithm was stopped after 31 days, after the 689th generation was processed. In this time, it has run slightly over 11.000 unique solutions, so it took around 4 minutes per solution to run the solution and calculate the results. As there are $6,3e^{57}$ possible solutions, only a very small portion of it has actually been evaluated.

The model was run on one laptop, due to the software license that was needed to run the passenger assignment model. This made it more difficult to distribute computations over more than one processor. The used software did not allow this easily either, so due to time constraints it was chosen to accept the longer calculation time. The used laptop had 16 GB of ram and a dual core 2,8 GHz Intel i7-4810MQ processor.

7.3 Results

This section presents the results from the Amsterdam case study. It successively discusses the costs, the performance of the services, the resulting network, the performance of the network, transfers, and walking.

7.3.1 Costs

Figure 7.7 shows the results from the Amsterdam case study per generation, broken down into the different cost components. It is clear that a large part of the improvements take place in the beginning of the process. After that, still better solutions are found, but the improvements are not very large. The final solution is found around the 620th generation and remains the best for another 70 generations.




Figure 7.7: Results Amsterdam case study

The costs of the first and final solution are shown in Figure 7.8, where they are split into the different components. The in-vehicle costs are the largest component in both solutions, followed by the walking costs. Waiting costs and operator costs contribute about the same to the total costs, while transfer costs play only a marginal role. There are no unsatisfied demand costs in both solutions, which means that all zones are served. This was to be expected, as both the public transport and walking networks of Amsterdam are very extensive, which makes it possible for almost all origins or destinations to be reached via multiple stations.




Figure 7.8: Cost components comparison current situation, first solution and final solution Amsterdam case study

The results show similarities with the numerical experiments in chapter 6. In those experiments, the waiting costs were the main contributor to the decrease in total costs, followed by in-vehicle costs. In the Amsterdam case, this is the case as well. The walking costs have decreased as well, and will be discussed at the end of this section, as knowledge about the new network is necessary to explain this.

The variable operator costs have increased significantly. In fact, they are double what they are in the current situation. The reason for this will become clear in the next section, where the network is shown. The same can be said for the transfer costs, which have risen slightly. Still, the decrease in other passenger costs more than make up for this.

It can be seen that for especially the waiting costs, the best solution of the first generation shows a large decrease in costs already. This is related to a higher frequency, which becomes clear from the lower in-vehicle costs and higher operator costs as well.

The proportions of the costs have remained more or less the same, opposed to what happened in the numerical experiments in the previous chapter. The reason for this is the network, which is more extensive in the Amsterdam case study than in the numerical experiments. This provided the opportunities to decrease walking costs, for instance, which was not possible in the numerical experiments.

In Figure 7.9, it can be seen that the best solution in the first generation is already better than the current situation. The difference is about 1%. However, the largest gains are achieved in the model, as the optimization reduces the total costs over 10% compared to the first generation, or almost 12% compared to the current situation. The model is thus capable of achieving significantly better solutions for an existing public transport network.




Figure 7.9: Total costs comparison current situation, first solution and final solution Amsterdam case study

7.3.2 Network

The final solution for the public transport network of Amsterdam is shown in Figure 7.10. It can be seen that quite some changes have taken place. They will not all be discussed, but the most important and impactful changes are noted. For a complete overview of the lines, the reader is referred to Appendix B, which contains the number of stops and the frequencies per line for both the original situation and the final solution.

The first striking features are the frequencies of particularly the metro lines, which have been increased to up to 60 times per hour. This is the same behaviour as was observed in the numerical experiment without infrastructure limitations, as waiting time is an important component of passenger costs. The relaxation of capacity constraints made these high frequencies possible, as they would not fit in the current infrastructure. It is important to keep this in mind in the interpretation of the result. The combination of the major effect of waiting time on the solution and the unlimited infrastructure capacity leads to unrealistically high frequencies, which need to be seen more as a guidance. The model did not take fleet size constraints into account, i.e. an unlimited amount of available vehicles is assumed. This is not the case in the real situation, where for this kind of frequencies many extra vehicles need to be acquired by operators, with the corresponding costs as a consequence.

The same happened to some lines that were originally regional stop services, though to a lesser extent. The lines that run mainly outside the study area, for instance to Eindhoven or Groningen, do not have this frequency increase. This is most probably caused by the operator costs, which are dependent on both run time and run distance. These lines are very expensive compared to for instance the metro and as the model focuses on the Amsterdam region, there are not many passengers that benefit from an increased frequency. Most passengers in the model have a shorter line as alternative, so frequencies of longer lines were sometimes not increased.

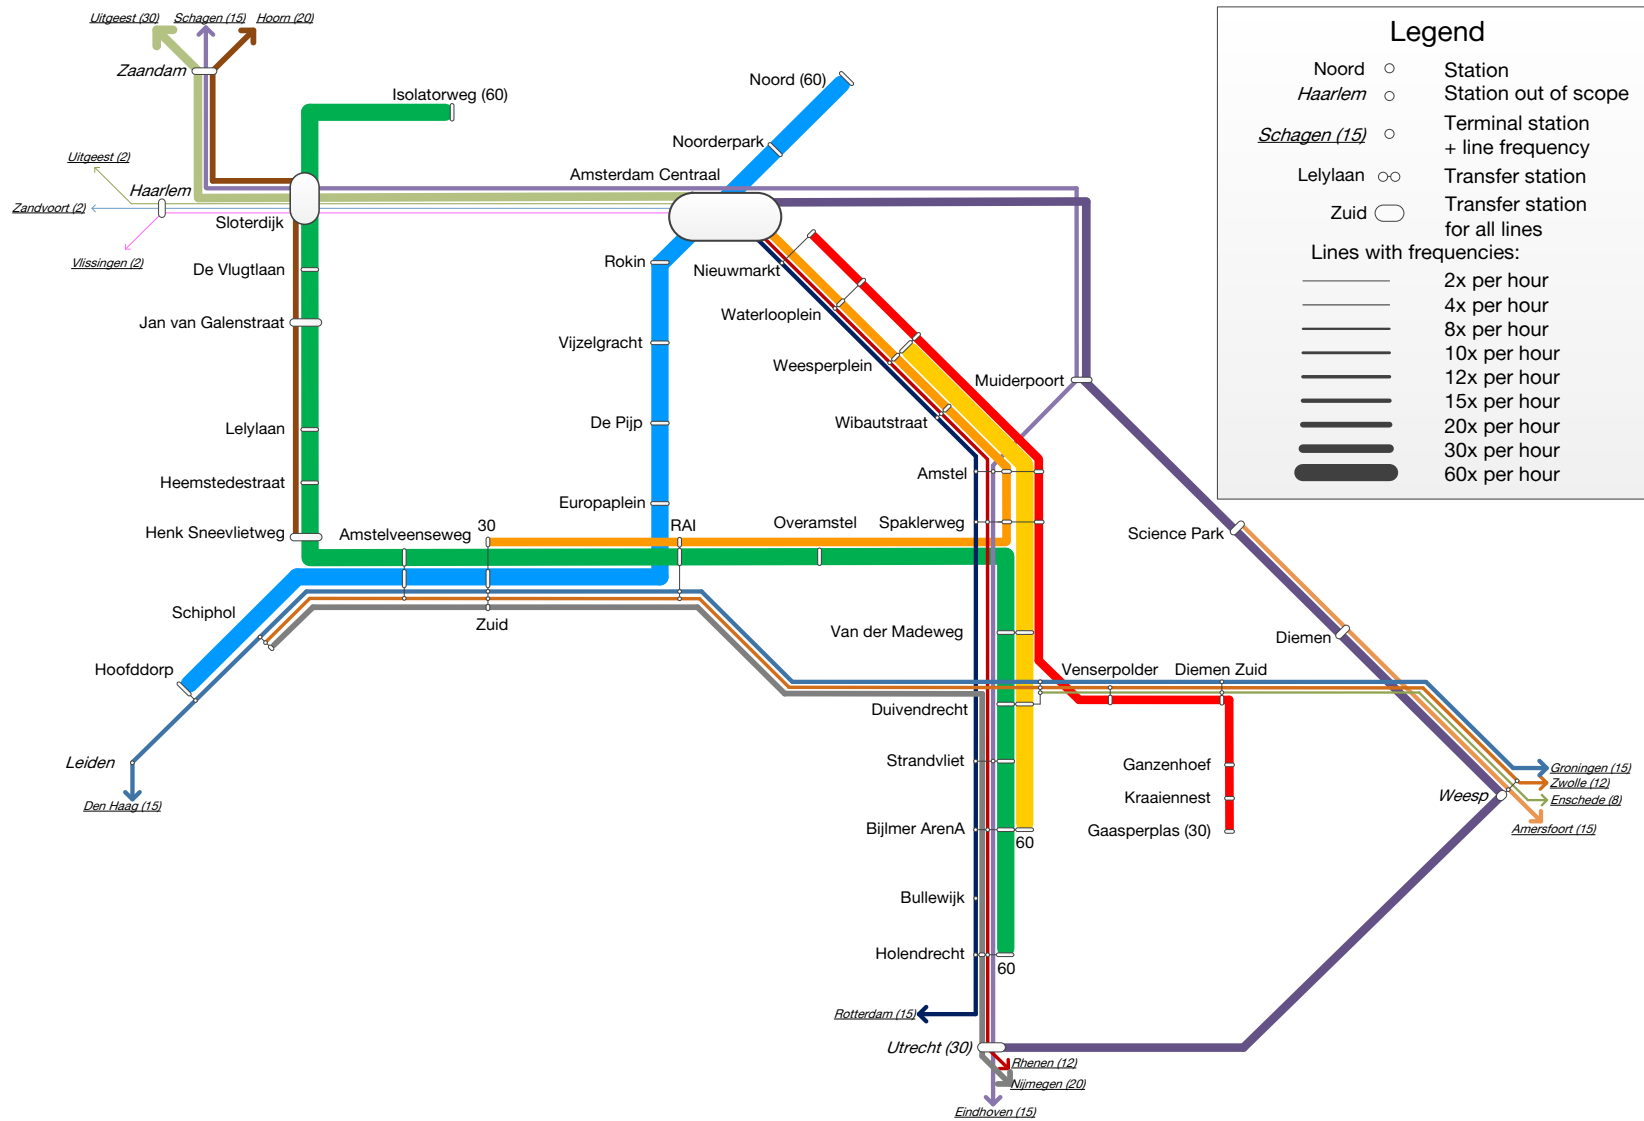


Figure 7.10: Network final solution Amsterdam case study

Another striking feature is the shortening of some lines. Again, the metro lines have changed the most. Allowing regional lines to enter the metro tunnel between Amsterdam Centraal and Amstel brought about quite some changes in the stopping patterns of the affected lines (see Figure 7.11).




