Ties Brands

Multi-objective Optimisation of Multimodal Passenger Transportation Networks

MULTI-OBJECTIVE OPTIMISATION OF MULTIMODAL PASSENGER TRANSPORTATION NETWORKS

Ties Brands

MULTI-OBJECTIVE OPTIMISATION OF MULTIMODAL PASSENGER TRANSPORTATION NETWORKS

PROEFSCHRIFT

ter verkrijging van
de graad van doctor aan de Universiteit Twente,
op gezag van de rector magnificus,
prof. dr. H. Brinksma,
volgens besluit van het College voor Promoties
in het openbaar te verdedigen
op donderdag 15 oktober 2015 om 14:45 uur

door

Tieske Brands

geboren op 28 april 1983
te Helmond

Dit proefschrift is goedgekeurd door de promotor en door de copromotor:

Prof. dr. ir. E.C. van Berkum (promotor)
Dr. ir. L.J.J. Wismans (copromotor)

Voorwoord

Eric van Berkum gaf mij in 2009 de kans te starten met een promotieonderzoek, waarvan het onderwerp zo goed bij mijn achtergrond aansloot (een combinatie van wiskundige optimalisatietechnieken en iets met openbaar vervoer!) dat ik daar wel op in moest gaan. Gelukkig ben ik nog steeds, zeker nu het traject zo goed als afgerond is, blij met die keuze. Een keuze die ik mede heb gemaakt nadat ik het er met Jaap Vreeswijk en met Wim Korver over gehad had, dank voor jullie advies.

Eric, ik ben blij dat jij mijn promotor bent. Een toegankelijke hoogleraar met verstand van zaken en goede humor (Ik zeg: doen!). Ook al heb ik je advies om latex te gaan gebruiken altijd in de wind geslagen, gelukkig heb ik wel geluisterd naar je advies om mijn wiskundige model aan de hand van het boek van Sheffi op te stellen. Ik heb erg veel van je geleerd, waar ik in de toekomst nog veel aan zal hebben.

Ook ben ik blij dat Luc copromotor is geworden in de loop van mijn promotietraject. Dit gaf extra ruimte en diepgang in de discussies over het vervolg van het onderzoek, en bracht in sommige gevallen een welkome medestander in discussies. Om in zijn voetbaltermen te spreken, we hebben een leuk een-tweetje gemaakt: mijn afstudeeronderwerp ging over optimaliseren van een verkeersnetwerk, daar had Luc wat aan tijdens de start van zijn promotieonderzoek, en ik heb veel gehad aan zijn onderzoek en kennis toen ik vervolgens zelf met promotieonderzoek bezig was. En: zo'n overleg is toch ook gewoon een goede reden om een goede koffie te pakken met een flink stuk chocola erbij?

Op de universiteit die ik al zo goed kende uit mijn studietijd heb ik vele kamergenoten gehad, onder andere door een verhuizing van de vakgroep naar de Buitenhorst. Allen bedankt voor de gezelligheid: eerst Anthony en Malte, toen de grote kamer met Mohammed, Sander, Wouter en Jaap en later op een kleinere kamer alleen Sander, goed voor vele kilo's pepernoten. Dorette, bedankt voor je hulp bij allerlei praktische zaken. Ook Kasper wil ik bedanken voor de hulp bij ICT perikelen en voor het enthousiasme waarmee we samen een case hebben opgesteld binnen het vak public transport: een case die net als mijn onderzoek ging over het netwerkontwerpprobleem.

Mijn onderzoek was onderdeel van een breder onderzoeksproject: 'strategy towards sustainable and reliable multimodal transport in the Randstad'. Hierdoor kwam ik eenvoudig in contact met mede-onderzoekers en maakte kennis met de kijk op dit onderwerp vanuit

andere vakgebieden. Ik wil het onderzoeksteam bedanken voor de interessante discussies die we onderling gehad hebben: Ingo Hansen, Gijs van Eck, Rob van Nes, Daniel Sparing, Rob Goverde, Luca Bertolini, Andrew Switzer, Piet Rietveld, Jos van Ommeren en Yuval Kantor. Mijn dank gaat uit naar NWO, het programma 'Duurzame Bereikbaarheid van de Randstad', voor het financieren van dit onderzoek.

Ik wil ook graag de beleidsmakers bedanken die de tijd hebben genomen om met mij (en met andere leden van het SRMT project) in gesprek te gaan over onze onderzoeken of om gegevens beschikbaar te stellen. Als het goed is zien jullie sommige ideeën daaruit terug in het proefschrift. In het bijzonder waren dat Suzanne Kieft, Constance Winnips, Menko Noordegraaf, Rob Koning, Bart van der Heijden, Gerald Zwarthof en Wouter van Beek.

Ik wil graag mijn commissie bedanken voor de tijd die zij hebben gestoken in het lezen van mijn proefschrift. Ik ben erg benieuwd naar jullie vragen tijdens de verdediging: Karst Geurs, Johann Hurink, Bert van Wee, Jos van Ommeren en Rob van Nes.

Mijn collega's bij Goudappel Coffeng: bedankt voor jullie begrip als ik eens geen tijd had omdat ik verder moest werken aan mijn proefschrift. Vooral Niels van Oort wil ik hier noemen: hij had als geen ander begrip voor mijn situatie als part-time promovendus en heeft mij van vele goede adviezen voorzien.

Mijn paranimfen Sander en Gerard. Sander, jij bent de eerste student civiele techniek die mij onder ogen kwam, we zijn tegelijk de universiteit binnengekomen. Toen ik na anderhalf jaar terugkwam op de universiteit was jij er nog steeds. Het toeval (of niet?) wil dat we nu ook tegelijkertijd de universiteit weer verlaten. Je creativiteit en positiviteit kunnen erg goed van pas komen tijdens mijn verdediging. Gerard, jouw interesse in de techniek en onderzoekende basishouding is overgeslagen op mij, dan wel niet in de chemie maar in de civiele techniek. Toch heeft duurzaamheid tot jouw grote tevredenheid een belangrijke plek binnen mijn onderzoek. Jouw ratio en rust kunnen erg goed van pas komen tijdens mijn verdediging.

Anke en Gerard, jullie hebben mij uiteraard het langst gesteund, vanaf het begin van mijn leven. Ik heb erg veel goede dingen van jullie meegenomen, ik geloof dat ik uit een perfect nest kom. Jullie staan altijd voor mij klaar, zelfs tot aan de afgelopen drukke periode aan toe.

Annemiek, allereerst bedankt voor het controleren van mijn proefschrift op de Engelse taal. Maar vooral dank je wel voor je liefde en geduld tijdens de afgelopen jaren en voor het bieden van een gezellige thuiskomst na een lange dag werken.

Ties Brands, Deventer, september 2015.

Contents

Chapter 1: Introduction

In this first chapter the context, problem statement and research objective are presented. A motivation is given for the need to invest in a more sustainable passenger transportation network. This is also the objective of the larger Sustainable Accessibility of the Randstad (SAR) research programme of which this research is part. The aim of the SAR programme is to increase the sustainability of the transportation network in the Randstad, the main urban area in the Western part of the Netherlands. A promising solution approach is to stimulate multimodal trip making by designing a network that facilitates travellers shifting from less sustainable to more sustainable modes of transport. This research aims to investigate the extent to which (re)designing the multimodal transportation network can contribute to improving sustainability, while taking the behavioural responses of travellers into account. The research approach is presented in this chapter as well. Sustainability has several aspects that are expected to be conflicting, so the network design problem is formulated as a multi-objective optimisation problem, with various aspects of sustainability as objectives. Solving this problem results in a Pareto set of possibly optimal solutions, from which valuable information can be derived to assess the solution direction proposed. The approach is further specified by defining the research scope. Finally the contributions of the research are given, as well as an outline that explains the structure of the thesis.

1.1. Context

Transportation of passengers and goods enables people to carry out activities. These activities result in trips, which may be related to economic activities, like employees commuting to work or companies transporting goods from the production location to shops. These trips and related activities may also have non-monetary social welfare benefits, like visits to family and friends or recreational activities. In short, increasing mobility throughout the years has brought society many benefits.

On the other hand, increased mobility has had some negative effects on society as well. These negative effects are usually called (negative) externalities. Some of these externalities of mobility affect more people than the traveller alone. Therefore, the traveller often does not take these negative effects into account when making the choice to travel. These effects may consist of climate impact due to CO_2 emissions, traffic injuries and fatalities, bad air quality due to local pollutants, use of urban space by infrastructure and parking of vehicles and / or noise disturbance. Furthermore, the development and maintenance of transportation networks is a large financial burden for the tax payer. All benefits and costs can be classified into one of the three main aspects of sustainability, also referred to as people, planet, profit (Fisk, 2010). When applying this concept to the transportation system, the system should (Zuidgeest, 2005):

- provide cost-efficient transport services and infrastructure capacity, be financially affordable and support vibrant, sustainable economic activity (economic sustainability; profit);
- meet the basic human needs for health, comfort, convenience and safety, allow and support development of communities and provide a reasonable choice of transport services (social sustainability; people);
- have little or no impact on the integrity of ecosystems, use energy sources that are essentially renewable or inexhaustible, produce no more emissions and waste than the transport system's carrying capacity and produce no more noise than an acceptable threshold of noise pollution (environmental sustainability; planet).

This variety of effects illustrates that aiming for a sustainable transportation system cannot be captured in one singe objective, resulting in various objectives to be taken into account, each making an aspect of sustainability operational. The best would be to improve all aspects simultaneously. If this is not possible, the challenge is to find a balance between these three main aspects.

Some collective body (typically a government) can make the transportation system more sustainable by influencing the (positive and negative) impacts of mobility on society by taking measures. The governments have objectives that improve the benefits (and reduce the costs) for society as a whole, in terms of total benefits / costs and / or distribution of these benefits / costs over society. The current balance between aspects of sustainability and therefore the aspired change in this balance varies considerably between different regions in the world. In developing countries the main focus is on the role of the transportation system in economic development and less priority is given to for example traffic unsafety, saving the natural environment, limiting climate impact or equity (Zuidgeest, 2005). In the developed world, a higher economic standard has resulted in more attention to other aspects of sustainability

related to the transportation system. For example, in the Netherlands the number of traffic related fatalities has been decreasing over the last decades, as a result of strong policy (SWOV, 2009). However, in the Netherlands still challenges exist. Economic activities are harmed by congestion: in 2013 a total of 42.9 million vehicle loss hours occurred in the Netherlands (Van Veluwen and De Vries, 2014). The air quality in some parts of the Netherlands is below the standards set by European legislation (Zanten et al., 2013). The transportation sector (excluding aviation) was responsible for over 20% of the CO_2 emissions in the Netherlands in 2013 (CBS, 2013a). The CO_2 emission reduction targets for the European union are a 40% reduction (compared to 1990) for 2030 and an 80-95% reduction for 2050 (European Commission, 2014). It is likely that the transportation sector has to contribute substantially to reach these targets on the overall level.

Improving the network of public transport (PT) can contribute to these challenges by causing a modal shift from private cars to PT, mainly for medium-distance and long-distance trips. This shift is expected to improve social and environmental sustainability and at the same time could improve economic sustainability. These improvements may already be realised on a medium-term planning horizon (a few years from now), but can last for decades, so also a long-term planning horizon is needed. Although improving bicycle networks can contribute to meeting the challenges, also in relation to the introduction of electronic bicycles, bicycle trips are only an alternative for short trips. These trips cover only a limited part of all traffic flow: trips below 10 km have a share of 18% of all distance travelled in The Netherlands and 19% in the urbanised western part of the country (CBS, 2013b). This is the reason not to focus on the bicycle network. The environmental performance of the car is expected to improve in the future, but the progress of this transformation is unsure and is therefore no solution for the coming years. Moreover, the challenges in the Netherlands concerning the sustainability of the transportation system are expected to grow in the future: Huisman et al. (2013) forecast that the number of residents (and consequently the traffic flows) in the Netherlands will continue growing until 2040, especially in the large cities. This expected growth is another reason not to focus on improving the car network, because even if the environmental performance of cars will improve drastically (e.g. an electric car charged with renewable energy), cars still need a lot of space, both for driving (on infrastructure) and for parking. Therefore, a solution direction is chosen to improve the PT network. However, investments in PT infrastructure require large financial resources, while the budgets available to mitigate these problems have become smaller in recent years due to the worldwide economic crisis. This is not expected to change in the near future. Therefore the challenge is to select the most cost-efficient measures from all possibilities to improve the PT network.

In the Netherlands the densities of activities and traffic flows are the highest in the Randstad area, so that is where the challenges concerning the sustainability of the transportation system are most urgent. The Randstad is the urbanised area in the West of The Netherlands (see Figure 1.1) and includes the cities of Amsterdam, Rotterdam, Den Haag and Utrecht. The area currently has around 7 million inhabitants, making it one of the largest urban agglomerations in Europe. The densities of residents and jobs are expected to grow further because of re-urbanisation (Huisman et al., 2013). However, the situation in the Randstad is different from most other large urban regions in Europe. Firstly, the Randstad is polycentric, with large areas of open land between urban centres. Secondly, the residential areas in the suburbs of the various cities in the Randstad mainly consist of low-rise buildings, resulting in lower densities

of residents compared to suburbs in most other large urban areas throughout Europe. This spatial typology is one of the reasons for the share of PT in the number trips in the Randstad to be small compared to the share of the car: 7% for PT (train, bus, tram and metro) and 45% for car (CBS, 2013b). Another reason is typical for The Netherlands: the share of the bicycle is high (27% of all trips). The situation is different in for example the metropolitan areas of London and Paris: the shares in London are 45% for PT and 33% for car (Transport for London, 2014) and in the region of Île-de-France are 33% for PT and 37% for car (Tregouët, 2010), with much lower shares for the bicycle. So although desirable from environmental sustainability point of view, the present PT system does not appear to provide a sufficiently attractive alternative for many people in the Randstad.

Figure 1.1: The Randstad area, consisting of the four large cities Amsterdam, Rotterdam, Den Haag and Utrecht, surrounded by smaller towns and with open area in between

Given the specific spatial situation in the Randstad and the solution direction to enhance the transportation network without large investments, the existing infrastructure can be utilised better by facilitating an easy transfers, especially from private modes (bicycle and car) to PT modes (bus, tram, metro, train). This transfer from private to public modes needs to be facilitated, by providing both transfer locations and a PT network that is sufficiently attractive. The definition of a multimodal trip, as given by Van Nes (2002), is that within a single trip two or more different modes of transport are used, between which travellers have to make a transfer. This can stimulate the use of PT without the need for investment in large infrastructure, because the strengths of both types of modes are used: PT for fast, high capacity, long-distance connections and private modes for flexibility for the shorter distances. Possible measures include opening park-and-ride (P&R) facilities, local train stations, express train stations or PT service lines. Private modes can then be used to reach PT, removing a

barrier of a traveller to use PT. This holds both for bicycle (Martens, 2007; Pucher and Buehler, 2009) and for car (Bos et al., 2004; Hamer, 2010). Multimodality as a solution direction is also recognised by the Dutch national policy document on infrastructure and spatial planning (Ministry of Infrastructure and the Environment, 2012).