Figure 7.11: Magnification Amsterdam Centraal - Amsterdam Amstel: some lines do not stop at Amsterdam Centraal anymore

Two of the metro lines do not continue until Amsterdam Centraal anymore, which has two main reasons. First, their function is partially taken over by the extra lines through the tunnel, which stop at multiple stations in the tunnel. This transfer-free connection to the city centre makes it unnecessary for all metro lines to go all the way to the central station. The high frequencies of some of the lines result in quite some saved operator costs when cutting a line off.

The second reason is that the zones around the central station have a limited attraction and production value of itself (see Figure 7.12), so there are not many passengers that absolutely have to be in the area. For most passengers that go to Amsterdam Centraal, the station acts as a transfer station. If this transfer can take place somewhere else, or another route can be used between origin and destination, there is no need to go to the station at all.

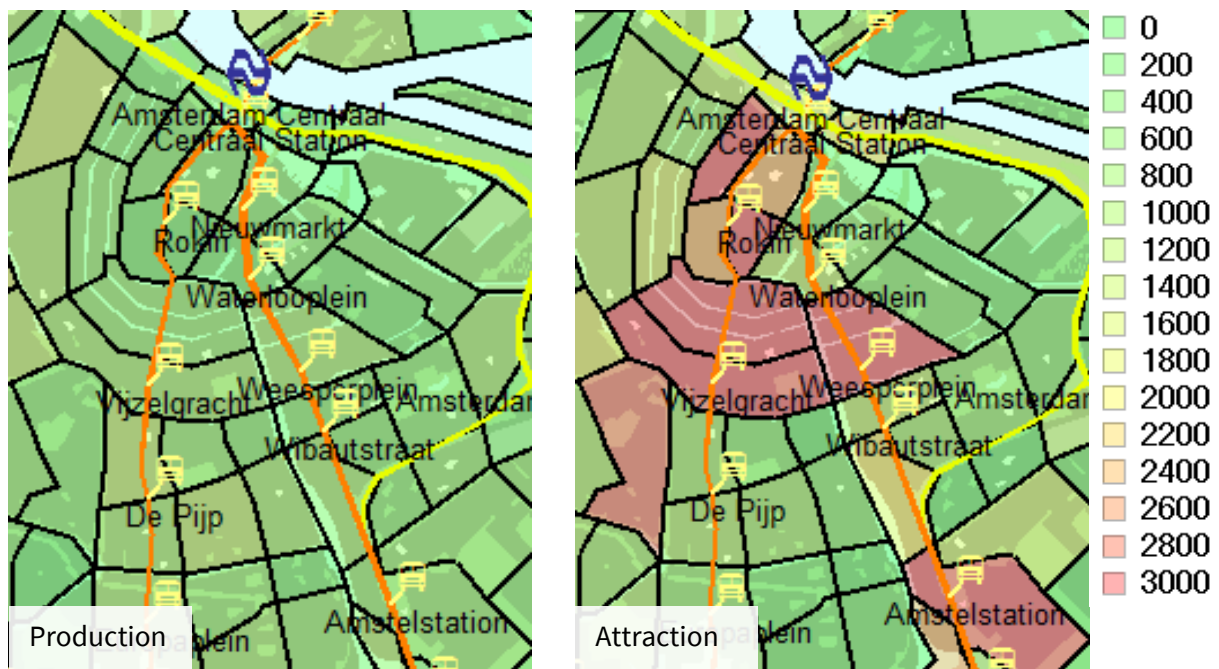


Figure 7.12: Magnification production (L) and attraction Amsterdam Centraal - Amsterdam Amstel for morning peak, with a concentration of attraction in the city centre

This is also the case for the line from Eindhoven to Schagen, which runs from the south to the north-west via Bijlmer ArenA, Amsterdam Centraal and Sloterdijk. The line does not stop in Amsterdam Centraal anymore, as the other stations where the line does stop provide all the relevant transfers. The high frequencies of local lines make transferring a lot less costly, as there is barely any extra waiting time, leaving only the transfer penalty.

Some metro lines are cut short in other places as well, for instance in the south-east towards Gein station. In the original situation, both metro lines (the boldest lines in Figure 7.13, running straight southwards) had a branch towards Gein station, running east from Holendrecht station. This branch has disappeared for both lines, with one line terminating at Holendrecht and the other line at Bijlmer ArenA. However, where in the previous case the functions of both lines were taken over directly by other lines, the area around Gein now is deprived from rail-bound public transport. Instead, it is now only served by the underlying network, which is in this case a bus line through the area. This bus line, which runs 6 times per hour, provides a good connection to both Holendrecht and Gaasperplas stations. Therefore, the operational benefits of cutting short the metro lines outweigh the higher passenger costs for the passengers to and from the area, a number that is already relatively low besides. This example shows that taking the underlying network into account can lead to different choices for the network subject to design and that neglecting this underlying network makes the analysis incomplete.

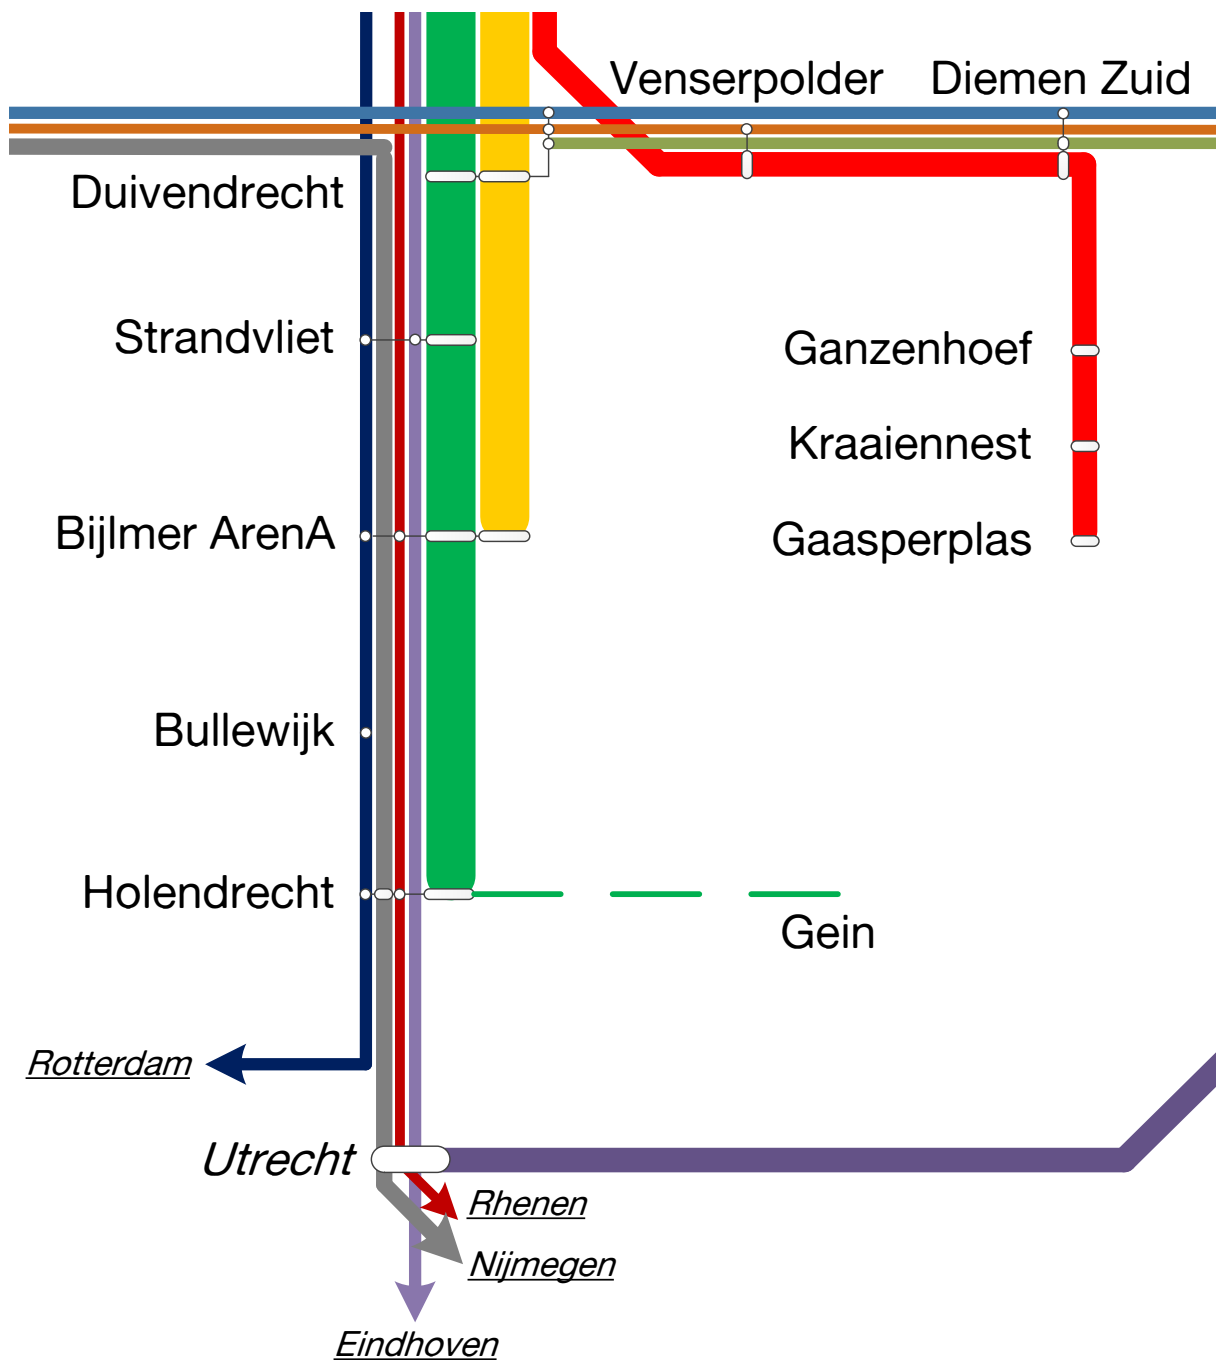


Figure 7.13: Magnification Gein, where the metro lines do no reach the Gein area anymore

Another station that is not served anymore is Postjesweg, between Jan van Galenstraat and Lelylaan. This is a minor station that is also served by buses, which could explain why it is skipped by metro and train lines. The same goes for Verijn Stuartweg between Ganzenhoef and Diemen Zuid. Overall, the trend is that lines serve as many stations as possible, unless the same route is already provided by another line. The best example of this is the metro line between Weesperplein and Bijlmer ArenA, which skips half of the stations along its route and is shortened at both ends as well. Other metro lines and the regional train lines through the metro tunnel between Amsterdam Centraal and Amstel provide the same service and take over the function of the metro line.