The next step is to determine what combination of measures related to multimodal trip making is the most effective in improving the various aspects of sustainability. The current transportation network is taken as a starting point: designing a transportation network from scratch is not realistic, because the existing infrastructure contains large investments from the past. The travellers respond to changes in the network by changing their travel behaviour, which determines the effectiveness of the measures. The geographical location of the measures needs to be determined as well, i.e. where the measures should be implemented in the existing transportation network. A network approach that includes the current transportation network is needed, because individual measures are influenced by the existing network and by each other. Some measures are related because of their locations in the network, for example a P&R facility may be more effective when the corresponding station is served by a PT service with a high frequency.

When such measures to improve the transportation network are planned, politicians make the decisions based on information provided by policy officers, i.e. government employees who support decision makers. These policy officers gather information by doing research themselves or by searching relevant results and information from academic research or from consultancy companies. In current planning practice the information on which the decision is based often consists of an assessment of a limited number of pre-defined network solutions (i.e. combinations of measures). The composition of these solutions is usually based on expert judgment, where each solution has a (quantitative or qualitative) score for multiple objectives. These multiple scores need to be combined to a final score per network solution. For that, multiple criteria decision making (MCDM) methods are used.

In current planning practice social cost-benefit analysis (SCBA) is often used as an MCDM method. In SCBA, several effects of a project are combined to one cost-benefit ratio or balance, by monetising and aggregating each effect. This can be seen as using one specific set of weight factors to combine all objectives into one overall score, whose aim is to maximise the social welfare effect of the project. The economic principle behind SCBA is that weight factors are derived from consumer preferences. These preferences result from revealed or stated preferences in choice situations presented to consumers. These values are reliable when the choices relate to daily choices that are easy to recognise by a respondent, like travel time savings, but become less reliable or do not exist at all when it comes to environmental and other effects that are not priced in markets or similar choice situations. Therefore alternative methods are needed, which contain more uncertainty (Schroten et al., 2014). According to Sager (2013), SCBA is subject to the comprehensiveness dilemma: a narrow SCBA makes good economic sense, but a comprehensive SCBA has uncertain economic content, because the methods used to derive weight factors are not based on consumer preferences, leaving the basic principle of SCBA. Although the principle behind SCBA is sound, this discussion illustrates that the weight factors used to prioritise the objectives are subject to debate, especially those related to environmental sustainability. This might result in neglecting or

emphasising environmental sustainability when too much emphasis is put on the final cost-benefit ratio during the decision making process.

Another issue in current planning practice is that the network solution that is selected as the best from the set of pre-defined solutions may not be the best overall. It can very well be that there exist better feasible solutions which have not been considered, because expert judgement was used to define the solutions. Therefore the best solution or solutions should be the result of an optimisation process, which is often done in academic literature. For this approach one or more objectives are quantified, resulting in measurable score(s). Most studies have maximum accessibility as a single objective, in some cases subject to constraints on externalities, such as an emission reduction target or a budget constraint (Farahani et al., 2013). This results in one optimal network solution, but it does not provide insight into the dependencies between objectives (the various aspects of sustainability), i.e. the extent to which the objectives are opposed or aligned. Moreover, no information is provided on the possibilities to improve the network further if the budget is slightly increased. Another common method is to combine a set of objectives using a weighted sum, where the weights represent the compensation principle between the objectives. However, setting these weights is not trivial (as illustrated before by the debate on weight factors to be used in SCBA). If the weights are determined in advance, uncertainty concerning these weighting factors is not incorporated and the sensitivity of the outcome to these factors is not known.

Given that sustainability cannot be captured sufficiently in one single (aggregated) objective and given the solution direction to stimulate multimodal trip making, there is a need for multi-objective optimisation of multimodal passenger transportation networks. This may contribute to improvements in all aspects of sustainability, but most likely it is not possible to optimise all aspects of sustainability simultaneously. In other words, environmental, economic and social sustainability are likely to be at least partly opposed (i.e. the optimal network design for economic sustainability is most likely not equal to the optimal network design for environmental sustainability). In policy documents various aspects of sustainability are often mentioned to be improved simultaneously. An example is the Dutch national policy document "Structuurvisie Infrastructuur en Ruimte" (Ministry of Infrastructure and the Environment, 2012). However, when it is not possible to improve all objectives simultaneously, choices have to be made. Insight into how the objectives relate and into the consequences of certain decisions is needed to enable decision makers to make a final decision. There has been little research that focuses specifically on several types of measures related to multimodal trip making, incorporating travellers' behaviour both on the car network and the PT network. This specific type of measures is combined with the need to optimise multiple objectives related to sustainability to come to a combination of measures to be implemented. It is especially a challenge to present the outcome of this optimisation procedure in a way that is easy to interpret for decision makers, and therefore can act as decision support information. Since these decisions are taken with a medium to long planning horizon, another relevant aspect is uncertainty concerning future developments. For example, in forecasts for socio-economic data, like the number of residents and households, large confidence intervals are used (Huisman et al., 2013). Therefore, the long-term robustness of the decision support information for these uncertain developments needs to be investigated.

1.1.1. Project context

This research is part of the second call of the Sustainable Accessibility of the Randstad (SAR) research programme (2008-2014). In the SAR programme a wide range of academic research groups in the Netherlands participates by doing research on what is needed to guarantee the accessibility of the Randstad in the long term (Verdus, 2014b). The SAR programme has been an initiative of the Dutch Ministry of Infrastructure and the Environment and the Dutch Ministry of Economic Affairs, Agriculture and Innovation. The programme has been administered by NWO (the Netherlands Organisation for Scientific Research). It focuses on the future, with a time horizon until 30 years from now. It concentrates on the internal and external accessibility of the Randstad, including its main ports and urban areas, with relation to people, goods and information flows and all within the broad framework of long-term trends, for example in economic development, demography, climate and energy.

One of the research projects within SAR is 'Strategy towards sustainable and reliable multimodal transport in de Randstad' (SRMT). In the project SRMT it is investigated how PT can act as a backbone of reliable transport chains and can contribute to a vital and accessible Randstad. The research presented in this thesis is project 3 out of the 5 research projects within SRMT and focuses on the design of multimodal passenger transportation networks. It has several relations with the other projects (in various disciplines) in SRMT (see Figure 1.2):

- In project 1 an equilibrium model addressing land use in a multimodal context is developed. This model can be used to determine the effects of multimodal network designs resulting from project 3 on urban land use from a spatial economic perspective. This changed land use can result in a changed transportation demand, which can again be input to come to new network designs, to finally come to an equilibrium between transportation network design and land use.

- In project 2 integrated transition strategies for the Randstad are developed, which can be used to enable realising the multimodal network designs resulting from project 3 in the complex policy context in the Randstad. The other way around, project 2 can indicate that certain measures are or are not feasible, with consequences for the definition of the design problem in project 3.

- In project 5 a timetable model that integrates rail and other PT is developed. This model can be used to further specify multimodal network designs resulting from project 3 in terms of timetables. These dynamic multimodal networks can then be assessed using the model developed in project 4. Furthermore, in project 5 a model to analyse the capacity utilisation of railways is developed, which can be used to test whether the rail services resulting from network designs are feasible given the current railway infrastructure.

- In project 4 a dynamic multimodal network assignment model is developed. This model can be used to assess multimodal network design resulting from projects 3 and 5 in more detail to determine the objective values with more accuracy.
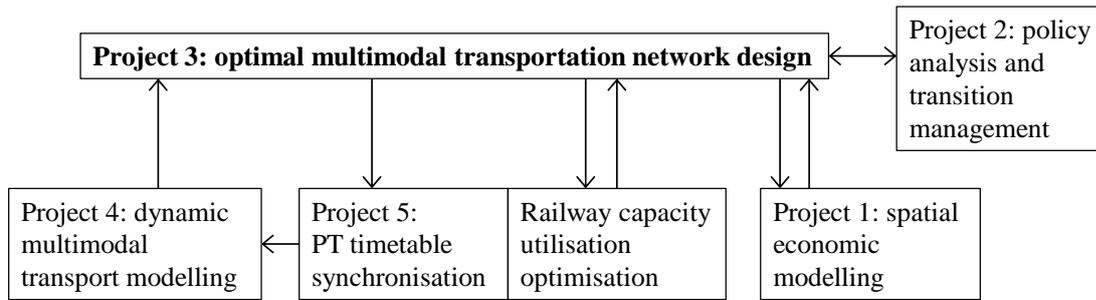
Figure 1.2: Relations between the five research projects within SRMT

1.2.   Problem statement and research objective

In Section 1.1 it was observed that transportation provides many benefits to society, but also comes with some negative externalities. Therefore the aim is to make the transportation system more sustainable, i.e. to improve all three aspects of sustainability: economic, social and environmental sustainability. Because medium-distance to long-distance car trips are responsible for a large share of the externalities of the present transportation system, shifting these trips from car to PT can contribute to this aim. The focus on the Randstad area (being a part of the SAR research programme) and the need for cost-efficient measures lead to stimulating multimodal trip making as a solution direction. Various measures in a multimodal passenger transportation network interact with the existing network and with each other, so the network needs to be (re)designed as a whole. Thereby, the behavioural response of travellers needs to be taken into account, to be able to correctly evaluate the effects of the measures. It is important to know how much improvement is possible using this type of measures, so more is needed than only choosing from a pre-defined set of network solutions. Given the fact that sustainability has multiple aspects that may not be in line with each other, it is very likely that no single network exists that is optimal for all aspects simultaneously. Instead, trade-off information is needed, so that depending on the priorities given to the various aspects, a decision can be made on which network is to be preferred. This leads to the following research objective:

The objective of this research is to determine which (re)designs of the multimodal passenger transportation network in the Randstad contribute best to improving various aspects of sustainability and to provide insight into the extent to which these aspects are improved and into how scores on these aspects and designs relate.

1.3.   Research approach

Designing a multimodal passenger transportation network is a specific example of a network design problem (NDP). Since multiple objectives are to be included to represent the various aspects of sustainability, more specifically the problem is formulated as a multi-objective NDP (MO-NDP). An MO optimisation problem has more than one objective to take into account explicitly. This enables the identification of trade-offs or compromises among conflicting objectives (Coello Coello, 2006). An MO-NDP typically involves determining the optimal values for a set of predefined decision variables, given certain constraints, by optimising (more than one) system objectives that depend on the behaviour of travellers in the

network. By defining the problem as an MO-NDP, a simplification of the real world is made: policy goals are quantified by objectives, decision variables represent measures in the network and a model predicts the behaviour of travellers in the network. Therefore, this definition of the problem is an important step.

Each aspect of sustainability may be put into operation in several ways and therefore not all objectives one can think of can be included. Furthermore, not all possible measures in the network are realistic within the context of an existing, heavily used urban region like the Randstad. Therefore, the objectives and decision variables / measures in the case study have been defined in cooperation with the stakeholders involved, resulting in the most important aspects to be taken into account explicitly. These stakeholders have been involved in the SRMT project in the form of a user group. In this user group policy officers from the Ministry of Infrastructure and the Environment and from provinces and city regions in the Randstad participated, as well as employees from the Dutch Railways (NS), from an urban PT operator and from a commercial property development company. Interaction with the user group was organised through meetings and workshops. This resulted in a definition of the MO-NDP for a case study in the Randstad, in terms of objectives, decision variables / measures and constraints.

In this MO-NDP each combination of measures represents one network design, also called a solution. Such a combination of measures corresponds to a set of values of decision variables. All possible combinations of values for the decision variables form the decision space. Each point or solution in decision space corresponds to a point in objective space, which consists of the values of the objective functions for that solution. When constraints are considered, which can be related to the values of decision variables or to objective values, only a part of solution space and / or only a part of objective space is considered to be feasible, resulting in the feasible set of solutions.

To be able to assess the objective values that correspond to a solution (i.e. to determine its scores or to evaluate the solution), the behaviour of travellers in the transportation network needs to be determined for a given transportation demand. Since it is not possible to observe this behaviour in the real world (the network designs do not exist yet), a transportation model is used to model this behaviour. The result of such a model are loads, speeds and travel times in the network, which can be used to calculate the objective values using effect models. The combination of a transportation model and effect models leads to a mapping of decision space to objective space: based on the values of the decision variables the objective values can be determined. It should be noted that it is not necessarily true that solutions which are close to each other in objective space are also close to each other in solution space (e.g. two totally different multimodal network designs may result in similar scores on the objective functions).