7.3.2.1 North-South metro line at Schiphol Airport

As mentioned earlier, the higher frequencies make the threshold for a transfer lower. This is also shown by the situation around station Schiphol Airport. Although it is one of the larger zones in terms of production and attraction of passengers, it is not served by the North-South metro line. In fact, it is the only station on that line that is skipped. Apparently, the solution with a stop at Schiphol was not assessed by the model, as it could only assess a limited amount of solutions.

Passengers to the city centre now have to board a train and change at for instance station Zuid, as is the case in the original situation as well. To check whether this really is beneficial for the overall system, or if the same solution with a stop at Schiphol was never evaluated, the solution is adapted to include Schiphol in the stopping pattern of the North-South metro line and run. The results are shown in Figure 7.14.




Figure 7.14: Comparison of costs final and adjusted final solution

It can be seen that all costs decrease or stay constant. The difference in operator costs is not noticeable. As expected, the transfer and waiting costs decrease relatively the most. The results make clear that the algorithm does not guarantee to find the optimal solution. Including Schiphol in the North-South line is one of the improvements that make the final solution even better, as it connects several zones that produce and attract relatively many passengers.

Therefore, the following results are from the solution with a stopping North-South line at Schiphol Airport, as this is the best solution that was found. The information provided previously is the same for both solutions, except for a small change in costs and the extra stop at Schiphol. Schiphol becomes a transfer station for all passing lines with this extra stop.

7.3.3 Service performance

A quick overview of the differences in vehicle kilometres, passenger kilometres and passenger travel minutes is provided in Table 7.3. Especially the amount of vehicle kilometres has increased significantly, which is the reason that the variable operator costs have increased so much.

Table 7.3: Comparison of indicators between original situation and final solution

	Original situation	Final solution	Difference (%)
Vehicle kilometres (designed lines)	11.852	44.085	+272
Vehicle kilometres (all lines)	30.103	61.868	+106
Passenger kilometres (designed lines)	5.676.204	5.904.760	+4
Passenger kilometres (all lines)	11.086.891	10.778.861	-3
Passenger hours (designed lines)	93.204	96.507	+4
Passenger hours (all lines)	179.542	171.590	-4
Kilometres per passenger (all lines)	11,8	12,3	-3
Minutes per passenger (all lines)	22,4	21,4	-4

The total vehicle kilometres increase more than the passenger kilometres or hours decrease. The amount of passenger kilometres even increases for the designed lines, which indicates that the increased attractiveness of these lines attract more passengers. However, as the overall amount of passenger kilometres decreases, it can be concluded that there is a shift from the fixed bus network towards the lines subject to design. As was observed in the cost analysis, waiting and walking costs were important factors as well, which explains the imbalance between the increase of vehicle kilometres and passenger kilometres and time decrease. So although pure trip length and duration per passenger decreased as well, the complete benefits for passengers (less walking, waiting, crowding) is not caught in these numbers.

The information above is split for the different modes and visualised in Figure 7.15, Figure 7.16 and Figure 7.17 for all lines. The figures show that the supply of services increases, while passenger are provided with faster and shorter routes between their origin and their destination.

The metro system changes relatively the most in terms of vehicle kilometres, which results in more usage by travellers. While all other services are used less, both in time and distance travelled, the metro system is used more. An attractive metro system can thus relieve other systems within Amsterdam.




Figure 7.15: Vehicle kilometres by mode in original situation and final solution for all lines




Figure 7.16: Passenger kilometres by mode in original situation and final solution for all lines




Figure 7.17: Passenger hours by mode in original situation and final solution for all lines

The modal split in terms of passenger kilometres and passenger hours for both the original situation and the final solution is shown in Figure 7.18 and shows that the train is the main mode in all cases. This is explained by the trip length for train passengers, which is much longer than the trip lengths of local services.




Figure 7.18: Modal split original situation and final solution in terms of passenger kilometres and passenger hours

It can be seen that the largest changes take place for the metro, which is used more in terms of both kilometres and hours in the final solution than in the original situation. This is mainly at the expense of the other local services, which decrease relatively the most. Although the share of kilometres travelled by metro doubles, the share of time spent in a metro vehicle only increases by a third. Therefore, it can be concluded that in the final solution the metro offers on average a faster journey for the travellers using the metro.

7.3.4 Network performance

Figure 7.19 shows the difference in loads on the links in the network between the original situation and the final solution of the model. The largest difference is seen between the lines to the city centre, where passengers have shifted towards the eastern metro tunnel and the North-South metro line, mainly in the southern direction. These highly frequent lines are more attractive than other lines.




Figure 7.19: Difference in passenger link loads between the final solution and the original situation

Zooming in on the city centre, it can be seen that both the North-South metro line and the lines through the tunnel towards Amstel attract more passengers. This is at the expense of the train lines and some bus lines, which are used less.




Figure 7.20: Passenger load comparison Amsterdam centre

This also explains the difference in passengers between Zuid and Duivendrecht, as passengers take a detour to use the North-South line. Another apparent difference is the use of the network in the south-eastern part of Amsterdam, where the shortening of some metro lines has caused a decreased use. Their function is taken over by the bus, which can be seen in Figure 7.21.




Figure 7.21: Passenger load comparison Gein area

There were also some changes on the western side of the network, where travellers from Leiden (south west of Amsterdam) travel via the North-South metro line. This is the reason for the shortening of the train between Hoofddorp and Hoorn, which was the only service to provide a direct connection between Hoofddorp and Sloterdijk. The service was cut at station Henk Sneevlietweg, so it does not reach Schiphol anymore. There is also a shift from Amsterdam Zuid to Amstelveenseweg, caused by passengers to and from VU Medical Centre, for whom that station is slightly closer.




Figure 7.22: Passenger load comparison Amsterdam south west

All in all, most passengers have stuck to their original routes, as can be seen in Figure 7.23. The only big difference is the more extensive use of the North-South metro line, especially in the south. This can be directly related to the extension of the metro line to Schiphol and Hoofddorp.

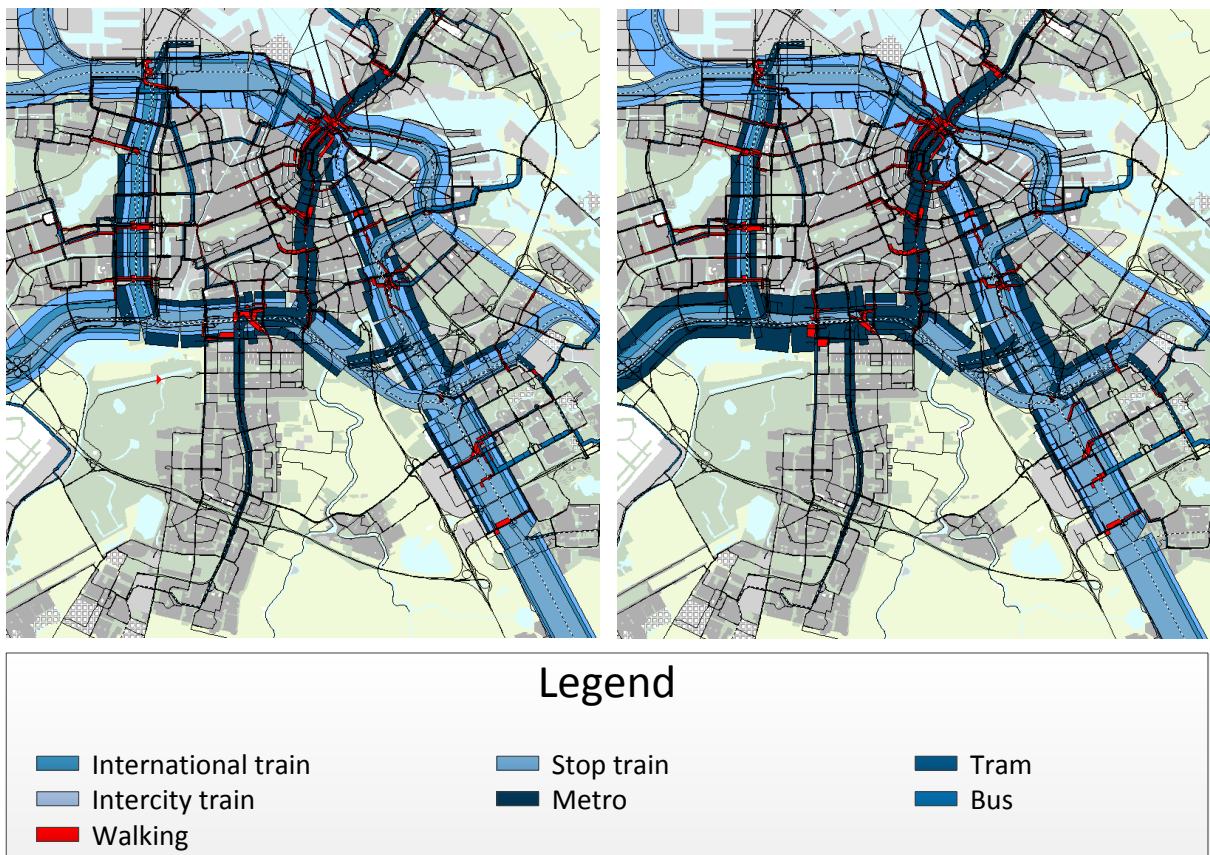


Figure 7.23: Comparison of passenger flows in the original situation (left) and final solution, which shows a shift from train to metro

7.3.5 Transfers

Another consequence of the increased use of the extended North-South metro line is a distribution of transfers over several stations. Appendix C contains the number of transfers per station in both the original situation and the final solution, which is also shown visually in Figure 7.24. The top 10 most changes in absolute numbers is provided in Table 7.4 and shows that a large part of the changes is caused by the change of the North-South metro line.

The other part of the top 10 is caused by changes in the stopping patterns of the eastern metro lines, where the transfers are divided over multiple stations. In the original situation, the transfers were concentrated in Amsterdam Amstel and Bijlmer ArenA, and while the amount of transfers at the former almost triples, the latter loses a lot of transfers. Some stations that are only served by metro lines in the original situation are served by regional lines in the final situation as well, making them suitable for transfers between local and regional lines. This is inseparable from the rerouting of regional train lines through the metro tunnel between Amstel and Amsterdam Centraal, as it is now favourable to use that tunnel to go to the city centre instead of travelling via Amsterdam Centraal.