To solve the resulting MO-NDP, heuristics are needed, because the NDP is an NP-hard problem (Johnson et al., 1978). In addition to that, the use of a transportation model to determine objective values implies that a heuristic is needed that uses a limited number of function evaluations, because these transportation models are computationally expensive in realistic networks. To this end the performance of two promising heuristics from literature is compared. The combination of the resulting heuristic and a transportation model leads to an operational modelling framework, which is able to approximate the solution of the MO-NDP.
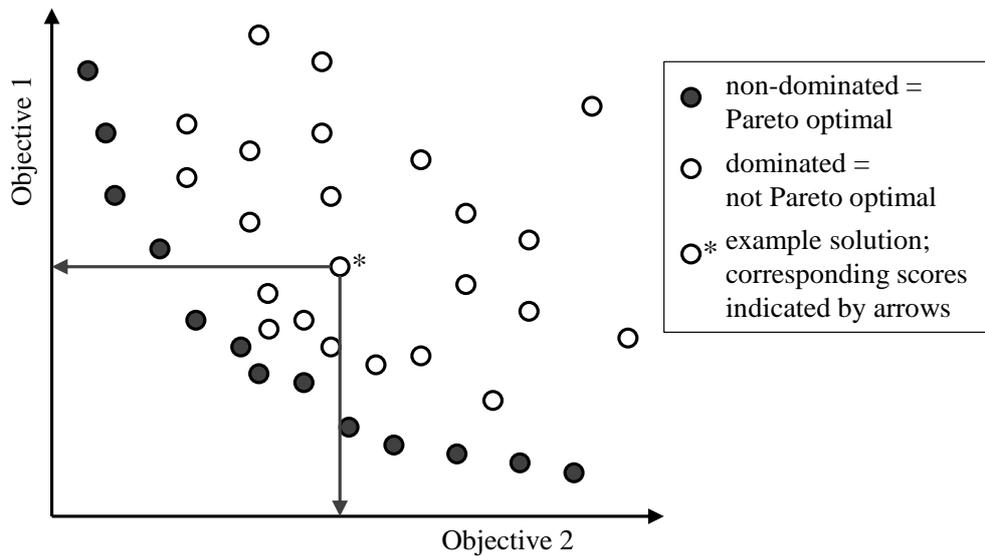
Figure 1.3: The concept of Pareto dominance

The outcome of MO optimisation is a set of solutions called Pareto set. This is different from the outcome of single objective optimisation, where generally one single optimal solution is found. The Pareto set consists of all solutions for which the corresponding objective values cannot be improved for one objective, without degradation of another. To determine the Pareto set, the concept of Pareto dominance is important (see Figure 1.3). Each solution is represented by a dot in the graph and has a score for 2 objectives. For an example solution (indicated by*) this is demonstrated by the two arrows to the two axes. In this example, both objectives are to be minimised. In the graph this means that moving to the lower left direction is desired. A solution that has a better (in this example lower) score for all (in this case two) objectives than another solution is said to dominate that solution: when all objectives improve, it is sure that this solution is better. In the example, all solutions in the lower left rectangular area (bordered by the two arrows) dominate the highlighted solution. Given this dominance relation, for the filled solutions in the graph no solution exists in the lower left direction: these solutions are non-dominated. All these solutions may be optimal in the end, depending on the preference a decision maker has for each objective. The set of these solutions is called Pareto-optimal set (or shorter: Pareto set) and the objective values of these solutions are called Pareto-optimal front (or shorter: Pareto front).

The resulting Pareto set contains valuable information which makes it possible to address issues like the level in which the objectives are conflicting or not and what kind of measures can be used to improve objectives (Wismans, 2012). This implies that the Pareto set contains knowledge on design principles when designing a multimodal passenger transportation network to optimise various aspects of sustainability. This knowledge may be revealed when analysing the Pareto set. It is important to present this information in such a way that it is easy to interpret for decision makers, because only in that case insight is provided. Analysing the Pareto set is also of interest when choosing one final compromise solution to implement. Choosing a compromise solution is related to MCDM in which the best solution is chosen considering multiple objectives (Wismans, 2012). The Pareto set can be used as input for a powerful, interactive decision tool, allowing the decision makers to learn more about the

problem before committing to a final decision. Analysis of the Pareto set and this choice is rarely addressed in MO-NDP literature, but necessary to select a combination of measures in the end. Also in this phase of the research stakeholders in the decision making process have been consulted: the preliminary results have been presented to the stakeholders, resulting in detailed feedback and discussion in an interview setting.

As mentioned before, a transportation model is used to come to the mapping between decision space and solution space. In this transportation model a transportation demand is taken as input data. However, this transportation demand in the future is a result of a demand prediction that contains uncertainty, mainly because this demand prediction is based on one specific scenario that represents the future socio-economic conditions, like numbers of residents and jobs. A different demand input for the transportation model results in a different mapping between decision and objective space, and therefore most likely to different optimisation results, while a decision maker likes to make a decision that is robust for future developments like the development of transportation demand. To investigate the impact of demand uncertainty, the modelling framework is run for several different demand input scenarios and the similarities and differences between the results are identified.

1.3.1. Research questions

The final research objective as formulated in the problem statement, is addressed by taking several steps, as outlined in the research approach. In these steps, several challenges are faced. These challenges are summarised in the form of the following research questions, which correspond to chapters (indicated between brackets):

- How can the problem be modelled mathematically? (Chapter 2 and Chapter 3)
  - Which aspects of sustainability are the most important and how can they be put into operation in objective functions?
  - Which decision variables should be included in the context of a multimodal passenger transportation network and how can they be defined?
  - How can the behavioural response of travellers be formulated and be made operational in a modelling environment, in order that a good trade-off between computation time and model accuracy is reached?
- Which solution approach is suitable for this multi-objective optimisation problem? (Chapter 4)
  - What solution method provides high quality solutions (i.e. a Pareto set) within a limited number of function evaluations, in the context of the necessity to model the behavioural response of travellers (that requires considerable computation time)?
  - Does the chosen solution approach have any implications on how to interpret the optimisation outcomes when used for decision support?
- How can the Pareto set be used to provide insight into the best performing network designs (in the form of decision support information)? (Chapter 5)
  - What methods (existing methods, methods to be enhanced or methods to be developed) make the Pareto set easier to understand and more useful to guide the decision maker to come to a final solution for implementation?
  - What information on the inherent structure of the problem can be derived from the Pareto set resulting from the case study?

  o How are the scores achieved by (re)designing the multimodal network
     compared to scores that may be achieved by different measures (that are out of
     the scope of the NDP)?
- To what extent are the resulting network designs robust for demand uncertainty due to
  various possible future socioeconomic conditions? (Chapter 6)

1.4.   Research scope

In this section the research scope is defined, i.e. the research is delineated. This is necessary,
because the topic of multimodal passenger transportation network design is very broad: it can
have many different focuses and it can be put into operation in many ways. The scope is made
clear by defining what will be included in the research and what will be excluded.

1.4.1.   Types of measures

Following from the context as described in Section 1.1, the types of measures involved are
related to developments in the multimodal passenger transportation network. Such a
multimodal network consists of several elements, like roads, intersections, bicycle paths,
railway infrastructure, PT service lines, stations and stops. The existing transportation
network is taken as a starting point. No measures are included to expand the car network or to
expand the bicycle network. Large-scale investments in the PT network are not included
either, because the coming years no large investment budgets are expected to be available. As
a result, a majority of the elements in the multimodal network is fixed, but a selection of
elements related to PT are decision variables during optimisation. All decision variables aim
to make the PT choice alternative more attractive for the traveller. First, local train stations
can be opened or closed: opening a station provides opportunities for travellers to or from the
station area, also because the bicycle can be used as access or egress mode, but travel time for
through passengers increases. Second, the status of existing stations can be changed from
express (express trains call at the station) to local (no express trains call at the station) or vice
versa, with similar effects. Third, frequencies of local train lines and frequencies of major bus
lines can be increased to reduce waiting times. Finally, P&R facilities can enable travellers to
combine car and PT, for example to avoid congestion in an urban area but still being able to
use the car as access mode in a rural area.

The PT service lines in the network are represented by their route over the infrastructure, their
travel times between stops and their frequency. P&R facilities are represented by indicating
for each PT stop whether it is or is not allowed to use the car as access mode. All decision
variables are captured in this representation. This level of detail corresponds to the long-term
planning horizon, where a strategic insight into the network designs is desired. This implies
that the timetable is not designed and the capacity of railway infrastructure is not checked for
the proposed services. Designing an optimal timetable for a set of heavy rail services on a
given network of railway tracks is an optimisation problem in itself, which is dealt with in
project 5 of SRMT (Sparing and Goverde, 2013).

When a distinction is made between so called push and pull measures, these measures are
examples of pull measures. Pull measures make the desired alternative more attractive and are
usually more easily accepted by the public. On the contrary, push measures make the

undesired alternative less attractive and are usually less easily accepted by the public (Soderholm, 2013).

A range of alternative measures is available to contribute to sustainability of the transportation system. Without being complete, a range of alternative (types of) measures is mentioned here. Although of interest, these measures are not included in the definition of the optimisation problem in this research, also because many other approaches are covered by other projects in the SAR programme. First of all, this research focuses on passenger transportation, so measures to influence freight traffic are not considered: designing a freight transportation network in a multimodal context is studied for example in Zhang (2013). This research further focuses on the use of the transportation system, rather than on technological measures to improve the environmental performance of vehicles. This could for example be the shift to emission free vehicles, which may be cars, PT or other types of vehicles, for example electric vehicles (Sierzchula, 2015) or hydrogen fuelled vehicles (Huijts, 2013). Furthermore, the following solution directions are excluded from this research, which may also contribute to an increased use of the PT network. The first alternative type of measures is changing pricing policies, which can be put into practice both as pull and as push measures. These policies can for example consist of road pricing, fares in PT or parking charges. Pricing policy as a solution direction is investigated in the I-Prism project (Verdus, 2014a). Other opportunities are provided by ICT developments, for example the possibility for teleworking, which is investigated in the project Synchronising network (Verdus, 2015a). The last mentioned alternative approach is to make more efficient use of existing infrastructure, for example by dynamic timetabling on the rail network (Wang, 2014) or by providing information to travellers, which is investigated in the TRISTAN project (Verdus, 2015b).

1.4.2. Behavioural response

As described in Section 1.4.1, measures are proposed to stimulate multimodal trip making. These measures will cause changes in the costs of choice alternatives in the multimodal network, resulting in changing behaviour of the travellers in the network. To find the best solutions, the modelling framework anticipate this changing behaviour by using a model to predict the behavioural response of travellers to changes in the multimodal transportation network. The behavioural model should include the possibility to combine private and public modes to mode chains (for example using car or bicycle to reach a train station), next to the traditional unimodal choice options (car-only trips or PT trips that only involve walking as access and egress modes). The proposed measures typically change the attractiveness of certain routes and modes, where mode chains are seen as separate modes. Therefore, it is important to include mode and route choice in the behavioural response model. On the other hand, total transportation demand is assumed to be fixed, because the measures are relatively small-scale: they do not provide completely new travel options. From the proposed measures providing train stations with express train status and introducing new PT service lines may in the long run attract additional activities and therefore also have distribution or generation effects. A P&R facility may enable a traveller to travel in the peak hour, when the traveller used to avoid congestion by driving off peak, influencing departure time choice. However, it is assumed that the main effects of the proposed measures are due to changed route and mode choice and not due to changed trip distribution, trip generation or departure time choice. Furthermore, a static model is used, because the infrastructural measures are not time-

dependent: when such a measure is taken, it can usually be used all day. Frequency of PT service lines can be higher during peak periods than during day time, but usually lines with a high frequency during the peak periods also have a high frequency during day time. Although a dynamic model would give more insight into the network dynamics in terms of congestion and the effect of scheduled departure times of PT vehicles, this comes with much larger computation times. A static model has enough quality to assess the effects of the proposed measures, against a fraction of the computation time of a dynamic model. Finally, because the focus is on measures related to PT, the car modelling component can be relatively simple, e.g. without junction modelling.

A combination of existing modelling techniques is used to come to a suitable model. Development of new modelling techniques is outside the scope of this research. Furthermore, no new experimental or real-life data have been gathered to improve existing models. Parallel to this research, in project 4 of SRMT, a dynamic assignment model for multimodal transportation networks is developed, where the choice among the full range of modes and mode combinations is modelled as a single choice (Van Eck et al., 2014).

To be able to model multimodal trip making in a realistic way (so including multiple access and egress modes), it is necessary to use a detailed PT network and a network for access and egress modes, in addition to a car network. To control computation time, the size of the study area is limited (to a part of the Randstad area instead of the whole Randstad) and the number of zones will be limited. Furthermore, for the study area a transportation model with a proper network needs to be available.

1.4.3. Aspects of sustainability: objectives

To capture the multi-facetted nature of sustainability at least one objective is chosen for each main aspect of sustainability (economic, social and environmental). Each of these main aspects of sustainability can be made operational in many ways. On the other hand, the more objectives are included, the more difficult the optimisation problem becomes. Also, analysing and interpreting the results becomes more complex. Furthermore, considering more (opposed) objectives probably means that a larger part of the feasible solutions is Pareto-optimal. For example, in a different case study for the MO-NDP Wismans (2012) showed that when increasing the number of objectives from 3 to 5, the share of Pareto-optimal solutions in the total set of assessed solutions during optimisation increases from 10% to 16%. As a result, the objectives should cover all aspects, but the number of objectives should be as low as possible.

An important condition for an objective to be included is that it is possible to quantify the objective based on the network data resulting from the transportation model used to model the behavioural response. Furthermore, for the optimisation process only objectives that are at least partly opposed are of interest: when two objectives are completely aligned one of the two can be omitted during optimisation, because these objectives result in the same solutions being optimal. The exact number of objectives and which objectives to include depends on the priorities given to these objectives by the stakeholders.

1.4.4. Transportation demand

Transportation demand for an average working day is used as input for the modelling framework. This means that day-to-day variability is not taken into account. Furthermore, only a one-hour period within the AM peak during such an average working day is modelled, because the AM peak is the busiest period of the day, so the transportation-related problems are most relevant during this period. Despite the notion that the impact of network designs also depends on other periods of the day and week, it is assumed that the main effects of the measures are captured in the results based on the AM peak period.

Another type of variation is the long-term development of transportation demand. For future transportation demand, existing forecasts are used, which are based on the Dutch national scenarios for economical and spatial developments, which are called WLO-scenarios (CPB et al., 2006). This expected geographical distribution of activities is incorporated in the national and regional transportation models in the Netherlands, from which the demand is derived. Because both the forecasts for economical and spatial developments and the demand that is derived from these forecasts contain uncertainty, the long-term robustness for different transportation demand developments is tested in the robustness analysis in Chapter 6.

1.5.   Contributions

The research that will be presented in the next chapters has several contributions to the existing literature.

Solution algorithms

The performance of two solution algorithms is analysed for the practical case study in multimodal passenger transportation network design. Literature provides little evidence on the performance of any of the available heuristics in this context, even though this information is needed, because every result is case-specific. The consequences of using a heuristic as a solution algorithm for the outcome of the optimisation process are also addressed in a novel way.

Effects of uncertainty in the demand input on optimisation outcomes

The transportation demand is important input for the MO-NDP. Analysis of the effects of uncertainty in the demand input on the outcomes of an optimisation problem is new for the multi-objective case. New indicators are developed to cope with the additional difficulty to compare sets of solutions instead of just single solutions.

Analysis of the results from multi-objective optimisation: decision support

A lot of literature exists on solving multi-objective problems, using a range of available solutions algorithms. However, most literature is devoted to methodology development: applications to real-world problems are limited. In those cases where practical applications are included in academic work, the optimisation outcome is mostly simply presented in the form of a Pareto set. In this research one step further is taken. Methods are (further) developed and applied that help the decision maker to choose a final solution for implementation based on the Pareto set.

Results concerning multimodal network design in the Randstad case study

The NDP is defined in such a way that is fits the circumstances of the Randstad. This specific combination of decision variables, objective functions and behavioural response is therefore unique. Application of the above-mentioned methods for decision support results in general problem knowledge, derived from the Pareto set. This provides insight into how objectives relate, what kind of measures related to multimodal trip making can be used to optimise certain objectives and the consequences of a certain network design in the Randstad case study.

1.6.   Outline

The structure of the thesis is summarised in Figure 1.4. After this introduction, more background information is provided in Chapter 2. First an overview is given of earlier research on the NDP, with focus on problems in the field of multimodal passenger transportation networks and on multi-objective problems. The second topic is modelling the behaviour of travellers in a multimodal network. Based on a set of model requirements, an overview of available techniques is provided to cope with the difficulties that arise when a multimodal network is studied. Together with the scope defined in this introduction, the background information leads to the formulation of a mathematical optimisation problem in Chapter 3. This includes the definition of objective functions, a translation of potential measures to decision variables and a model for combined mode and route choice, which includes multimodal trip making. This model is used to determine the behavioural responses and therewith to assess the network designs in terms of objective values. Furthermore, the case study in the Randstad is introduced and defined in detail. The solution approach in Chapter 4 contains two possible heuristics that are used to construct an approximation of the Pareto front. The performance of the two algorithms when applied to the case study is compared, using performance indicators. Besides, the fact that a heuristic is used is given attention, by testing the amount of variation in the optimisation outcomes that is caused by the method itself. Application of the solution approach to the case study in the Randstad leads to optimisation outcomes in Chapter 5. Since these optimisation results are too complicated to be interpreted directly by decision makers, this involves several methods to derive problem knowledge from the Pareto set and to reduce the number of solutions in the set. The outcomes of the methods provide insight into the extent to which measures that enable multimodal trip making can contribute to the sustainability objectives in the study area (either simultaneously or by making choices among objectives) and which (types of) measures can contribute to these objectives. The long-term robustness of the optimisation results is tested in Chapter 6, by investigating to what extent different scenarios for future transportation demand influence the optimisation results. The latter three chapters lead to conclusions to be drawn in Chapter 7. Furthermore, research directions for further research are identified.
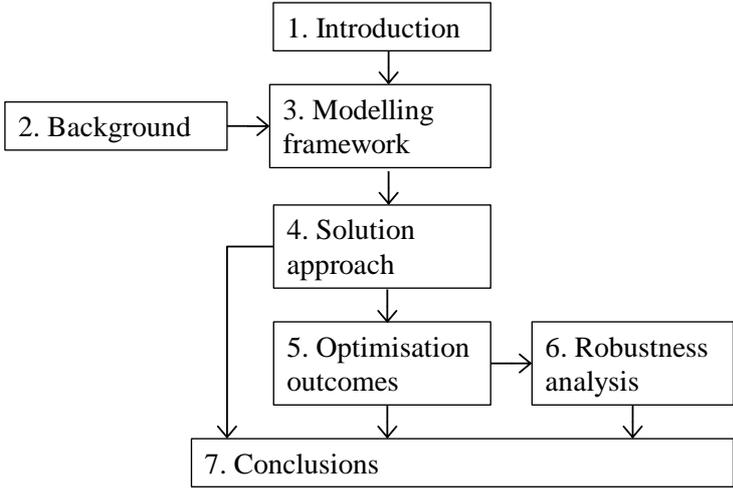
Figure 1.4: Structure of the thesis

Chapter 2: Background

In Chapter 1 the objective of this research was defined, including the context of the research and the research scope. This chapter provides literature research, giving the necessary background information that leads to the detailed formulation of the optimisation problem in Chapter 3.

This thesis employs mathematical optimisation techniques to find better transportation networks. In literature, this approach is referred to as the network design problem. More specifically, this thesis employs the design of a public transport (PT) network in relation to the multimodal passenger transportation network: adding or removing multimodal facilities (e.g. park-and-ride facilities and train stations) and adjusting PT lines. These kind of network adjustments are chosen to make the transfer from one mode to another easier, or, in other words, to enable multimodal trips.

In section 2.2 the reader is provided with literature on the topic of network design problems, focusing on multi-objective problems and on applications concerning multimodal networks. It appears that the combination of both is rare. In Section 2.3 attention is given to methods to cope with multimodal transportation networks. When behaviour of people who travel through such a network is modelled, networks of different modes need to be connected (forming a multimodal network), in such a way that a modelled traveller is able to combine more than one mode in a mode chain, resulting in a multimodal trip. This leads to a choice for a method that makes best use of existing modelling techniques and is efficient in terms of computation time in Section 2.4.

2.1. Introduction

The network design problem (NDP) aims to find an optimal network configuration, given a single or multiple objectives. It has many applications in transportation, but may also be applied to other types of networks, for example electricity grids or gas pipe networks. The multimodal passenger transportation NDP aims to find the optimal values for certain predefined decision variables, which optimise various system objectives, taking users' travel behaviour into account.

The system objectives represent the policy goals of the authority. In the case of transportation networks, the authority is usually a government, who have goals to improve society as a whole. A system objective can be any indicator that can be measured in the network, but in the context of the multimodal passenger transportation NDP these system objectives are likely to be related to one of the main aspects of sustainability (economic sustainability, social sustainability or environmental sustainability, see Chapter 1).

The predefined decision variables can be any measure related to the supply of infrastructure or services in the multimodal network, for example road or rail infrastructure, park-and-ride (P&R) facilities, frequencies of PT services, speeds of PT services, speed restrictions for roads, fare of PT or pricing of road links. Constraints may relate to availability of physical space to build infrastructure, legislation or budget. Costs of investment or operation may be included in such a constraint, or taken into account by including it as a system objective.

The transportation NDP is often solved as a bi-level optimisation problem (Farahani et al., 2013), to correctly incorporate the reaction of the transportation system users to network changes, as is argued by dell'Olio et al. (2006) and Tahmasseby (2009), see Figure 2.1. The upper level represents the behaviour of the network authority, optimising system objectives. In the lower level the travellers minimise their own generalised costs (e.g. travel time, cost), by making individually optimal choices in the multimodal network, considering variety in travel preferences among travellers. The network design in the upper level interacts with the behaviour of the travellers in the network: the lower level. The lower level is a constraint for the upper level, since the upper level cannot dictate the behaviour of the users in the lower level. Any network design the network authority chooses, results in a network state (e.g. travel times and flows), from which the system objectives can be derived. The bi-level linear programming problem is NP-hard (Gao et al., 2005), so any bi-level problem is NP-hard as well. Therefore, heuristics are needed to solve the bi-level NDP for larger networks. The huge number of feasible solutions and the non-convexity of the objective function necessarily requires the adoption of metaheuristic algorithms (Wismans et al., 2012).
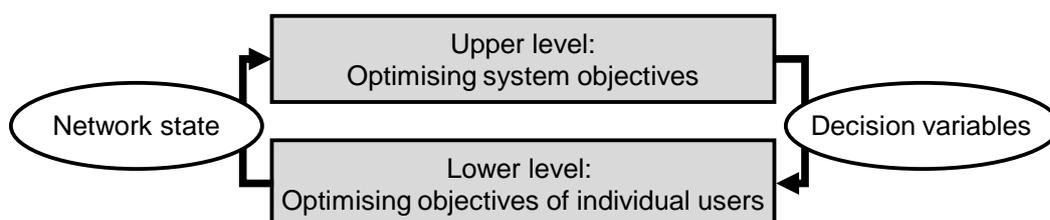
Figure 2.1: The bi-level optimisation problem

2.2.   Earlier research on NDPs

NDPs have received a lot of attention in literature, in many different versions. The basic version of the problem is the subclass of unimodal road NDPs, which has been studied for a long time, as reviewed by Yang and Bell (1998). They identify finding a global optimal search algorithm that can guarantee optimality in a computationally efficient manner and application to realistic road network examples as most important challenges. A more recent literature review of road NDPs (Wismans, 2012), shows that important steps are taken with respect to both challenges, but research on search algorithms has shifted from exact algorithms to heuristics. Wismans (2012) identifies several problem types, making distinctions between discrete, continuous and mixed decision variables, between fixed, stochastic and elastic transportation demand and between static and dynamic modelling of the traffic system. Further, attention is given to the distinction between single objective (SO) NDPs (one objective function is optimised, resulting in a single optimal solution) and multi-objective (MO) NDPs (more than one objective function is optimised simultaneously, resulting in multiple possibly optimal solutions, a Pareto set). It is found that in most MO cases where different solution algorithms are tested and compared, genetic algorithms (GAs) outperformed the other algorithms. Furthermore, GAs are often used in other research on various definitions of the NDP.

Another subclass of problems is the PT or transit NDP, which has been studied in various ways, as reviewed by Guihaire and Hao (2008). Their review confirms that the PT NDP in general is computationally intractable, preventing to guarantee overall optimality. However, the development of metaheuristic methods have made it possible to tackle large size problems more efficiently. These methods are flexible to adapt to any type of constraints and objectives and therefore can be applied to almost any PT planning problem. Furthermore, Guihaire and Hao (2008) identify a need for more realistic applications, since most studies in the review focus on theoretical problems. They also state that multiple path assignment and PT demand responsiveness are important to achieve a sufficient level of realism, because these aspects influence the objective function and planning results.

Since this thesis involves a multi-objective, multimodal passenger transportation NDP, the literature provided in the chapter focuses on studies involving more than one mode and / or more than one objective (see Table 2.1). Hereby, the definition of decision variables and the operationalisation of the lower-level model representing behaviour in the transportation network are important. Furthermore, the solution method is relevant, which has a relation with the size of the case study, since large computation times are expected. As mentioned, the objective of this thesis is to provide insight into how the transportation network of the Randstad can become more sustainable, so the method should be suitable for a realistic case study. Each of these aspects is given attention in this section.

The most important observation in Table 2.1 is that only Miandoabchi et al. (2012) apply MO optimisation to a multimodal NDP. They consider both new street construction / lane additions and redesign of bus routes to be decision variables. The aim is to efficiently invest the available budget and therefore investigate to what extent to improve roads or to improve bus routes in an integrated NDP. In the lower-level model, car and bus are distinguished as separate modes, so the traveller's choice is between either car or bus, so combining these

modes into multimodal trips is considered to be impossible. The focus of the research is on the development of the metaheuristics to solve the problem, where the GA is designed in such a way that the PT line routes are directly coded by the genetic string, so this result is only proofed to be valid for this specific problem definition.

Table 2.1: Examples of multimodal and / or multi-objective NDPs

Reference	Decision variables	Lower-level model	Objective(s)	Solution algorithm	Case study
Miandoabchi et al. (2012)	Road link capacity and bus routes	Mode choice between car and bus and route choice	User benefits; passenger share of the bus mode; service overage; average generalised travel cost for a bus mode trip	Hybrid GA; hybrid clonal selection algorithm	7 test networks of 9-40 nodes and 12-66 links, containing 2-5 bus lines
Hamdouch et al. (2007)	Pricing of private and public links	Mode choice between car and car combined with metro and route choice	Total travel time	Analytic	Test network of 11 nodes, 20 links, containing 4 metro lines
García and Marín (2002)	Parking capacity and parking pricing	Mode choice between car, metro and P&R and route choice	Total travel cost	Simulated annealing	4 test problems of 13-1002 nodes, 44-1678 links, containing 3-47 metro lines
Uchida et al. (2007)	Frequency of PT lines	Probit-based using a hyper-network; choice between all possible combinations of private and public modes	Sum of the total disutility of passengers and costs caused by the frequency setting	A local linear approximation of implicit function programming	One test hypernetwork with 11 nodes, 18 links and 4 PT lines
Chen et al. (2010)	Road links	Route choice	Travel time; construction costs	GA	Sioux Falls network (24 nodes, 75 links)
Cantarella and Vitetta (2006)	Urban road links	Elastic demand and route choice	Travel time; CO emissions	Heuristic multi-criteria technique based on GAs	One test network with 26 nodes and 34 links
Yin and Lawphongpanich (2006)	Congestion pricing scheme	Route choice	Travel time; CO emissions	GA	Test network with 6 nodes and 7 links
Sumalee et al. (2009)	Road user charging scheme	Route choice	Social welfare improvement; revenue generation; equity (Gini coefficient)	Dynamic self-adaptive penalty GA; Non-dominated sorting GA II (NSGA-II)	Case study in the city of Edinburgh, 350 links and 25 zones

Sharma et al. (2009)	Road link capacity	Route choice	Travel time; travel time variance	NSGA-II	Nguyen Dupius test network; Mumbai in India, 17 nodes, 56 links
Wismans et al. (2011)	Dynamic traffic management measures	Dynamic route choice	Noise; climate; congestion	NSGA-II, the strength Pareto evolutionary algorithm 2 (SPEA2), and (SPEA2+)	Urban test network with 3 routes
Mauttone and Urquhart (2009)	PT routes and their frequencies	Route choice in PT	Minimisation of operation hours; Minimisation of total travel time for PT users	GRASP (Greedy Randomised Adaptive Search Procedure)	2 test cases: Mandl (15 nodes, 21 links) and Rivera (84 nodes, 143 links)
Fan and Machemehl (2006)	PT route network design	Route choice in PT, with explicit choice between shortwalk and longwalk choice options	Aggregated costs (user costs; operator costs; unsatisfied demand costs)	Simulated annealing	1 test network with 95 zones, 160 nodes and 418 links (PT lines to be designed)
Beltran et al. (2009)	Bus routes and frequencies	Mode choice and route choice	Aggregated costs (operator's costs; user costs; external costs)	GA	1 test network, with 21 zones, 49 nodes, 90 links and 32 existing bus lines
Alt and Weidmann (2011)	PT routes, headway and vehicle size	Route choice: headway-based stochastic multiple routing	Aggregated costs (passenger costs; service provider costs)	Ant colony optimisation; GA	Mandl's test problem and case study in Winterthur
Gallo et al. (2011)	Frequency of rail lines (2 lines in small case and 14 lines in real case)	Mode choice between car and PT and route choice	Aggregated costs (operator costs, (rail and bus); PT user costs; car user costs; external costs)	Scatter search; GA	Test network (8 zones, 13 nodes, 32 links, 5 PT lines); real network in Campania, Italy (91 zones, 305 nodes, 862 links, 61 PT lines)
Cipriani et al. (2012)	PT routes and frequencies	Elastic demand and route choice	Aggregated costs (operator's costs; users' costs; external costs)	Heuristic route generation algorithm; GA	Real case study (450 zones, 1,300 nodes, 7,000 links, 200 bus lines)
Chakroborty (2003)	PT routing and the PT scheduling	PT route choice	Average travel time	GA	Mandl's test problem
Zhao and Zeng (2006)	PT network layout and headway	PT route choice	Total user cost	Combined GA and simulated annealing method	Mandl's test problem; Case study around Miami (2804 nodes, 4300 links, 81 bus lines)

Size of case studies in literature

In Table 2.1 it can be observed that several NDP studies include applications to real case studies of considerable size, instead of demonstrating optimisation techniques by tests on small artificial networks. Alt and Weidmann (2011) present a case study in Winterthur, Switzerland that is of realistic size. Another similarity is the availablitiy of a regional transportation model that enables detailed modelling of both the car and the PT network. However, the paper does not present any results concerning the outcome of the optimisation problem in terms of values of decision variables: the actual PT network that is designed is not presented. Other real-life applications (Gallo et al., 2011; Cipriani et al., 2012) do present optimisation results in terms of decision variables. However, in all these real-life applications car and PT are seen as competitors rather than modes that can be combined to one multimodal transportation network and all costs are aggregated into one single objective function. Since application on real networks exist for the PT NDP, it is expected that a real case study (on the scale of the Randstad area) is possible.