Table 7.4: Top 10 changes in number of transfers

Station	Number of transfers in original situation	Number of transfers in final solution	Change (%)
Station Zuid	21528	14264	-34%
Amstelveenseweg	0	5358	-
Hoofddorp	0	5262	-
Amstelstation	2644	7454	+182%
Station RAI	854	4851	+468%
Strandvliet	0	3425	-
Spaklerweg	1	3187	+239492%
Schiphol	6089	3543	-42%
Henk Sneevlietweg	0	2370	-
Station Bijlmer ArenaA	2509	312	-88%

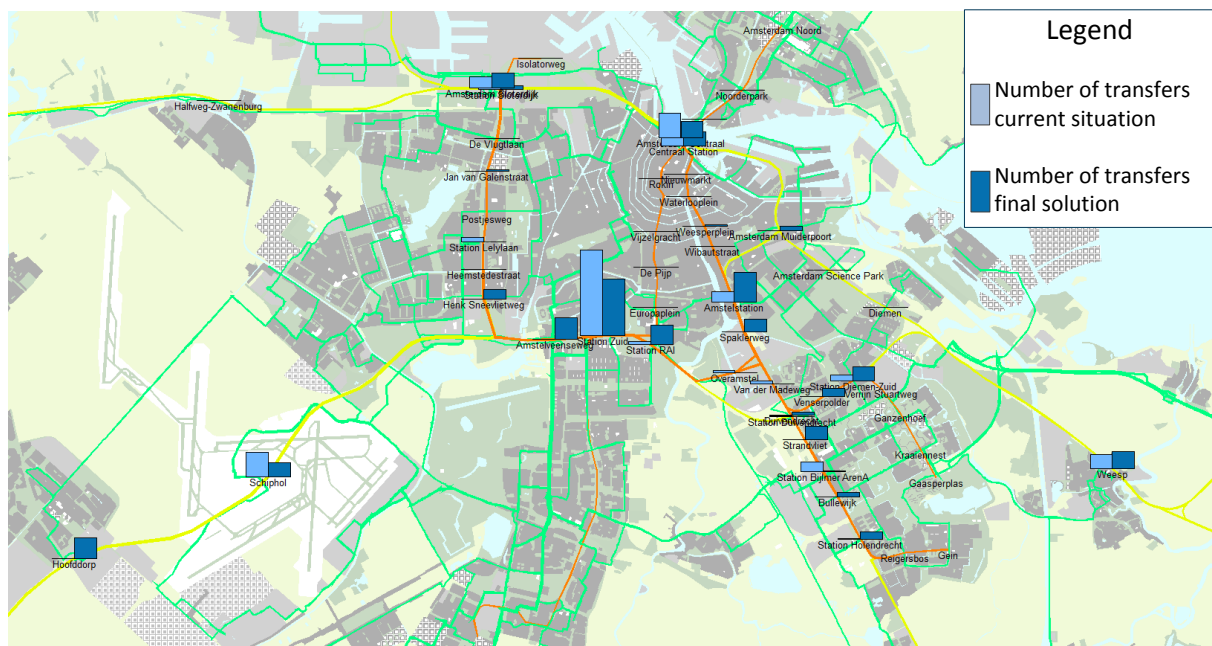


Figure 7.24: Number of transfers per station in original situation and final solution. Amsterdam Centraal and Sloterdijk are modelled separately for train and metro and are summed in the tables.

Especially Amsterdam Zuid is a gate for all passengers using the North-South metro line, as it facilitates by far the most transfers of all stations. In the final solution, its function is partially taken over by Hoofddorp and Amstelveenseweg, as these stations provide a transfer to the North-South line as well (see also Figure 7.25). Extending the North-South metro line can thus alleviate the pressure on Amsterdam Zuid, provided that the transfer time at those stations is favourable for passengers.

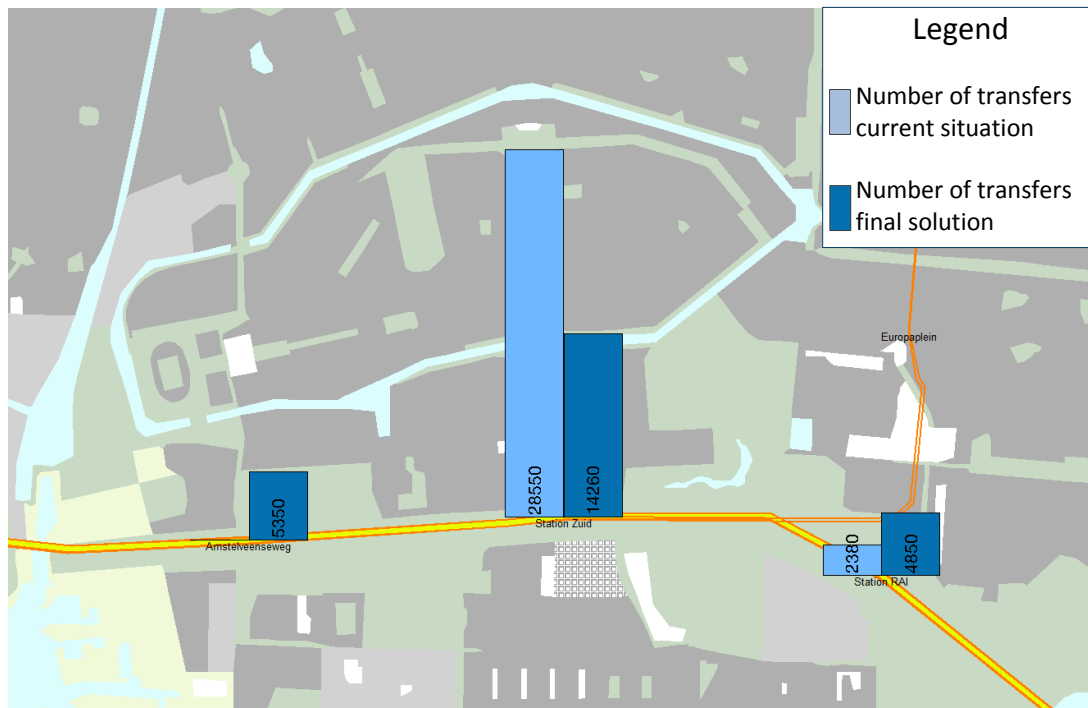


Figure 7.25: Number of transfers per station in original situation and final solution for Amsterdam Zuid area

The total amount of transfers at Amsterdam Centraal remains reasonably constant, although there is a shift from the train station to the metro station (see Figure 7.26). These stations were modelled separately, in which “Amsterdam Centraal” is connected to all train lines, except those rerouted through the metro tunnel between Amstel and Centraal, and “Centraal Station” is connected to all metro lines and the train lines through the aforementioned metro tunnel.



Figure 7.26: Number of transfers in original situation and final solution for Amsterdam Centraal

The shift in transfers from the train station to the metro station enforces the notion that the local system plays a more important role in the final solution than in the original situation. Several regional and national lines do not reach Amsterdam Centraal anymore, or do not stop there. Passenger that still need to transfer to a line at Amsterdam Centraal and previously used one of these lines, now use the local lines to reach Amsterdam Centraal or transfer at another station.

7.3.6 Walking

As mentioned at the start of this section, the reduced walking costs will be discussed now the new network has been shown extensively. There are several places where walking costs have decreased and they all have the same reason: passengers can use a station closer to their origin or destination, as it is now served by a line it wasn't served by previously.

An example is the centre of Amsterdam, which is shown in Figure 7.27. The figure shows that passengers walk to the stations served in the original situation by metro. In the final solution, the eastern metro branch (including stations Nieuwmarkt and Waterlooplein) is served by regional trains as well. This causes the shift from Amsterdam Centraal to the aforementioned stations.

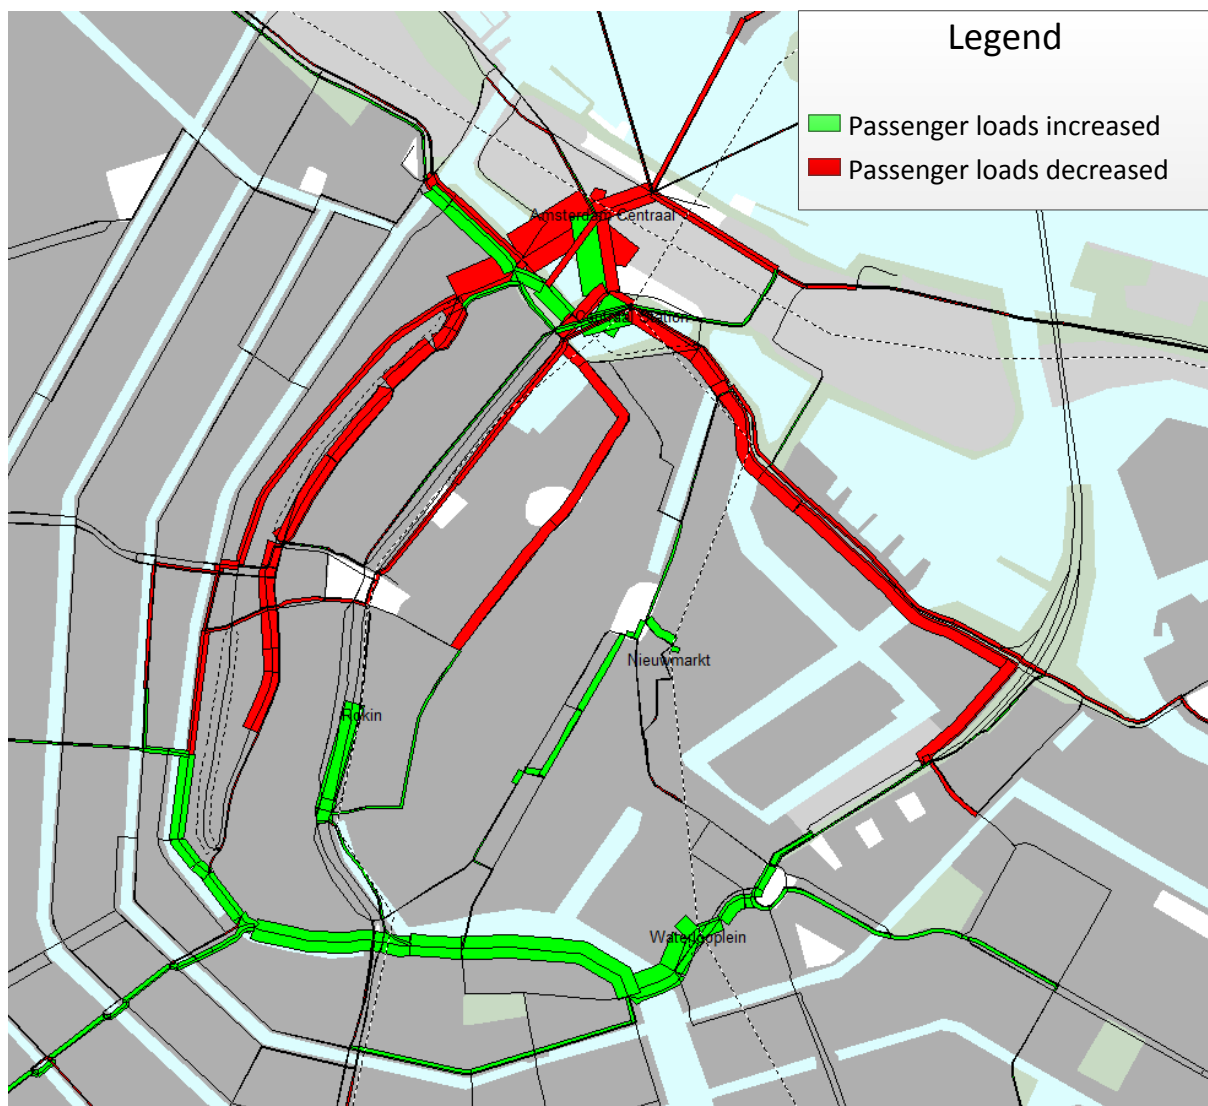


Figure 7.27: Difference walking in centre between original situation and final solution