Separation between modes

In most studies in Table 2.1 that contain more than one mode, these modes are completely separated and only used to determine mode choice. Only Uchida et al. (2007) use an advanced lower-level model, where a single hypernetwork is constructed from the networks per mode. This allows for multimodal trips to be made, i.e. apart from unimodal trips by car or by PT, a traveller may change from car to PT or vice versa. For the available multimodal routes, a probit stochastic user equilibrium condition holds, which leads to considerable computation times. The formulation allows for extending to more modes, for example to include the bicycle. However, the NDP is given limited attention: it is addressed only on a primitive network with the frequency of 4 PT lines as decision variables. It is noted that for the network under study the same optimal solution is found for multiple starting points, indicating a stable result, but this gives no guarantee for other cases. Also García and Marín (2002) allow for P&R as an option in mode choice, but their approach is less flexible, since a combination of modes should be predefined as a mode chain, resulting in a mode choice model that distinguishes between P&R, car and metro. An important advantage of this simplification is that it allows for application of the method to a larger test network.

Solution approaches

Concerning solution approaches, almost all studies use heuristics, where (variants of) GAs are used most. As mentioned earlier for Miandoabchi et al. (2012), algorithms are usually designed for a specific problem. This is most relevant when PT routes are decision variables, since one route can already imply large combinatorial complexity. Mauttone and Urquhart (2009) solve this problem (i.e. find a local optimum) by applying GRASP (Greedy Randomised Adaptive Search Procedure), which first serves the heaviest OD (origin-destination) pairs directly. Consecutively, PT routes between OD pairs are added until every OD pair is served according to the constraints: a certain fraction of the demand should be served directly and a certain fraction of the demand should be served with maximally one change. Cipriani et al. (2012) cope with this problem by applying a heuristic route generation algorithm (HRGA), which first generates a large and rational set of feasible PT routes, by applying different design criteria and practical rules. Thereafter, a GA finds an optimal (or

near-optimal) sub-set of PT routes and their frequencies. This is comparable with Beltran (2009), in which a GA is used to optimise a bus network by choosing a fixed number of bus routes from a route set. This route set is generated beforehand and contains direct routes between important OD pairs (travellers' perspective), routes that are part of a hierarchical structure (service provider's perspective: main and feeder lines) and finally routes from the present network. The frequencies are determined as a part of the lower-level model, which describes modal split and route choice for car (including congestion) and PT travellers. Using a greedy algorithm to solve the problem or to limit the solution space in these cases worked well to keep the computation time within limits.

Furthermore, some comparison studies are relevant to identify promising solution approaches. Gallo et al. (2011) compare scatter search and a GA, where scatter search achieved better results, because it uses the shape of the objective space to find a local minimum, but the authors indicate that this result is case-specific. Furthermore, scatter search needed more computation time, which may lead to problems in a larger case study. Wismans et al. (2011) compare three evolutionary algorithms, where in the same computation time, better results are achieved by NSGA-II and SPEA2+. Sharma et al. (2009) provide another example of a succesfull implementation of NSGA-II in the field of transportation network design. Finally, Miandoabchi et al. (2012) compare the performance of a hybrid GA and a hybrid clonal selection algorithm, where the GA performs best.

2.2.1. Conclusion

Literature does not provide an approach that can be used directly to solve the problem at hand. However, there are elements that can be used. Examples of NDPs applied to real-life case studies can be found. Although computation times are considerable in these studies, based on these studies it can be concluded that it is feasible to set up an NDP case study for (a part of) the Randstad area in the Netherlands, as long as the lower-level model is kept computationally lean enough. An NDP with explicit modelling of multimodal trips in the lower level (with all possible mode combinations) is only found once, where high computation times in the lower-level model lead to an application that is limited to a small test network. Another NDP study showed that predefining mode combinations in mode chains helps to make the lower-level model faster and enables application to a larger case study. This shows that predefining mode chains is a promising approach to include multimodal trips, while maintaining a feasible computation time. GAs are mostly used as a solution algorithm for the upper level and comparison studies showed that (specific variants of) GAs perform best when solving MO or multimodal NDPs. Since GAs need a considerable number of function evaluations, the computation time of the lower-level model should be limited. On the other hand, the lower-level model should be realistic enough to provide a reliable prediction of the objective function values.

2.3. Lower-level model

Multimodal transportation networks are characterised by a diversity in infrastructure (road, rail), transport services (PT lines, bicycle renting), and modes (train, bus, tram, metro, car, bicycle and pedestrian). A simple example of a multimodal network is shown in Figure 2.2. This multimodality has to be accounted for when assessing the performance of the multimodal

transport system, and even more so, when evaluating the effect of adding or removing typical multimodal facilities. This performance depends on the behavioural response of travellers to these changes in the network.

In a multimodal network, travellers may choose from a range of private and public modes. Further, they may combine these modes through transfers between modes, so a large number of travel alternatives is available. Such multimodal travel alternatives might have complex combinatorial structures, so planning such a trip involves multiple decisions to be made, related to the various modes, transport services, boarding and alighting locations, transfer options, and parking facilities. As the variation in characteristics (such as time, costs, and distance) among alternatives is expected to be higher than in the unimodal case, personal preferences, knowledge and perception may play an important role in making this decision. These preferences can relate to the utility provided by different modes of PT, different PT stops or different services (i.e. a fast or slow route, or a route with or without a transfer). Moreover, from the traveller's point of view, not all these alternatives may be seen as completely distinct travel options, as there are large similarities in terms of infrastructure usage, modes of transport, and PT services. Finally, transfer movements are of major influence on the attractiveness of a multimodal trip. A transport modelling tool that assesses multimodal infrastructure measures should be able to model mode and route choice adjustments due to these measures. It should take the mentioned aspects into account as much as possible to adequately describe multimodal travel behaviour. However, it acts as a lower-level model in the bi-level optimisation problem, so computation time should be as low as possible. Further, developing such a model is not the main objective of this thesis, so existing methods should be used as much as possible.

To this end, in the following section first requirements are defined for the lower-level model that is needed. Following these requirements, the next section provides an overview of modelling approaches that describe decision making of the traveller, to model flows through such a multimodal transportation network. This leads to a conclusion on what method is most suitable to use.
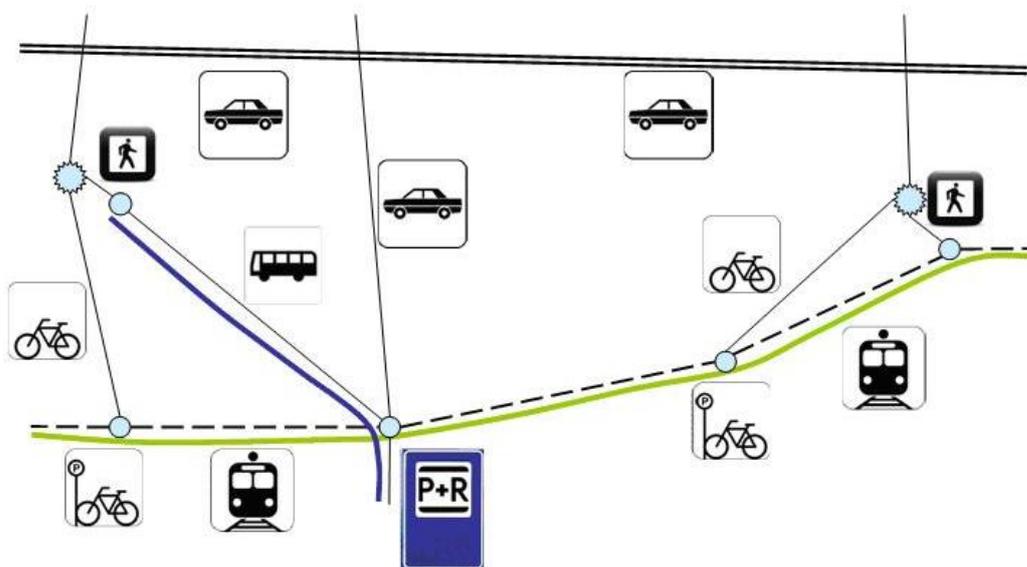
Figure 2.2: An example of a multimodal transportation network

2.3.1. Model requirements

The following four requirements are defined for the lower-level model.

1. *The model should adequately describe the behavioural response to measures taken, i.e. it should enable multimodal trip-making*

Following Chapter 1, this thesis focuses on policy measures that make multimodal trips more attractive. The lower-level model is necessary to evaluate this type of measures, which only make minor changes to the transportation network. This has consequences for the relevance of each of the steps from the traditional 4-step model (generation, distribution, mode choice and route choice). It is expected that the measures have minor influence on the long-term decisions in steps 1 and 2 of the 4-step model. This leads to the delineation that the total OD demand matrix is fixed. However, the third step is important to take into account, because the measures change the attractiveness of modes and especially of combinations of modes, which has consequences for the mode choices that are made by travellers. Therefore, the first requirement is that the lower-level model has to include the ability for (modelled) travellers to transfer from one mode to another mode in the network: to make a multimodal trip. Such a multimodal trip consists of several modes (also called mode chain). This trip may replace a trip that was formerly made by a single mode, so mode choice should be included in the model. Further, the measures influence the attractiveness of routes (including multimodal routes) and enable new routes, so the fourth step is also relevant. Note that the fourth step is also called assignment, comprising of route choice and network loading. Here attention is only given to the route choice problem, since network loading can be easily done when the route fractions are known (resulting from the route choice problem).

2. *The model should determine car and PT flows*

The system objectives will be largely determined by the flows on the car and the PT network, rather than the flows on the bicycle and pedestrian network. Firstly, car travel times depend on car flows, because flows in most cases are close to capacity, resulting in congestion. Bicycle and pedestrian flows are expected to be much lower than the capacity that is available for these modes, so their travel times do not depend on the flow. Secondly, also emissions and urban space used by parking are influenced by car flows, but not by bicycle and pedestrian flows. Thirdly, PT flows directly influence the revenue of the PT system, while bicycle and pedestrian flows do not. Concluding, the model should accurately model the car and PT flows, to enable calculation of system objective values from the model results. The bicycle and pedestrian flows are not needed, since the objective function values do not depend on these flows.

3. *The model should be able to deal with networks of realistic size*

It appears in literature (see Section 2.2) that the NDP is NP-hard and therefore a heuristic is needed. GAs are widely used and perform well in comparable NDPs. This leads to the use of a GA in the upper level, implying that many function evaluations are needed to reach enough convergence in the GA. Since the lower-level model is used to correctly determine objective function values, the computation time to solve to model should be low enough to be able to repeatedly run the model on networks of realistic size.

4. The model should be available in an existing software environment

The focus of this thesis lies on improvement of the solution algorithm, implications of the method for the modelling results and on interpretation of these modelling results, rather than on the development of a state-of-the-art multimodal assignment model. Therefore, components for the lower-level model to be used should already exist in a software environment, so that these components can be combined to an operational modelling framework.

2.3.2. Earlier research on modelling multimodal trips

Following requirement 1, this section presents methods that model mode and route choice, including routes that contains multiple modes. Mode and route choice of travellers depend on a range of factors. Paulley et al. (2006) identify system characteristics like travel times, fares and quality of service and individual's characteristics and lifestyle, like income, driving license, car ownership, and environmental concerns. Beirão and Sarsfield Cabral (2007) add the type of journey (purpose, related to luggage to be taken) and service performance perception of each transport mode to this. This thesis focuses on influencing sustainability objectives by improving the multimodal network, i.e. changing the system characteristics, while travellers' characteristics are constant. Variation in travellers' characteristics results in different modes and routes to be taken on a single OD pair, since different travellers have different preferences for the system characteristics of these modes and routes. This aspect is important to fulfil requirement 2, and is the main reason to include multiple routing. A distinction between various trip purposes and / or user groups would enable a more detailed modelling of the influence of travel behaviour of travellers with different characteristics on the network flows, to the cost of more data requirements and more computation time.

In a multimodal network multiple modes of transport and PT services, as well as transfer facilities connecting them, are available to the traveller. According to Fiorenzo-Catalano (2007), multimodal transport models should be able to predict the usage of the full range of modes and mode chains (to meet requirements 1 and 2), but this may lead to problems regarding requirements 3 and 4. This section contains two main approaches to cope with mode and route choice in a multimodal network, based on literature. The first approach, the supernetwork approach, fully integrates the car, PT, pedestrian and bicycle networks and results in simultaneous mode and route choice. The second approach makes a main distinction between car-only trips (unimodal trips) and trips that involve PT (combined with access and egress modes these trips are always considered to be multimodal). The section ends with some issues that are relevant for both approaches (capacity of PT vehicles and transfers). This provides the necessary context to explain in the next section which model is suitable to be used in the remainder of this thesis.

Supernetwork approach

The term 'supernetwork approach' indicates the family of models that includes three main modelling components: construction of an integrated multimodal network representation, the supernetwork (Sheffi, 1985), generation of the route set and simultaneous mode and route choice modelling (Benjamins et al., 2002). In this subsection each of these components is

discussed concisely. A recent application of the supernetwork approach can be found in Van Eck et al. (2014). The model is summarised in Figure 2.4.