7.4 Feasibility of the results

The solution as it emerged from SPAFOM is not feasible within the current infrastructure. A few of the reasons have been mentioned already, the most important reason being the limited capacity of the current infrastructure. The stations are the bottlenecks, as stopping takes time and letting a new vehicle stop every minute on only one track is simply not possible. The resulting network should therefore be interpreted as a guide to which lines are important for the network and on what part of the network they are important.

The North-South metro line, for instance, is an important connection over its entire length. In the north, it provides a fast transfer to the city centre for people arriving by bus from towns to the north, while in the south it connects Hoofddorp and Amsterdam Zuid to the city centre. When the metro line is connected to Schiphol Airport as well, the value of the line rises even more as it connects two of the major zones directly. This does not mean that the metro line should run every minute, but rather that a highly frequent connection adds value to the network.

The tunnel to Schiphol, which is now used solely by trains, could be used by metros as well. The termination of the intercity line to/from Enschede implies that not all train connections to Schiphol are needed. However, one should keep in mind that the used assignment model is focused on the metropolitan region of Amsterdam, which underestimates travellers to and from outside the region, especially if destinations are further away. A national model should be used to study the consequences of cutting short national rail lines.

Another reason why the frequencies from the solution are high, is that the amount of available vehicles has not been taken into account. The difference in frequency between the original situation and the solution is so high, that a large amount of extra vehicles need to be bought. These costs were not taken into account in the model, as the number of vehicles for 2040 is not known. Only operational costs were taken into account and these were relatively low due to the short length of the metro lines compared to regional and national train lines.

Allowing regional trains to use the metro tunnel between Amsterdam Centraal and Amstel proved to be a small success, as each station in the tunnel got one train stop. This was at the expense of some metro stops, two of which were cut short before reaching Amsterdam Centraal. It became clear that Amsterdam Centraal is not a destination for many passengers and that it is mainly used as transfer station in the regional and national network. Passengers using the metro often need to be in the city centre, so the capacity of the tunnel now dedicated to metro could partially be shifted to regional trains. This provides extra direct connections for passengers from the region around Amsterdam to the centre, while freeing up space in the train part of the station.

One of the outcomes of the model that will probably never happen is the complete abandonment of the rail infrastructure between Holendrecht and Gein. As the infrastructure is already there, it will lead to both societal and political resistance if the metro lines are cut short before reaching the area. However, the results show that it is certainly not needed to run two metro lines to the area, so a variant in which only one line to Gein might be possible. The results of the model suggest this to be the line to Amsterdam Zuid and Isolatorweg, if there is a possibility to transfer to another line towards the centre with a high frequency.

7.5 Limitations

The results of the Amsterdam case study have been influenced by limitations within the model, which result in that the outcome of the model will deviate from what would happen if the result of the model would really be implemented. The most important limitations will be discussed in this section.

Optimal solution not guaranteed

As a brute force method is not possible within the time constraints of this thesis, not all solutions can be assessed. Therefore, it cannot be guaranteed that a solution is optimal. This was clarified by the example of adding a stop of the North-South metro line at Schiphol airport.

Capacity constraints

The most important limitation is the lack of capacity constraints, as the choice was made to run SPAFOM without the constraints on infrastructure capacity to search for the network that would be really optimal. The numerical experiments have shown that including infrastructure capacity constraints shapes a large part of the solution, especially around busy stations. It is safe to say that the final solution of the model would never fit within the current infrastructure of Amsterdam, for instance because of the two metro lines that both run every minute on the branch towards Utrecht. Even if headways of one minute between two metro vehicles could be achieved, all four available tracks would be needed for the metro, leaving no room for any of the train lines. Moreover, this is without taking stopping into account, which takes even more time.

It would be interesting to see what kind of solution would come from a capacity constrained SPAFOM run. Unfortunately, in this thesis there was no time to run the model again, because of the long run time of SPAFOM for Amsterdam. However, a run with capacity constraints is not straightforward, as the model was run for the year 2040. Technological advancements could influence the capacity of rail infrastructure significantly, so the current infrastructure capacity cannot simply represent a future capacity.

Fixed origin-destination matrix

Another limitation of the model is the fixed origin-destination matrix, which causes people to always travel to the same destination, regardless of the network. Also in this case, time constraints are the reason for this, as including destination choice in the model would increase its run time to an unworkable amount of time of up to two days per solution.

Limited amount of lines

The lines that were subject to optimization were fixed as well, as SPAFOM is not able to construct new public transport lines. This limited the options of SPAFOM, especially for lines that run outside of the study area. SPAFOM optimized the network with the lines that were provided, which makes it even less likely that really the optimal network would be found. The optimal network would then have to be a combination of lines already in the network. On the other hand, the fixed origin-

destination matrix is partially a result of this network, as the network determines the travel costs for a passenger to a destination and therefore influences the destination choice of the passenger.

Average waiting time

An assumption that was already discussed in the numerical experiments is that of average waiting time for every line a passenger uses. Especially with today's possibilities to plan routes, passengers can easily plan their journey to avoid unnecessary waiting. However, to more accurately calculate waiting time, a schedule would be needed to calculate for instance the waiting time during a transfer. This level of detail is not included in the model, with the consequence that waiting time could be overestimated and frequency increases have a large impact on the passenger costs.

No operator distinction

In SPAFOM, the operator costs are assumed to be borne by one operator. For the numerical experiments, this was a safe assumption, as it was a fictional network. However, in the case of Amsterdam, there are multiple public transport operators. Even if the overall operator costs decrease, the costs for one operator can increase a lot, while those of another operator decrease. This might be unacceptable, if the former operator cannot compensate the extra costs with revenues or the latter operator loses more revenue than costs are decreased. In the case of Amsterdam, it can be seen that the frequencies of metro lines, which are operated by GVB, are increased a lot, while some train lines, operated by NS, stay behind. This also induced a shift in passengers from train to metro lines, which means that NS loses revenue, while GVB wins revenue. However, to provide the services that resulted from the model, GVB has to invest in more vehicles, which could cost more than is gained in revenue.

Scale of the study area

The scale of the study area plays a role as well and relates to the previous limitation. The further a line runs outside the study area, the less accurate its occupancy is, often underestimating it. This is because VENOM is the regional model of Amsterdam and although it is accurate for quite a distance outside Amsterdam, it does not extend to Groningen or Eindhoven. Relations outside Amsterdam, such as Eindhoven-Utrecht, are not taken into account in for instance the determination of the frequency of the train between Eindhoven, Amsterdam Centraal and Schagen. This is a result of the scope demarcation, and can in further research be overcome by assigning a minimum frequency to certain lines.

7.6 Conclusions of the Amsterdam case study

The results of SPAFOM for the case study of the Amsterdam metro and train network show that an increased frequency on all lines is beneficial for the system. The reduced waiting costs lower the threshold for a transfer, which results in a more extensive use of the local metro lines. An important role can be played by the North-South metro line, especially if it is extended towards Schiphol.

The trend in the final solution is that passengers from outside Amsterdam, with a destination in the city centre, travel by train or bus to the nearest station connected to a metro line and use the metro line from there. This is made easier when frequencies of the metro lines are high, as waiting time plays a large role in the transfer threshold. Therefore, an overall distinctive role for train and metro is still present within the Amsterdam area, with trains providing the regional and national services and metros providing the service for the last mile(s) to passengers with a destination within Amsterdam.

The shorter a line is, the higher its frequency becomes. This is a direct result of taking operator costs into account, as longer lines are more expensive for operators. Another reason is that the model focuses on the region of Amsterdam, so only a very small share of the passengers that would profit from a higher frequency on national lines are taken into account. This resulted in high frequency metro lines, medium frequency regional lines and low frequency national lines.

Some metro lines were shortened, especially around Amsterdam Centraal. Here, regional lines took over some of the functions of the metro lines, so the high frequency metro lines could be shortened to save operator costs. This not only happened at Amsterdam Centraal, as there were several areas that were served sufficiently well by other public transport lines that metro lines could be shortened.

The production and attraction of passengers of zones plays a large role in this too. Again, Amsterdam Centraal is the best example, as the area around it does not produce or attract many passengers compared to for example the city centre. Providing direct connections to the city centre via regional lines made the metro partially redundant, which showed in the reduction of stops of the metro in the centre.

Despite this reduction in the number of stops per line, increasing travel speed, the largest contributor to the reduction of passengers costs were the walking costs. In the original situation, passengers from outside Amsterdam had two options to reach the centre: walk from Amsterdam Centraal or transfer to another public transport service. As some regional lines provided direct access to the city centre, passengers can egress closer to their destination, without the need for a transfer. Due to the high attraction value of the city centre, this small reduction in walking time affected many passengers, resulting in a relatively large reduction in passenger costs.

Still, the metro lines are the most important means of transport for passengers within Amsterdam. Especially the North-South metro line plays an important role, as it serves areas that produce and attract many passengers. Furthermore, the line provides a direct connection to Amsterdam Zuid, which is served by many lines and facilitates a large part of all transfers.

8 Conclusions and recommendations

The final chapter of this thesis contains the conclusions of the thesis, which are separated in scientific conclusions (section 8.1) and practical conclusions (section 8.2). This separation is parallel to the separation of chapters 6 and 7. The chapter concludes with a summation of the limitations of this research in section 8.3 and guiding for future research in section 8.4.

8.1 Conclusions

This thesis has developed a model to optimize stopping patterns and frequencies of a given public transport network regarding both passenger and operator costs, with the choice to enforce or relax infrastructure capacity constraints. The model is called Stopping Pattern and Frequency Optimization Model, or SPAFOM in short. Its goal is to answer the following research question:

Given a public transport passenger demand, a network with multiple services and a set of possible lines, what are the optimal stopping patterns and frequencies for each, while satisfying capacity constraints?

The model uses a genetic algorithm to improve a start set of solutions, based on the sum of walking costs, in-vehicle costs, transfer costs, unsatisfied demand costs, fixed operator costs and variable operator costs. This combination, together with the inclusion of an underlying network and the extensive passenger assignment model, makes a comprehensive analysis of potential solutions possible.

This approach shows some similarity with the work of Arbex *et al.* (2015), which also used a genetic algorithm to assess both passenger and operator costs. Lin and Ku (2014) used a genetic algorithm as well, but only to assess operator costs. Both applied their algorithm to a network of limited size, similar to the test network used in this thesis.

In order to answer the research question, several sub questions have been researched and answered.

How should the set of lines, for which the stopping patterns and frequencies will be optimised, be determined?

The decision was made to use the existing lines in a network, as the changes in the network can then be related to the existing situation. For the Amsterdam case study, the train and metro lines were chosen, to limit the amount of decision variables. Papers that did not have route set choice as part of the optimization chose for the existing lines as well, for instance Goossens *et al.* (2006).

The start set of solutions is not only determined at random. It proved to be beneficial for the speed of the model to include designed lines, lines that were constructed beforehand. This way, the genetic algorithm converged faster, as the constructed solutions were found to be better than the first set of random solutions. The random solutions are still needed, though, to ensure enough diversity throughout the process. Otherwise, the model would almost certainly come up with a combination of the designed solutions, due to the crossover mechanism.

The combination of stopping patterns and frequencies was not found in literature, which mostly focused on either of the two. In this literature, frequencies were often optimized together with route sets (Ceder and Wilson, 1986; Fan and Machemehl, 2004; Borndörfer *et al.*, 2008; Schmidt, 2014; Arbex *et al.*, 2015), while in this thesis the set of routes was already fixed.

Are existing, traditional services, with a clear distinction between long distance and short distance services and according stop spacing, optimal for passengers and operators, or is another stopping pattern more desirable?

Each solution contains the stopping pattern and frequency of each line that is chosen to be included in the optimization. Moreover, for each line a choice can be made which station(s) need to be included in the optimization. The frequency is determined for a whole line and can be set to zero, excluding the line from the network.

The original stopping patterns were not altered very much in the numerical experiments. If changes were made, these were mostly caused by a lack of capacity, not by a saving in passenger or operator costs. In the Amsterdam case study, the resulting stopping pattern showed some more changes compared to the original situation. Some regional and national lines gained one or more additional stops, while in the case that two lines served the same demand, the amount of stops decreased for one of those lines.

Based on the results of the test network and the Amsterdam case study, the distinction between long distance services with few stops and short distance services with many stops are justifiable, as long as the services with the most stops run with the highest frequency. Each extra stop reduces the walking time of passengers, which often outweighs the longer in-vehicle time for passengers that have an extra stop along their journey. It also proved better to serve an origin-destination relation with two local services with high frequency, which includes a transfer, than to serve the relation with one long distance service with a low frequency.

In urban areas such as Amsterdam, this is what often happens, as tram and metro services stop at almost every station they pass. However, there are very few cities where local and regional/national services really run together, as not only their infrastructure is often separated, but the services are operated by different operators as well. Karlsruhe is one of the examples of a city where both services run on the same infrastructure.

Lin and Ku (2014) found the same results for their case study of a part of the Taiwanese rail network, where the local trains were preferred in cases of lower demand. Only when demand increased, express services were used as well to serve long distance demand. Frequencies were not part of the optimization in this paper. Other papers focus on the performance of the algorithm instead of the actual network results and can therefore not be compared in this regard.

What frequencies are needed for the different services on the lines to satisfy passenger demand?

Unfortunately, enforcing a vehicle capacity constraint in the model would have taken too much computation time, which was already considerable. Instead, a crowding factor was applied to the in-vehicle time experienced by passengers, which was adapted from Cats and Hartl (2013). As waiting time has a large influence on passenger costs, the frequencies of all lines were very high in the model results. This reduced not only the waiting time, but the crowding costs as well. The frequencies should therefore be increased to reduce passenger costs, until the infrastructure capacity is reached or the extra operator costs outweigh the reduced passenger costs.

Could the current infrastructure accommodate the desired stopping patterns and the required frequencies for a future passenger demand?

The capacity constraints, which were enforced in the numerical experiments and relaxed in the Amsterdam case study, act for both the stations (nodes) and the line infrastructure (links) of the network. The usage of a link or node is determined by summing the minimum headways of the lines using the link or node. For nodes, the minimum stopping time is added as well if the line stops at the node. Decelerating and accelerating is taken into account in this stopping time.

In literature, the capacity constraints most often apply to the vehicle, for instance vehicle capacity (Lin and Ku, 2014; Arbex *et al.*, 2015), a maximum number of available vehicles (Ceder and Wilson, 1986), or both (Fan and Machemehl, 2004; Schmidt, 2014; López-Ramos *et al.*, 2017). Link capacity (Borndörfer *et al.*, 2008) or station capacity (Yue *et al.*, 2016) are taken into account much less.

The results of the test network showed that it is beneficial for the total costs to run as many vehicles as possible within the limitations of the infrastructure capacity, as this reduces both waiting time and perceived in-vehicle time due to crowding. Although the downside is that operator costs increase significantly, passenger costs can be decreased even more. Furthermore, a better service can attract more (paying) passengers, which was not taken into account in this thesis, but could be an important factor in the decision making process of operators.

However, the capacity of the infrastructure often limits the possibilities to increase service frequencies. This can be partially overcome by adapting stopping patterns, so fewer lines use the same infrastructure, allowing a higher frequency on all lines. Although this reduces the amount of direct connections, and thus increases the amount of transfers, the reduction in waiting time makes up for the extra transfer. For Amsterdam, capacity constraints were relaxed, but it is clear that the resulting frequencies from SPAFOM are not possible within the current infrastructure capacity limitations.

Given a public transport network with multiple services and a set of possible lines, what are the optimal stopping patterns and frequencies for each line serving a given passenger demand, while satisfying capacity constraints?

To answer the main research question: in all scenarios in the numerical experiments, as well as for the Amsterdam case study, extra stops were beneficial for the total costs. Express services with a limited amount of stops are only beneficial if there is a local service with many stops along the same route and the demand justifies an extra service. Frequencies should be as high as capacity constraints allow, as this reduces waiting and crowding costs for passengers. The influence of operator costs is very limited, if fixed costs are ignored. In the numerical experiments, the result was influenced more by capacity limitations than by operator costs.

8.2 Practical implications

The results for the test network imply that operators should focus on reducing waiting time and, to a lesser extent, walking time. They also imply that these two factors are more important than pure in-vehicle time. The way to do this is by running services with many stops at a high frequency. This higher frequency reduces crowding in the vehicles, which is part of the perceived in-vehicle time, and compensates for the extra in-vehicle time due to the extra stops.

On the level of scale of a city or metropolitan area, this means focusing on the metro network, as the metro has a relatively short stop spacing and a relatively high frequency. This was also the outcome of the SPAFOM run for the metropolitan region of Amsterdam, where the highest frequencies were all found at the metro lines.

Although the Amsterdam case study was run without infrastructure capacity limitations, it can be expected that when infrastructure capacity constraints are enforced, a similar outcome can be observed. In all capacity constrained test scenarios a higher frequency for lines with short stop spacing was favoured over having fast lines with wide stop spacing. This implicates that it is better to have fewer lines with higher frequencies and more need for transfers, than more direct connections with lower frequencies and fewer transfers.

The run time of the model itself was very long, about a month. This is not practical for planning, as the model requires a lot of input that influence its results. A lot of improvement is needed to make the model applicable in practice. Suggestions for this can be found in section 8.4.1.

8.3 Limitations

The limited time available for this research has resulted in several choices to make the research possible in the available time. A selection of the choices is listed in Table 8.1, grouped together depending on the type of limitation: parameter limitation, optimization limitation or assignment limitation. The expected impact on the results is listed as well, which represents how likely it is that the results of SPAFOM would be different if another choice would have been made. The choices, their influence and the expected impact on the results are explained more extensively after the table.

Table 8.1: Limitations

Group	Choice	Influence	Expected impact on results
Parameter	Average waiting time	Waiting time might be overestimated	Large
	Stopping time independent of number of passengers	The stopping time is the same for the same vehicle, while in reality it depends on the number of passengers alighting and boarding the vehicle	Low
Optimization	Only existing lines	Resulting network depends on original situation	Medium
	No operational constraints	Resulting network might not be possible in practice	Low
	No differentiating in operators	Even if total operator costs are decreased, one operator might see an unacceptable increase in costs	Medium
	Fixed vehicle types	Vehicle types do not change with different stopping patterns	Low
Assignment	Only two networks	Unclear if other networks yield other results	None
	Combined frequencies	Each line is considered separately for a journey, even if more lines offer the exact same journey between two stations	Medium
	No disruptions	Network robustness is not taken into account	Medium

Average waiting time

One of the most important choices was to use the average waiting time of passengers when calculating their waiting costs. Because of this, the waiting costs in SPAFOM might be overestimated, as passengers in the model do not time their arrival. Although this is not realistic, it is not possible to calculate waiting time more accurately without a timetable. This timetable could have been constructed for the current situations of the test case (where it would be a fictional current situation) and the Amsterdam case, but it was not possible to include scheduling in SPAFOM.

This would need to be a fully automatic process for each evaluated solution and both the construction of the process and the process itself while running the model would take too much time. By capping the waiting time to a maximum, the results were made somewhat realistic, but still the influence of the waiting time on the final solution was large.

Stopping time independent of number of passengers

The time a vehicle needs per stop does not depend on the number of passengers that alight or board the vehicle, nor on the number of passengers inside the vehicle. All three can potentially have influence on the stopping time, especially if it concerns many passengers. Including this in the model would require several iterations per solution, as more stopping time per vehicle means fewer vehicles fit within the same infrastructure capacity. This would affect the network, which would in turn change passenger flows, leading to a large increase in run time. As only the test network included infrastructure capacity constraints and the effects of this were expected to be quite low, making stop time dependent on the number of passengers was not judged to be worth the extra calculation time.