The supernetwork representation (see Figure 2.3) contains a layer with a directional graph for each mode, usually PT, car, bicycle and walking. Instead of allowing (infrastructure) links to be used by multiple modes, these network links now appear in multiple layers (corresponding to modes). In PT layers links are duplicated for every PT service traversing this link. The layers are integrated into a single network by adding artificial transfer links (vertical lines in Figure 2.3), connecting the pedestrian network to the remaining layers (all layers are connected with each other by the pedestrian network). These transfer links represent transfer possibilities and hence their cost function is based on attributes such as fixed PT costs, parking time, and parking costs. Private transport layers (car and bicycle) are connected to the pedestrian layer at locations where these vehicles can be parked, for example, centroids, car parks and P&R facilities. At stops the pedestrian layer is connected to PT layers.
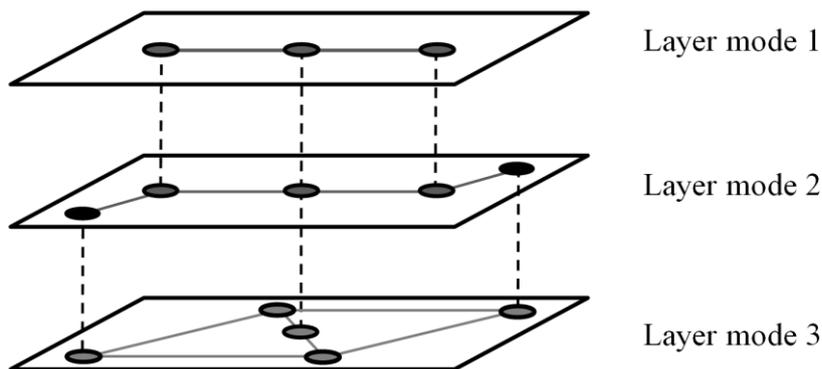
Figure 2.3: The supernetwork representation: each mode has its own network layer. The vertical, dashed lines represent transfer links between the modes / layers

The next step is to generate the route set. Routes in which several modes are combined can be feasible alternatives and as such need to be considered. The route set is generated by first extracting a subset of available routes, which is usually done through repeated shortest-path searches with randomised attribute values and preference parameters. In addition, enough non-car alternatives should be included in the route set to represent travellers without driving licenses and / or without cars. These routes may not appear if they are not explicitly added, because OD relations exist where the car is by far the fastest travel option. On the other hand, not every mode chain is feasible (Fiorenzo-Catalano, 2007), for example a trip composition train-car-train is very unlikely, since usually no private vehicle is available at the train station. Such infeasible travel alternatives should be excluded from the set of considered travel alternatives (realism of the route set). To improve the computational efficiency, the set of feasible routes is filtered further by logical constraints. Finally non-competitive routes and routes showing large overlap with more attractive alternatives are excluded from the set. These filtering procedures prevent the model from assigning travellers to illogical or behaviourally doubtful alternatives.

```
┌──────────────┐      ┌──────────────────┐
│ Supernetwork │─────▶│ A-priori choice set│
└──────────────┘      │    generation     │
                      └──────────────────┘
                                │
┌──────────────┐      ┌──────────────────┐
│Transportation│─────▶│     Mode and      │
│   demand     │      │   route choice    │
└──────────────┘      └──────────────────┘
```
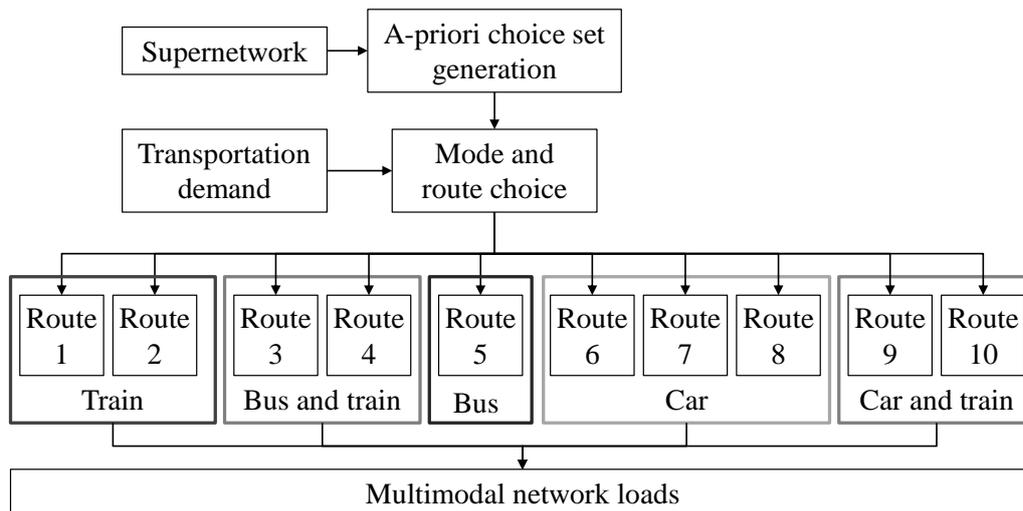
Figure 2.4: Combined mode and route choice in the supernetwork model for an example choice set

After such a route set is generated, simultaneous mode and route choice remains. The mode and route choices may be seen as a simultaneous choice as they are heavily correlated (Arentze and Molin, 2013). Travel alternatives describe both the routes and the (chains of) modes that are used. Increased complexity of the network implies larger differences in travellers' knowledge and perception of travel alternatives and their attributes as well as more variation in travellers' preferences, so it is more realistic to model the route choice in a stochastic way rather than performing a deterministic assignment, especially when modelling multimodal trips. Another relevant issue when the choice between travel alternatives is explicitly modelled is the following. Mode, services or space related unobserved route attributes could have rather complex correlation patterns (Hoogendoorn-Lanser et al., 2006). Travellers may have intrinsic preferences for certain modes that are not represented by any observed attributes. Furthermore, the attractiveness of a certain part of a trip can depend on the trip composition as a whole (Fiorenzo-Catalano et al., 2003). For example, the costs for using a bicycle as access or egress mode are usually different at the home-side than at the activity-side of a trip. Where a bicycle is usually available for free at the home-side, a bicycle needs to be rented or parked in advance at the activity-side of a trip, implying extra costs. Similarly, the attractiveness of travelling by train is higher when one can board or alight at an intercity station (Bovy and Hoogendoorn-Lanser, 2005). This is independent of the hierarchic level of the train service (stop, regional, or Intercity) being used.

The supernetwork approach contains detailed, conceptually sound behavioural rules and is therefore very well capable of modelling combinations of modes and producing correct flows on the car and PT networks (requirements 1 and 2). On the solving the model other hand, solving the model is computationally expensive (requirement 3), especially the generation of the choice set, and it is not yet easily available (requirement 4).

Multiple access and egress modes and multiple routing in PT

The second approach makes a main distinction between car-only trips (unimodal trips) and trips that involve PT (combined with access and egress modes these trips are always

considered to be multimodal). Bicycle-only and walk-only trips are not considered, based on requirement 2: the flows of bicycle and walk do not have influence on the system objective values. To take heterogeneous preferences and perceptions among travellers into account (requirement 2 as well), multiple routing is needed. In car traffic assignment, multiple routing is induced by assuming user equilibrium, because travel times depend on flows, which for most OD pairs results in multiple routes with equal, minimum travel times being used. For trips that involve PT, one aspect of multiple routing is modelled by introducing multiple access and egress modes by pre-specifying mode chains as a specific combination of access mode and egress mode, with PT in between. When both the access mode as the egress mode are fixed (and only shortest access and egress routes are assumed between centroids and stops), the problem is reduced to a PT assignment problem, which is addressed well in literature. Access and egress stop choice is then included in the PT assignment. To limit computation time the number of candidate stops for access and egress should be kept limited by defining constraints (e.g. for walking the distance should not be too far or for car a P&R facility should be available). The main mode choice is between car and PT, but within the PT option several mode chains are available as a mode choice option for the traveller. This is summarised in Figure 2.5. The remainder of this subsection reviews PT assignment methods that include multiple routing and relates these methods to the requirements in Section 2.3.1.
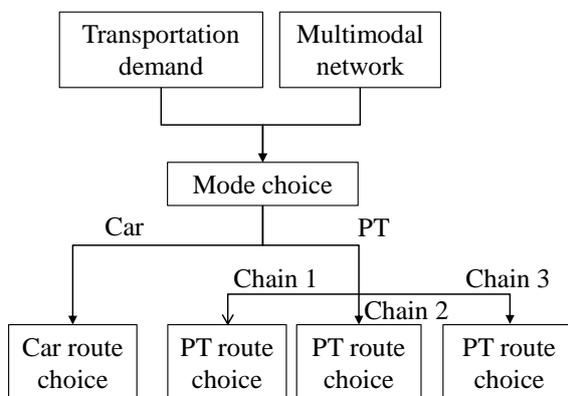
Figure 2.5: Mode and route choice when combinations of an access mode, PT and an egress mode are defined as mode chains

Multiple routing through the PT network can be achieved in three ways. The first is stochastic assignment (also called probit), where preference parameters (such as value of time) or attribute values (such as link travel time) are drawn from a distribution, and the shortest path is sought in the modified network multiple times (Nielsen, 2000). This makes it possible to describe passengers' different preferences towards different sub-modes and against transfers. This may also include dependencies of choices through chains of sub-modes and differences in the distribution of travel- and waiting times for different sub-modes. This methodology can describe route choices in public transport very well, both due to the model's ability to describe overlapping routes and due to the many different coefficients, error components and distributions that make it possible to calibrate the model (meeting requirement 2). The many parameters are a weakness of this approach, since this complicates the calibration. Furthermore, drawing from a distribution involves large computation time, making it

impossible to apply in a large-scale NDP, as is stated by Uchida et al. (2007), so requirement 3 is not met.

The second way to achieve multiple routing is also a probabilistic approach, but analytically calculates the likelihood that any alternative is the shortest path, resulting in smaller computation times when the model is solved (meeting requirement 3). The method of optimal strategies (Spiess and Florian, 1989) is the first implementation of this approach which assumes random arrival of passengers and services and assumes that a passenger boards the first arriving vehicle that belongs to a line in the attractive set. This leads to a distribution among the attractive lines proportional to frequency. The concept is implemented in several software packages and used extensively (meeting requirement 4). A strategy specifies the set of PT lines at a stop that bring the traveller closer to the destination. A more formal definition of the strategy is named hyperpath, which is a so called unique acyclic support graph and represents all paths that may be attractive for the traveller. Specifically for PT assignment the hyperpath framework is described by Nguyen and Pallottino (1989) and Spiess and Florian (1989). This theory is further developed in Nguyen et al. (1998): a logit model is introduced to calculate the distribution over the different paths in the hyperpath based in generalised costs instead of only using the frequencies of the PT lines in the attractive set. Lozano and Storchi (2002) include personal preferences, like expected travel time and number of transfers, to choose the hyperpath that is optimal for a specific user. Florian and Constantin (2012) also introduce a logit model for splitting passengers among options, specifically to get a better distribution of passengers over routes and to improve modelling of access to the PT network by walking. These improvements result in more realistic distribution of travellers among routes (to better meet requirement 2).

The third possibility that naturally incorporates multiple routing is dynamic assignment, taking the complete schedule into account, for example see Wilson and Nuzzolo (2008). In schedule based modelling, each individual PT vehicle is modelled explicitly, making it possible to exactly calculate transfer times and waiting times, dependent on the departure time. Different departure times then naturally imply a spread over different routes, since for each time the optimal route can be determined and very realistic routes are found (meeting requirement 2). On the other hand, this method is data-intensive and computationally expensive, so it does not meet requirement 3.

Congestion and crowding

When travel demand is high, congestion occurs on the road network and the high occupancy level of PT services has an impact on travel time and comfort level experienced by travellers. The consideration of alternatives in which multiple (private and PT) modes are combined introduces correlations in travel times and flows between different modes of transport. Higher travel times on the car network might impact the share of P&R alternatives: the quality of PT travel options that have car as access or egress mode are influenced by the flows on the car network. Another issue related to capacity is the available number of parking places at bike and ride and P&R facilities. More travellers using such a facility increase walking times to and from the parking location at first (Hensher and King, 2001). Eventually, no more vehicles can be parked when the available capacity is met. Furthermore, it is possible to introduce

crowding to represent a less comfortable trip when a traveller has to use an overcrowded PT vehicle, representing congestion in the PT system, see for example Cepeda et al. (2006).

Introducing capacity in PT and / or parking lots makes the costs of using them dependent on the flow, which makes it necessary to introduce an iterative equilibrium assignment. This would increase computation times too much for application in this thesis (requirement 3), so PT and parking capacity are not taken into account in the PT assignment, given the limited congestion in the PT network in the Netherlands. However, congestion on the car network is included when the travel times and costs of P&R alternatives are determined.

Transfers

A transfer can be made between a private mode and a PT mode, or within the PT network. Making a transfer usually involves costs, for example parking costs, parking time, walking time, discomfort or the risk of missing a PT connection. These transfer-related costs are a substantial part of the total generalised costs of a full trip (Bos et al., 2004; Hoogendoorn-Lanser et al., 2006). To fulfil requirement 2, it is therefore necessary to predict the impact of (changes in) transfer attributes.

2.3.3. Conclusion

This section showed two methods to model trips through a multimodal network. The first method uses a supernetwork that combines the networks of all public and private modes to one network. In this supernetwork, a route set is generated that contains all relevant routes, including unimodal and multimodal routes. Traffic is assigned to these routes depending on their attractiveness, taking overlap between the routes into account. However, no real-world application of this model with limited computation time was yet available during this research. So although the supernetwork approach meets requirements 1 and 2 very well, it does not meet requirements 3 and 4. This leads to the second best method to be used. The chosen method involves dividing trips into trips that have car as the only mode and PT. For car-only trips standard methods are available and applied. For trips that include PT, several mode chains are distinguished (requirement 1) that consist of multiple combinations of access and egress modes (with PT always acting as main mode). Route search within one mode chains is modelled by a PT assignment algorithm that is known from literature and is applied in practice in limited computation time (requirements 3 and 4). To accurately model car and PT flows (requirement 2), different preferences of travellers are taken into account by multiple routing. The resulting transportation model therefore contains both unimodal and multimodal trip options during mode and route choice and is described in more detail in Chapter 3.