Only existing lines

The use of only existing lines to optimize a network limits the model in finding the optimal solution, although this could not be guaranteed anyway. Just like scheduling, public transport line construction was not within the scope of this thesis, as it would not have been able to include it in the model within time. Moreover, it would expand the solution space of the model, which would lead to much longer running times, both because of the extra operations and the extra generations needed for convergence.

No operational constraints

The Amsterdam SPAFOM run resulted in the shortening of several lines, especially on lines with high frequencies, where such a shortening resulted in a substantial reduction of operator costs. Although the reduced operator costs are welcomed by operators, there might be complications in operating such lines. For example, trains and metros are allowed to change directions at each station, while it could be that this is not possible there due to infrastructure limitations. The strategical level of the model did not allow taking the effect of heterogeneity on the network's capacity into account, either. The impact on the results is not expected to be high, as this is irrelevant for the numerical experiments and capacity was neglected in the Amsterdam case study anyway. However, if both operational and capacity constraints would be enforced, the result would probably differ significantly. This will be elaborated on in the next section, where directions for future research are proposed.

No differentiating in operators

The choice not to evaluate operator costs per operator has been made specifically so this would not influence the results. It was not possible to make a trade-off about what operator is more important or research how unequal results for operators should have been taken into account, as the focus was on the different aspects of passenger costs. Differentiating between operators would require a level of detail that does not match with the rest of the thesis, such as more exact cost figures of vehicles.

Fixed vehicle types

The vehicle types were bound to the transit lines, independent of what suits the line. For example, it might be beneficial to change from a train to a metro vehicle if a line is shortened, stops more, and/or increases in frequency. This is not taken into account in the model, leading to less efficient networks. However, the impact on the results is probably quite low, as the biggest difference between the two vehicles is the minimum headway, which is not relevant in the Amsterdam case study.

Only two networks

Another choice that was made was to run the numerical experiments for different scenarios in the same network. Although this provides a good view of what happens if a network is used differently, which is what happens in reality, it does not show what the outcome of the model would be in different networks. The results of the Amsterdam case study do imply that the results are somewhat similar, partially because of the way waiting time was incorporated, but the large differences between the two networks in both size and complexity make it hard to compare them. The choice for only one test network has not resulted in different results, but it has limited the interpretation of the results.

Combined frequencies

On some links in both test network and the network of Amsterdam, multiple lines provide an option to travel between two stations. In the used passenger assignment model, these lines were treated as completely separate choices for passenger routes, while in truth the lines could be seen as one on that link. This allows aggregating the frequencies of both lines, but this is not possible in the model. Not allowing lines to ‘work together’ could have quite some impact on the results, as waiting time was an important factor in the determination of the results. Lines now competed for space on the network, instead of using it efficiently together.

No disruptions

The model assumes a perfect situation, in which no disruptions take place and every link is always available. In practice, a blockage could cut off any line at any time. If network robustness becomes a factor in the design of a new solution for the network, it might very well be possible that another solution comes out as the best.

8.4 Future research

The limitations of this thesis open up directions for further research. These directions can be discerned in two categories: further research to improve the model and further research using the results of the model. They will be discussed in that order.

8.4.1 Model improvements

SPAFOM has been developed on a strategical level, which had as result that the level of detail of the results was not very high. All aspects of the model can be improved and a few of those aspects are listed below. They mostly come forth from the limitations listed in the previous section.

Computational efficiency

The main computational limitations of the model were caused by the assignment model. Although the assignment model had already been stripped from a lot of redundant links and nodes, this might be improved even further. This will make the path choice for each origin-destination pair easier, reducing the run time per assignment. The run time reduction could also be achieved by aggregating even more zones, but the effect on the accuracy of the assignment model should then be monitored carefully.

The optimization model used the OmniTRANS database, which consists of tables, to calculate the different costs needed to calculate the objective function value per solution. When more than one table needs to be consulted for information, lengthy iterations take place. These iterations can be avoided if all information per cost component can be found in one table. However, this might result in a complete database reconstruction and software changes, which might take more time than can be gained.

The code that was used is certainly not optimal, as it was written with only basic programming knowledge and a very basic knowledge of the Ruby language. Improvements can most probably be made, as the optimization process is very repetitive and could be parallelised for an important part. The latter also depends on the available hardware, as the runs in this thesis were performed on a medium level laptop. Using advanced desktop computers or even clusters could improve the model run time substantially, provided that the code is adjusted accordingly. Again, this also depends on whether the OmniTRANS software allows for this.

Model on tactical level

The model can be improved by taking the level of detail a step further, to the tactical level. On this level timetables are generated, which make it possible to calculate costs such as waiting costs more accurately. This would also improve the infrastructure capacity constraint mechanism, which now relies on average headways, and operational constraints could be taken into account as well.

These operational constraints could also take network robustness into account, for instance by adding disruptions to the network. The passenger assignment model could be able to cope with disruptions, after which the optimization part would take this into account as well. However, the effect should be monitored closely, as the capacity constraints in this thesis have shown that limitations in the network have a large influence on the result.

More detailed input

For the input of the model, most of the time average values have been taken, for instance for dwell times of vehicles in stations. Operational costs could be given extra detailing as well, for instance extra energy costs if an extra stop is being made.

More networks

As mentioned in the limitations, the model could be used to explore the impact of different networks on the resulting stopping patterns and frequencies. This could be combined with one or more of the model improvements mentioned above.

8.4.2 Usage of results

The results from SPAFOM are exactly that: model results. They cannot be used directly to improve public transport networks, although they do provide insights in the direction that could be taken. This section provides some directions that could be used for further research.

Results on tactical level

On the tactical level, simulation software could be used to run stochastic simulations of solutions, for instance including the dependence of stopping time on the number of passengers that want to alight or board a vehicle. This is not possible now, due to the high run times. However, SPAFOM does not only output the final best solution, but all other assessed solutions as well. A selection of the best solutions can then be further optimised.

Implications for operations

The implications of the results for operating the public transport lines as resulting from SPAFOM have only been explored superficially, as the focus of the model was on passengers. The network optimization model could also be explored from an operator point of view, as the model has shown to be capable of reducing operator costs as well.

Bibliography

- Amstelveenlijn | Waarom vernieuwen we de Amstelveenlijn?* (2015). Available at: <http://amstelveenlijn.nl/waarom-vernieuwen-we-de-amstelveenlijn/> (Accessed: 13 July 2017).
- Arbex, R. O. et al. (2015) 'Efficient transit network design and frequencies setting multi-objective optimization by alternating objective genetic algorithm', 81, pp. 355–376. doi: 10.1016/j.trb.2015.06.014.
- AT Osborne (2017) *Afwegingskader exploitatiekosten*.
- Borndörfer, R., Grötschel, M. and Pfetsch, M. E. (2008) 'Models for line planning in public transport', in *Lecture Notes in Economics and Mathematical Systems*, pp. 363–378. doi: 10.1007/978-3-540-73312-6_18.
- Bus | GVB (2017). Available at: <https://www.gvb.nl/vervoer/bus> (Accessed: 5 September 2017).
- Cats, O. and Hartl, M. (2013) 'Modelling public transport on-board congestion: comparing schedule-based and agent-based assignment approaches and their implications', pp. 1–21.
- Cats, O., West, J. and Eliasson, J. (2016) 'A dynamic stochastic model for evaluating congestion and crowding effects in transit systems', *Transportation Research Part B*, 89, pp. 43–57. doi: 10.1016/j.trb.2016.04.001.
- Ceder, A. and Wilson, N. H. M. (1986) 'Bus network design', *Transportation Research Part B: Methodological*, 20(4), pp. 331–344. doi: 10.1016/0191-2615(86)90047-0.
- Centraal Planbureau and Planbureau Voor De Leefomgeving (2015) *WLO-Cahier Mobiliteit*. Available at: http://www.wlo2015.nl/wp-content/uploads/PBL_2015_WLO_Mobiliteit_1686.pdf (Accessed: 9 March 2017).
- Cipriani, E., Gori, S. and Petrelli, M. (2012) 'Transit network design: A procedure and an application to a large urban area', *Transportation Research Part C: Emerging Technologies*. Elsevier Ltd, 20(1), pp. 3–14. doi: 10.1016/j.trc.2010.09.003.
- CROW (2015) 'Kostenkengetallen regionaal openbaar vervoer 2015'.
- Desaulniers, G. and Hickman, M. D. (2007) 'Public Transit', *Handbooks in Operations Research and Management Science*, 14(C), pp. 69–127. doi: 10.1016/S0927-0507(06)14002-5.
- Van Eck, G. (2010) 'Ontwerp van stedelijke openbaar vervoer netwerken Herontwerp van het Utrechtse buslijnnennet met behulp van een genetisch algoritme'. Available at: http://www.verdus.nl/upload/documents/Ontwerp_van_stedelijke_openbaar_vervoer_netwerken_-_2010_-_Van_Eck.pdf (Accessed: 10 July 2017).
- Fan, W. and Machemehl, R. (2008) 'A tabu search based heuristic method for the transit route network design problem', *Computer-aided Systems in Public Transport*, pp. 387–408. doi: 10.1007/978-3-540-73312-6_20.
- Fan, W. and Machemehl, R. B. (2004) 'Optimal Transit Route Network Design Problem : Algorithms , Implementations , and Numerical Results', 7(2), p. 268.
- Financieel Dagblad (2016) *ProRail: verleng Noord-Zuidlijn naar Schiphol via ons spoor | Het Financieele Dagblad*. Available at: <https://fd.nl/economie-politiek/1165105/prorail-verleng-noord->

zuidlijn-naar-schiphol-via-ons-spoor (Accessed: 13 September 2017).

Financiën in het kort | NS 2016 (2017). Available at: http://www.nsjaarverslag.nl/jaarverslag-2016/s1234_financininhetskort2/a1091_Financiën-in-het-kort (Accessed: 5 May 2017).

Garis, H. De (2004) *Book Review: Introduction to Evolutionary Computing*, New York. doi: 10.1162/evco.2004.12.2.269.

Gemeente Haarlemmermeer (2005) *Eindrapportage regionale vervolgstudie openbaar vervoer Corridor Amsterdam-Schiphol-Haarlemmermeer*.

Givoni, M. and Rietveld, P. (2014) 'Do cities deserve more railway stations? The choice of a departure railway station in a multiple-station region', *Journal of Transport Geography*. Elsevier Ltd, 36, pp. 89–97. doi: 10.1016/j.jtrangeo.2014.03.004.

Goossens, J.-W., van Hoesel, S. and Kroon, L. (2006) 'On solving multi-type railway line planning problems', *European Journal of Operational Research*, 168(2), pp. 403–424. doi: 10.1016/j.ejor.2004.04.036.

Govers, B. (2017) 'Europees Spoor in drie scenario's'. Available at: <http://www.railforum.nl/wp-content/uploads/2017/03/Presentatie-Bas-Govers.pdf> (Accessed: 23 March 2017).

Gu, W., Amini, Z. and Cassidy, M. J. (2016) 'Exploring alternative service schemes for busy transit corridors', 93, pp. 126–145. doi: 10.1016/j.trb.2016.07.010.

Guis, N., Keizer, B. De and Nes, R. Van (2012) 'Lijnvoering van de toekomst voor het Nederlandse spoornetwerk', *Tijdschrift Vervoerswetenschap*, (3), pp. 84–98.

Ibarra-Rojas, O. J. et al. (2015) 'Planning, operation, and control of bus transport systems: A literature review', *Transportation Research Part B: Methodological*. Elsevier Ltd, 77, pp. 38–75. doi: 10.1016/j.trb.2015.03.002.

Keizer, B. De, Fioole, P.-J. and van 't Wout, J. (2013) 'Optimalisatie van de lijnvoering op Railnetwerken', *Colloquium Vervoersplanologisch Speurwerk*, (november). Available at: http://www.cvs-congres.nl/cvspdfdocs_2013/cvs13_067.pdf.

Kouwenhoven, M., de Jong, G. and Koster, P. (2014) 'New Values of Time and Reliability in Passenger Transport in the Netherlands', *Research in Transportation Economics*, 47, pp. 37–49. Available at: http://eprints.whiterose.ac.uk/84748/7/VOTVORpassengerpaper_v1_3.pdf (Accessed: 19 May 2017).

LeighFisher (2014) 'ERTMS in Nederland'.

Lin, D. Y. and Ku, Y. H. (2014) 'Using genetic algorithms to optimize stopping patterns for passenger rail transportation', *Computer-Aided Civil and Infrastructure Engineering*, 29(4), pp. 264–278. doi: 10.1111/mice.12020.

López-Ramos, F. et al. (2017) 'Integrated approach to network design and frequency setting problem in railway rapid transit systems', *Computers & Operations Research*, 80, pp. 128–146. doi: 10.1016/j.cor.2016.12.006.

Mauttone, A. and Urquhart, M. E. (2009) 'A route set construction algorithm for the transit network design problem', *Computers & Operations Research*, 36, pp. 2440–2449. doi: 10.1016/j.cor.2008.09.014.

- Metro* | GVB (2017). Available at: <https://www.gvb.nl/vervoer/metro> (Accessed: 3 July 2017).
- Ministerie van Infrastructuur en Milieu (2017) 'Nationale Markt- Markt en Capaciteitsanalyse 2017 (NMCA) Hoofdrapport', 2017, pp. 1–62.
- Mitchell, M. (1996) 'An introduction to genetic algorithms', *Computers & Mathematics with Applications*, 32(6), p. 133. doi: 10.1016/S0898-1221(96)90227-8.
- Nayeem, M. A. K., Rahman, M. and Sohel, R. (2014) 'Transit network design by genetic algorithm with elitism'. doi: 10.1016/j.trc.2014.05.002.
- van Nes, R. (2002) *Design of multimodal transport networks: A hierarchical approach*, Proefschrift. doi: 90-407-2314-1.
- van Nes, R. (2015) 'PTND 1 Introduction and urban networks 2015'.
- van Nes, R. and Bovy, P. H. L. (2000) 'Importance of Objectives in Urban Transit-Network Design', *Transportation Research Record*. Available at: <http://trrjournalonline.trb.org/doi/pdf/10.3141/1735-04> (Accessed: 13 March 2017).
- NS (2017) 'Intercity (VIRM) fabrikant'. Available at: http://www.ns.nl/binaries/_ht_1495619211105/content/assets/ns-nl/over-ns/2017/virm-poster.pdf (Accessed: 3 July 2017).
- OmniTRANS - DAT.Mobility* (2017). Available at: <http://www.dat.nl/nl/producten/omnitrans/> (Accessed: 3 May 2017).
- Ortúzar, J. de D. and Willumsen, L. G. (2011) *Modelling Transport*. Fourth Edi. Chichester, UK: John Wiley & Sons, Ltd. doi: 10.1002/9781119993308.
- ProRail (2017) *Netverklaring 2018*. Available at: https://www.prorail.nl/sites/default/files/netverklaring_2018_versie_1.3.pdf (Accessed: 5 July 2017).
- RET Jaaroverzicht 2015 - Jaarverslag 2015* (2016). Available at: <http://jaaroverzicht2015.retjaarverslag.nl/> (Accessed: 5 May 2017).
- Ruby Programming Language* (2017). Available at: <https://www.ruby-lang.org/en/> (Accessed: 25 May 2017).
- Schafer, A. et al. (2000) 'The future mobility of the world population', *Transportation Research Part A*, 34, pp. 171–205. Available at: www.elsevier.com/locate/tra (Accessed: 10 March 2017).
- Schmidt, V. (2014) 'Hybrid large neighborhood search for the bus rapid transit route design problem', *European Journal of Operational Research*, 238, pp. 427–437. doi: 10.1016/j.ejor.2014.04.005.
- Tram* | GVB (2017). Available at: <https://www.gvb.nl/vervoer/tram> (Accessed: 5 September 2017).
- Yue, Y. et al. (2016) 'Optimizing train stopping patterns and schedules for high-speed passenger rail corridors', *Transportation Research Part C: Emerging Technologies*, 63, pp. 126–146. doi: 10.1016/j.trc.2015.12.007.
- Zeeburgereiland: nieuwe stadswijk aan het water - Gemeente Amsterdam* (2017). Available at:

<https://www.amsterdam.nl/projecten/zeeburgereiland/> (Accessed: 9 September 2017).

Appendix A: Definitions

A couple of terms need to be clarified beforehand, as they do not have the same meaning to everyone. The definitions given below will be used throughout the whole document and the thesis report. Some definitions will be refined in the thesis report.

Infrastructure: the physical rails and stations. There will be no distinction between infrastructure for train or metro, but instead it will be assumed that any vehicle can run on any track. Although this is not the case right now, it is technically possible to realise this.

Line: a set of consecutive links and nodes that is served by one type of train. A line consists of at least two nodes and a link between each node.

Link: representation of a stretch of rails between two nodes. Its most important property is its capacity as well.

Link capacity: the number of trains that can traverse a link in a certain time interval. This depends on the speed and the minimum distance between vehicles.

Node: representation of a train or metro station where trains and metros can stop. The most important property of a node is its capacity, both for stopping and traversing vehicles.

Node capacity: the number of trains that can stop at a node in a certain time interval. If one or more lines do not stop at a node, the capacity of the node is still reduced if a platform track is needed to let the train pass the node.

Passenger capacity: the maximum number of passengers that can be carried over a link or node in a certain time interval. This depends on the capacity of the trains and the number of trains running on a line.

Service: a vehicle with certain properties (speed, passenger capacity, comfort). Traditional services are:

(Intercity) train: high capacity train with focus on comfort, traditionally used for medium to long distances between major city stations.

Metro: medium capacity train for short distances within an urban area.

Stop train, tram, and bus are not part of the optimization, but can be part of the underlying network. Trains that stop at (almost) every station they pass, which are often referred to as stop trains, are modelled with as intercity train vehicles.

Stop: a vehicle halting at a station to let passengers in and out.

Appendix B: Solution data current and final solution

Line	Number of stations on line	Number of stops		Frequency (veh/hr)	
		Original network	Final solution	Original network	Final solution
Isolatorweg – Gein	20	20	15	5	60
Zuid – Amsterdam Centraal	10	10	8	4	30
Hoofddorp – Noord	11	8	10 (11)*	15	60
Gaasperplas – Amsterdam Centraal	14	14	10	8	30
Gein – Amsterdam Centraal	15	15	4	4	60
Vlissingen – Amsterdam Centraal	2	2	2	2	2
Zandvoort – Amsterdam Centraal	2	2	2	2	2
Uitgeest – Amsterdam Centraal (via Haarlem)	2	2	2	2	2
Uitgeest – Amsterdam (via Zaandam)	2	2	2	4	30
Eindhoven – Schagen	11	3	4	4	15
Hoofddorp – Hoorn	9	3	3	4	20
Groningen – Den Haag Centraal	8	3	6	6	15
Zwolle – Hoofddorp	8	5	6	6	12
Enschede – Schiphol	8	3	2	2	8
Utrecht – Amsterdam Centraal	4	4	4	4	30
Amersfoort – Amsterdam Centraal	4	4	2	2	15
Nijmegen – Schiphol	9	3	3	6	20
Rotterdam Centraal – Amsterdam Centraal	13	5	9	2	15
Rhenen – Amsterdam Centraal	13	5	6	4	12

*In the final solution of the Stopping Pattern and Frequency Optimization Model, the line did not stop at station Schiphol. However, including the station in the stopping pattern decreased all passenger costs.

Appendix C: Change in transfers per station, sorted

Station	Number of transfers in original situation	Number of transfers in final solution	Change (%)
Station Zuid	21528	14264	-34%
Amstelveenseweg	0	5358	-
Hoofddorp	0	5262	-
Amstelstation	2644	7454	182%
Station RAI	854	4851	468%
Strandvliet	0	3425	-
Spaklerweg	1	3187	239492%
Schiphol	6089	3543	-42%
Henk Sneevlietweg	0	2370	-
Station Bijlmer ArenA	2509	312	-88%
Station Diemen-Zuid	1500	3625	142%
Venserpolder	0	2068	-
Station Holendrecht	204	1800	782%
Amsterdam Centraal	8971	7826	-13%
Bullewijk	0	1053	-
Amsterdam Muiderpoort	0	1046	-
Station Lelylaan	997	0	-100%
Amsterdam Sloterdijk	3876	4720	22%
Duivendrecht	142	915	543%
Van der Madeweg	686	4	-99%
Weesp	3632	4297	18%
Overamstel	528	0	-100%
Weesperplein	1	430	45661%
Jan van Galenstraat	0	337	-
Station Duivendrecht	97	0	-100%
De Vluchtlaan	0	56	-
Waterlooplein	0	51	-
Heemstedestraat	0	24	-
De Pijp	2	6	156%
Vijzelgracht	2	5	189%
Rokin	1	4	405%
Wibautstraat	0	3	-
Europaplein	2	0	-100%
Amsterdam Science Park	0	1	-

Noorderpark	0	1	-
Isolatorweg	0	0	-
Diemen	0	0	-
Nieuwmarkt	0	0	-
Postjesweg	0	0	-
Verrijn Stuartweg	0	0	-
Ganzenhoef	0	0	-
Kraaiennest	0	0	-
Gaasperplas	0	0	-
Reigersbos	0	0	-
Gein	0	0	-
Amsterdam Noord	0	0	-