2.4.   Conclusion

The NDP has been studied widely in literature. However, the specific combination needed in this thesis (to fulfil the objective of the thesis in the context of a sustainable development of the accessibility of the Randstad area in the Netherlands), has not yet been found. This combination includes application of the NDP to a multimodal passenger transportation network, including multiple objectives resulting in a Pareto front, with an application on a

network large enough to cover a considerable part of the Randstad area. Nevertheless, literature provides elements that are combined to a suitable model. GAs make it possible to solve large-size problems efficiently, if a fast enough lower-level model is available. For this lower level, various aspects are relevant in the context of multimodal networks. Many researchers studied components of a multimodal trip through such a network related to mode and route choice. Applications on large-scale networks are limited, since computation times are usually extensive. A model that covers all relevant aspects and at the same time can be solved fast enough is not yet available. However, elements in literature can be combined in such a way that the requirements for the lower-level model are met as much as possible by including elements that provide a large model accuracy, like multiple access and egress modes and multiple routing. At the same time, computation time to solve the model is limited (needed for the application of this model within the multimodal NDP at hand) and the models are almost ready to use. The resulting model is mathematically formulated in Chapter 3. From earlier research it can be concluded that GAs are effective and efficient in NDP studies. The most promising GAs are used to solve the problem at hand, as is further described in Chapter 4.

Chapter 3: Modelling framework

After providing background information in Chapter 2, in this chapter the problem is formulated mathematically. First, the formulation of the upper-level optimisation problem is given, including network definitions. Thereafter, the models used in the lower level are described, which include multiple routing for public transport (PT) and for car and mode choice, including multiple access and egress modes. The next section provides the effect models used to calculate the values of objective functions. Finally, the optimisation problem is further specified for the cases: networks are described and decision variables are defined for two practical case studies in the northern part of the Randstad area in the Netherlands. These case studies are used in the experiments that follow in further chapters.

This chapter is partly based on:

Brands T., E. C. van Berkum (2014) Performance of a genetic algorithm for solving the multi-objective, multimodal transportation network design problem, *International Journal of Transportation* 2(1), pp. 1-20

Brands T., E. de Romph, T. Veitch, J. Cook (2014) Modelling public transport route choice with multiple access and egress modes. *Transportation Research Procedia* 1, pp. 12-23

3.1.   Mathematical optimisation problem

The multi-objective, multimodal passenger transportation network design problem is formulated as a mathematical optimisation problem. In this section first the necessary definitions are given, followed by the formulation of the optimisation problem in the form of a mathematical programme and a formal definition of Pareto-optimality.

3.1.1.   Network and demand definition

The multimodal passenger transportation network is defined as a directed graph G, consisting of node set N, link set A, line set L and stop set U (see also the example in Figure 3.1). For each link a one or more modes that can traverse that link are defined. If cars are allowed, a free flow speed v_a^0 and capacity x_a^{max} are defined. Origins $r \in R$ and destinations $s \in S$ are subsets of N, i.e. trips start at a node in R and end at a node in S. Note that nodes may be in both sets: nodes may act as both origin and destination. A node in R and / or S is also called a centroid. Total fixed transportation demand $\overline{q}_{rs}$ is stored in a matrix with size $|R| \times |S|$. PT service lines $l \in L$ are defined as ordered subsets $A_l \subseteq A$: a PT line traverses several consecutive links in the graph. This implies that a route that uses a PT line may only consist of links in A_l that follow the PT line in the correct direction. Each line l belongs to a PT system / vehicle type $T_l \in B$ (i.e. bus, tram, metro, stop train, express train). Furthermore, links that can be traversed by non-car access and egress modes (walking and bicycle) are defined as subsets $A_l \subseteq A$. For each PT line and each non-car access / egress mode, (in-vehicle) travel times per link tv_{al} are defined, which are assumed to be independent of the flow. In addition, for each PT line a frequency F_l is defined. PT flows can only traverse PT service lines or access / egress links. PT stops $u \in U$ are defined as a subset of N. Consequently, a line l traverses several stops. A stop represents a train station, a metro station, a tram stop or a bus stop, depending on the PT lines that serve the stop. The term 'stop' is used to indicate all these types of stops, except for cases where the context is clearly related to train / metro, in which case 'station' is also used. Access / egress modes and PT are only connected through these stops: at a stop the traveller may transfer to other PT lines or other modes. Whether a line calls at a stop u or not, is determined by D_u^T: a subset from B of types of PT lines that serve that stop. Whether a stop contains a park-and-ride (P&R) facility or not is indicated by a binary stop property D_u^P. All together, the transportation network is defined by $G(N, A, L, U)$. For each OD pair, both (unimodal) car routes $k \in K_{rs}$ and routes that contain PT $k \in \hat{K}_m^{rs}$ are defined. The latter routes are multimodal, in which a combination of using a specific access mode, PT and a specific egress mode is defined as a mode chain $m \in M$. Depending on mode chain m, these routes may contain car links (on which travel time depends on car flow) or may only contain non-car links (on which travel time does not depend on flow).

In Figure 3.1 a small example network is shown with 8 nodes (of which 2 centroids), 7 links, 3 stops, and 1 PT line, to illustrate the objects that define the network. An access route from the origin to the boarding stop is shown that consists of a connector link (dashed) and a link that allows walking. Similarly, an egress route is shown from the alighting stop to the destination. The PT line traverses three links, i.e. links that allow PT. The PT line passes three stops.
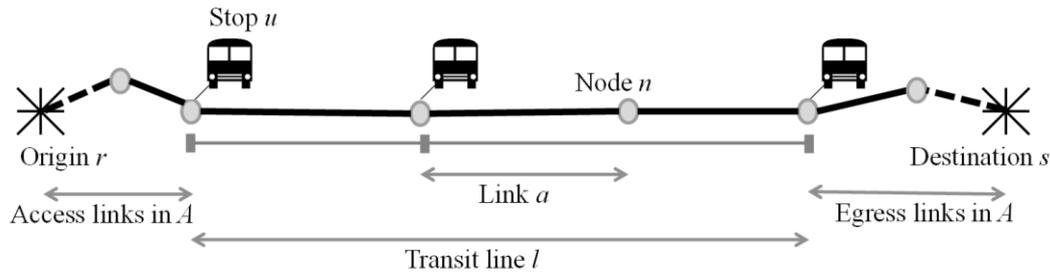
Figure 3.1: Example network, showing 8 nodes (including origin and destination), 7 links (including 2 dotted connector links and 2 other access and egress links), 3 stops and 1 PT line

3.1.2. Formulation multimodal MO-NDP

A decision vector $\underline{y}$ is defined, which consists of V decision variables: $\underline{y} = [y_1, \cdots, y_v, \cdots, y_V]$. Y is the set of feasible values for the decision vector $\underline{y}$ (also called decision space). The objective vector $\underline{Z}$ (consisting of W objective functions, $\underline{Z} = [Z_1, \cdots, Z_w, \cdots, Z_W]$ that are defined in Section 3.3) is a function of the decision vector $\underline{y}$. $\underline{Z}(y)$ belongs to the so-called objective space. In principle $\underline{Z}$ may be any value from $\mathbb{R}^W$, but depending on its meaning, an objective function may be subject to natural bounds (for example, travel time can never be negative).

The multi-objective optimisation problem (MOOP) is defined in programme 3.1-3.15 as a bi-level problem. The lower-level programme in 3.3-3.15 is based on the formulation in Sheffi (1985). All candidate lines and stops are included in the sets L and U, where the frequencies F_l and the properties D_u^T and D_u^P of candidate lines and stops may change dependent on the decision vector $\underline{y}$ (the relation between $\underline{y}$ and F_l, D_u^T and D_u^P is defined for the case studies in Section 3.4). The frequencies and dwelling properties of other lines and stops do not change, i.e. are not dependent on the decision vector $\underline{y}$. These characteristics define the multimodal transportation network G, where a line l does not exist when $F_l = 0$ and a stop u does not exist when $D_u^T = \{\}$. A fixed passenger transportation demand $\bar{q}_{rs}$ is assigned to this transportation network, such that satisfies the constraints 3.2 and 3.3. Feasibility constraint 3.2 should be satisfied to define a transportation network that is physically within reach. Constraint 3.3 represents the lower-level optimisation problem, which optimises modal split and flows in the network from the travellers' perspective and has Equations 3.4-3.15 as constraints. In Eq. 3.3 the lower-level objective z is optimised. The first term represents the user costs in the car network and the second term represents the perceived costs in the PT network. This formulation results in user equilibrium, i.e. to the situation where all used car routes have equal costs, that are also equal to the perceived costs in the PT network. Freight is assumed to take the shortest route on the free-flow network, so the link based freight loads can be directly derived from freight demand $\tilde{q}_{rs}$ in constraint 3.16. The costs in the car system depend on the freight traffic represented by $\tilde{x}_a$, which is assumed to be a fixed pre-load (expressed in Passenger Car Equivalents).

Constraint 3.4 ensures that the sum of car route flows (expressed in vehicles using the average car occupation Ω) and PT demand equals total demand. Constraint 3.5 is a non-negativity constraint for car route flows. Note that PT route flows are nonnegative by the definitions in

constraints 3.6 and 3.7. 3.6 ensures positive PT demand. 3.7 distributes PT demand over various mode chains and PT routes, using the fractions in constraints 3.12 and 3.13: 3.13 defines the route fractions for PT trips per mode chain: the PT model contains multiple routing, where trips are distributed among sensible routes using logit models for line choice and stop choice (see Section 3.2.1). Constraint 3.12 defines the mode chain fractions for PT trips. Constraint 3.10 calculates the OD based generalised costs per mode chain as a weighted average of route costs, to correctly take into account combined waiting costs of multiple PT lines. Constraint 3.9 calculates the logsum for the mode chains, to represent OD based PT costs. Constraints 3.8 and 3.11 define the relation between link costs and route costs, for car and for PT. Note that PT generalised costs not only depend on which links / lines are used, but also on which combination of links / lines, to correctly calculate waiting time and penalty costs (see Section 3.2.1 for generalised costs calculation). Finally, constraints 3.14-3.16 define the vehicle and passenger link flows based on car, freight and PT route flows. Note that PT routes may contain car links, in the case of the use of a P&R facility in the route (see constraint 3.14). On the other hand, car routes never include PT passenger links, as can be observed in constraint 3.8. When a PT route contains walking and / or bicycle, the use of these links is included in $\hat{t}_{al}$, where l now represents walking or bicycle instead of a PT line.

$$\min_{\underline{y}} \underline{Z}(\underline{y}, \underline{x}^*, \underline{\hat{x}}^*, \hat{q}_{rs}^*) \qquad , s.t. \tag{3.1}$$

$$\underline{y} \in Y \tag{3.2}$$

$$\underline{x}^*, \underline{\hat{x}}^*, \hat{q}_{rs}^* = \arg\min_{\underline{x}, \hat{\underline{x}}, \hat{q}_{rs}} z(\underline{y}, \underline{x}, \hat{\underline{x}}, \hat{q}_{rs}) = \sum_{a \in A} \int_0^{x_a} t_a(\omega, \tilde{x}_a, \underline{y}) d\omega + \sum_{r \in R, s \in S} \int_0^{\hat{q}_{rs}} \left(\frac{1}{\theta_1} \ln \frac{\omega}{\overline{q}_{rs} - \omega} + \hat{c}_{rs} \right) d\omega$$

$$\tag{3.3}$$

$$s.t. \quad \Omega \sum_{k \in K_{rs}} f_k^{rs} + \hat{q}_{rs} = \overline{q}_{rs} \qquad\qquad \forall r, s \tag{3.4}$$

$$f_k^{rs} \geq 0 \qquad\qquad \forall r, s, k \in K_{rs} \tag{3.5}$$

$$\hat{q}_{rs} \geq 0 \qquad\qquad \forall r, s \tag{3.6}$$

$$\hat{f}_{km}^{rs} = \psi_m^{rs} \zeta_{km}^{rs} \hat{q}_{rs} \qquad\qquad \forall r, s, m, k \in \hat{K}_m^{rs} \tag{3.7}$$

$$c_k^{rs} = \sum_{a \in A} t_a(x_a, \tilde{x}_a, \underline{y}) \delta_{a,k}^{rs} \qquad\qquad \forall r, s, k \in K_{rs} \tag{3.8}$$

$$\hat{c}_{rs} = -\frac{1}{\theta_2} \ln \sum_{m \in M} e^{-\theta_2 \hat{c}_m^{rs}} \qquad\qquad \forall r, s \tag{3.9}$$

$$\hat{c}_m^{rs} = \sum_{k \in \hat{K}_m^{rs}} \zeta_{km}^{rs} \hat{c}_{km}^{rs} \qquad\qquad \forall r, s, m \tag{3.10}$$

$$\hat{c}_{km}^{rs} = \sum_{a \in A} \left(t_a(x_a, \tilde{x}_a, \underline{y}) \breve{\delta}_{a,km}^{rs} + \sum_{l \in L} [\hat{t}_{al}(\underline{y}) \hat{\delta}_{al,km}^{rs} + \sum_{l' \neq l} \sum_{a' \neq a} \hat{t}_{a'l'al}(\underline{y}) \hat{\delta}_{a'l'al,km}^{rs}] \right)$$

$$\forall r, s, m, k \in \hat{K}_m^{rs} \tag{3.11}$$

$$\psi_m^{rs} = \frac{e^{-\theta_2 \hat{c}_m^{rs}}}{\sum_{m \in M} e^{-\theta_2 \hat{c}_m^{rs}}} \qquad \qquad \forall r,s,m \qquad \qquad (3.12)$$

$$\zeta_{km}^{rs}(\hat{c}_{km}^{rs}; \hat{c}_{k'm}^{rs} \forall k' \in \hat{K}_m^{rs}, k' \neq k; \theta_3; \theta_4) \qquad \forall r,s,m,k \in \hat{K}_m^{rs} \qquad (3.13)$$

$$x_a = \sum_{r \in R} \sum_{s \in S} \sum_{k \in \{\hat{K}_{rs}\}} f_k^{rs} \delta_{a,k}^{rs} + \Omega \sum_{r \in R} \sum_{s \in S} \sum_{m \in M} \sum_{k \in \{\hat{K}_m^{rs}\}} \hat{f}_{km}^{rs} \breve{\delta}_{a,km}^{rs} \quad \forall a \qquad (3.14)$$

$$\hat{x}_{al} = \sum_{r \in R} \sum_{s \in S} \sum_{m \in M} \sum_{k \in \hat{K}_m^{rs}} \hat{f}_{km}^{rs} \hat{\delta}_{al,km}^{rs} \qquad \qquad \forall a,l \qquad (3.15)$$

$$\tilde{x}_a = \sum_{r \in R} \sum_{s \in S} \tilde{q}_{rs} \tilde{\delta}_a^{rs} \qquad \qquad \forall a \qquad (3.16)$$

Based on the resulting vehicle and passenger flows in the network G, an objective vector is calculated for a decision vector $\underline{y}$ in a general formulation (3.1). The case-specific formulation of objective functions is given separately in Section 3.3.

3.1.3. Multi-objective optimisation: Pareto-optimality

Since the problem is formulated as an MO optimisation problem (with objectives that are at least partly opposed), no single optimal solution exists. Instead, the result of solving the formulated problem consists of a set of possibly optimal solutions, the Pareto set. It depends on the decision maker's preferences which solution from the Pareto set is chosen as the best solution to be implemented. These preferences are not taken into account in the optimisation problem, since this is considered to be part of the post-optimisation analysis.

Mathematically, the concept of Pareto dominance is as follows. Assuming two decision vectors $\underline{y}_i, \underline{y}_{i'} \in Y$, $\underline{y}_i$ is said to strongly dominate $\underline{y}_{i'}$ if and only if $Z_w(\underline{y}_i) < Z_w(\underline{y}_{i'}) \ \forall w$, or mathematically $\underline{y}_i \prec \underline{y}_{i'}$, following the notation used by Deb (2001). Additionally, $\underline{y}_i$ is said to cover or weakly dominate $\underline{y}_{i'}$ if and only if $Z_w(\underline{y}_i) \leq Z_w(\underline{y}_{i'}) \ \forall w$, or mathematically $\underline{y}_i \preceq \underline{y}_{i'}$. All solutions that are not weakly dominated by another known solution are possibly optimal for the decision maker: these solutions form the Pareto-optimal set P. When the set of all known solutions is defined as ϕ, there is no $\underline{y}_i \in P \ | \ \exists \underline{y}_{i'} \in \phi : \underline{y}_{i'} \preceq \underline{y}_i$.

3.2. Multimodal network modelling

In this section the models used to put the mathematical formulation of the lower level in the previous section into operation are presented in more detail. As defined in Section 3.1.1, the concept of mode chains is used, so a trip is either a unimodal car trip, or a multimodal trip with PT as main mode. First, the PT route choice model is explained. As explained in Chapter 2, this model puts multimodal trips into operation by including multiple access and egress modes (walking, bicycle and car are relevant for the Dutch situation) in mode chains. Such a mode chain consists of an access mode (used to reach the PT network), a part of the trip travelled in PT and an egress mode (used to reach the destination when leaving the PT network). The PT part may contain multiple PT service lines, which can be of the same or of a different PT system (i.e. train, bus, metro). If multiple PT services are used, a transfer is made between them, for which a short walk may be required. The model contains multiple routing

within one mode chain, to take heterogeneous preferences and perceptions among travellers into account. Next, the model that is used for (unimodal) car assignment is described. The section ends with the mode choice model, which brings both assignment models together and concludes the lower-level model.

3.2.1. Public transport route choice

In this section the public transport route choice algorithm is explained. This route choice algorithm calculates the routes from all origins to one destination all at once. First, an example is used to introduce the main steps of the algorithm. The step numbers correspond with the steps in the pseudo code that summarise the PT assignment algorithm at the end of the section. In the middle of the section the concept of generalised costs is formalised and the steps in the algorithm are explained in more detail.

Consider the illustrative example in Figure 3.2 (more details on each step are given later on in this section). The algorithm starts at the destination and builds a tree backwards from the destination to all origins in the network. Only stops that are in the egress candidate set U_{sm}^{E} are considered (step 1). All PT lines that serve these stops (in this case 3 lines) are now considered. For each stop upstream of this line, the costs to reach the destination using that line are calculated (yet without waiting costs, step 2). When more than one line serve a stop, the costs of the individual lines are combined to stop costs by a line choice model, resulting in costs per stop, including waiting time (step 3). Finally, these costs per stop are used in a stop choice model (taking access costs to reach the access stop from the origin into account) to determine the distribution among access stops from the access candidate set U_{rm}^{A} (step 6). Note that in small this example step 4 (walk transfers) and 5 (iterate) are not needed.
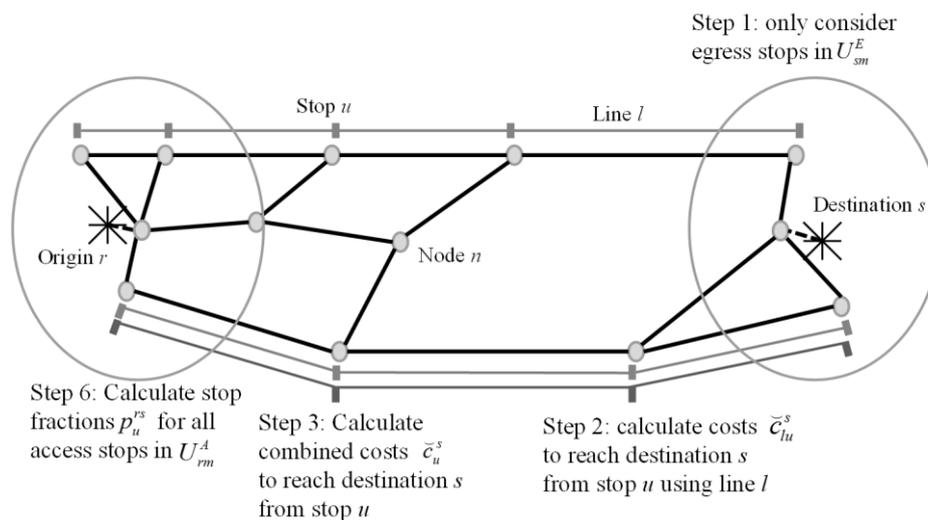
Figure 3.2: Illustration of steps in the PT route choice algorithm. Lines are visualised (in gray) alongside links (black), while in fact the lines are defined on the links. Stops are visualised on the lines, while in fact the stops are defined on the nodes. For simplicity, the example network does not contain a transfer and mode chain index *m* is here omitted

Generalised costs

Route search is based on generalised costs, which consist of several components. PT, walking and bicycle costs are assumed to be flow-independent: train, tram and metro lines have dedicated infrastructure. It is assumed that if severe congestion occurs in the car network (in regular situation), bus lines have dedicated infrastructure as well, which is mostly true for the cases considered in this research. Footpaths and cycling lanes usually have flows that are far below capacity, so no congestion is assumed on walking and bicycle links. Therefore, for mode chains that do not start with car access or end with car egress, the travel time is assumed to be independent of the flow (Eq. 3.17). Generalised costs contain distance (to represent monetary costs), travel time (both in the PT vehicle as when using an access or egress mode and dependent on PT or access / egress vehicle type *b*), waiting time and transfer penalty. For mode chains that start with car access and / or end with car egress, distance and time travelled by car are included (Eq. 3.18): those travel times depend on the flow.

$$\hat{c}_{km}^{rs} = \sum_{b \in B}[\hat{\alpha}_{1b}\hat{d}i_{kmb}^{rs} + \hat{\alpha}_{2b}\hat{t}v_{kmb}^{rs} + \hat{\alpha}_{3b}\hat{t}w_{kmb}^{rs}] + \hat{\alpha}_4\hat{n}t_{km}^{rs} \qquad \forall m \notin M_R \cup M_S \quad (3.17)$$

$$\hat{c}_{km}^{rs} = \sum_{b \in B}[\hat{\alpha}_{1b}\hat{d}i_{kmb}^{rs} + \hat{\alpha}_{2b}\hat{t}v_{kmb}^{rs} + \hat{\alpha}_{3b}\hat{t}w_{kmb}^{rs}] + \hat{\alpha}_4\hat{n}t_{km}^{rs} + \alpha_1 di_{km}^{rs} + \alpha_2 tv_{km}^{rs} \qquad \forall m \in M_R \cup M_S \quad (3.18)$$

Waiting time depends on the PT line and the stop where the PT line is boarded. The waiting time is a function of the headway (inverse of the frequency) of the PT line (see Eq. 3.20). The transfer penalty is used to model any extra cost (i.e. inconvenience) which is not addressable as travel time or waiting time. This takes into account that transferring between two PT lines is perceived as an inconvenience: given the same travel time and total waiting time, such an option is perceived as less attractive than an alternative with no transfer.

Step 1: access to and egress from the PT network

A multimodal network was already defined in Section 3.1.1. This means that origins and destinations are linked to stops utilising an underlying network. On this network, links may only be accessible by a subset of modes and speeds on the links are different for going by foot, bicycle and car. For most networks this implies that in theory from a centroid all or almost all stops in the network are reachable using the underlying network, resulting in large combinatorial complexity.

Therefore, in the first step of the PT assignment algorithm the number of choice options is limited. For each origin *r* a set of stops is identified where a PT trip from that origin may start, when using mode chain *m* (a candidate set U_{rm}^A). Because this set of stops is only a small subset of the total number of stops in the network, the path finding becomes a lot more efficient, while maintaining enough diversity to model variety among travellers. In Figure 3.3 an example network is given to clarify the criteria that are used to determine the candidate set. This network has a centroid in the middle of the network and contains 10 stops. Three of these stops give access to the train system, indicated by 🚊. Two stops have P&R facilities, indicated by Ⓟ. The following criteria for accessing the PT network are used, where a distinction is made between access modes. The candidate set is complete when every criterion is satisfied.

1. Distance radius. This criterion includes all stops within a certain radius around the centroid into the candidate set. For walk and bicycle access this criterion is used.

2. Type of PT system reached (i.e. bus, train). This criterion specifies that each centroid should have access to at least a minimum number of stops of a certain, typically higher order, PT system. For example, to include at least 1 train station in the candidate set, because walking to the nearest train station may be faster than taking a low frequency bus service to that (or another) train station.

3. Type of stop. This criterion specifies that the origin should have at least a minimum number of stops of a certain type in the candidate set. For car as access mode, this criterion defines a list of stops that have P&R facilities available.

4. Minimum number of stops. This criterion specifies the minimum number of stops in a candidate set.
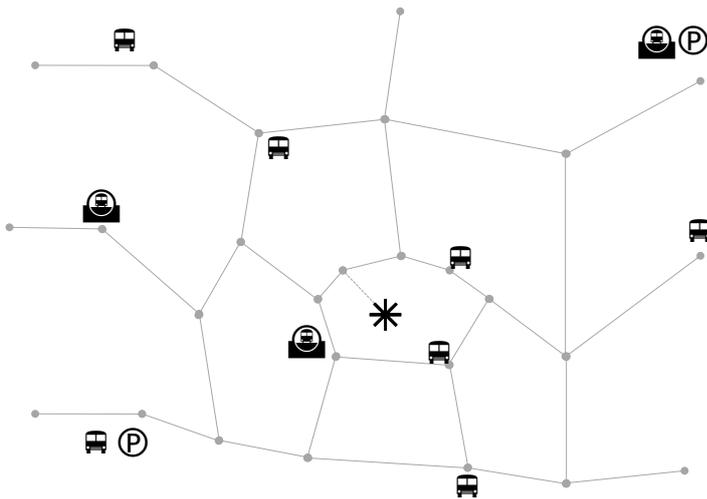
Figure 3.3: Example access to the PT network

For the example network, in Figure 3.4 (left) the candidate set for walking as access mode is shown. Only the first criterion is used: search radius. Since walking distances are smaller than cycling or car distances, a limited search radius is applied, resulting in the three stops just around the centroid to be included in the candidate sets U_{rm}^{A} (for mode chains m that have walking as access mode) for this origin r.

In Figure 3.4 (middle) the candidate set is shown for bicycle as access mode. The access radius is larger, which in the example results in two more bus stops to be included in the candidate set. Furthermore a cyclist prefers to access a higher order PT system, i.e. a train station. For this access mode the second criterion is added, such that at least two train stations are included in the candidate set

For car as access mode, the search radius is less relevant, since car is a fast access mode and the radius to be defined would be very large. However, the facilities at the stop are important, since only stops with P&R facilities can be used. So for car the third criterion is used, with P&R as a stop property. In Figure 3.4, the selected stops are indicated by Ⓟ.
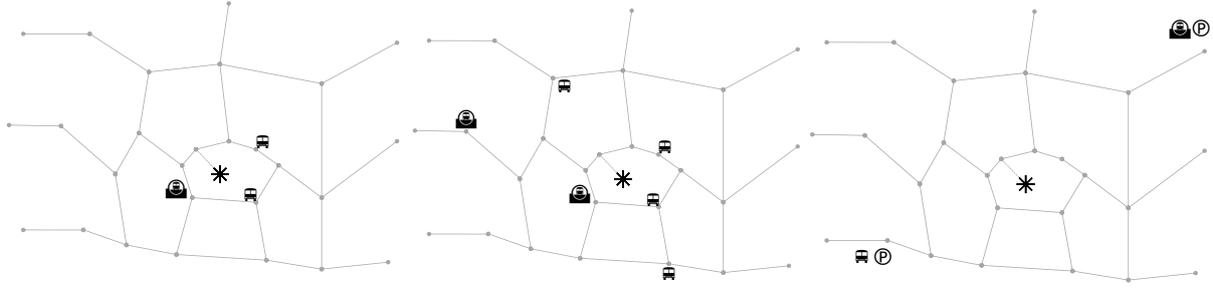
Figure 3.4: Candidate set for mode walk (left), bicycle (middle) and car (right)

The path finding from the centroid to the set of access stops is done with a traditional Dijkstra shortest path algorithm based on generalised cost, which includes distance and (for car possibly congested) travel time.

When alighting the PT network the exact same criteria and algorithms are used as for accessing the PT network. In practice this does not necessarily mean that the access stop set of a centroid is equal to the egress stop set of the same centroid (with equal access and egress modes), due to different characteristics of the links in opposite directions, such as one-way streets. When multiple mode-chains are evaluated these stop sets are determined for each relevant mode. The resulting set of stops that is found to be relevant for a certain destination zone s is referred to as the candidate set U_{sm}^E.

Step 2 and 3: line choice model

As defined earlier, for every origin r and destination s in the network a set of candidate stops where a PT trip may begin (the first boarding stops) U_{rm}^A and the set of candidate stops where a PT trip may end (the last alighting stops) U_{sm}^E are determined a priori. This does not necessarily mean that all stops in these sets are actually used: this depends on the attractiveness of the connection by PT between each pair of stops.

From the set of stops at the destination U_{sm}^E all PT lines that serve these stops are followed backwards towards the beginning of these PT lines. At every stop u upstream along the PT line the generalised costs $\breve{c}_{lum}^s$ are calculated, representing the costs to reach destination s from the moment immediately after boarding line l at stop u using mode chain m (so excluding waiting time).

When these generalised cost are calculated for the stops at all lines reachable from the candidate set of stops around the destination U_{sm}^E, the probabilities (fractions) p_{lum}^s for boarding the PT lines that serve stop u are calculated using Eq. 3.19, using generalised costs $\breve{c}_{lum}^s$ and using the frequency of line l. The frequency is relevant to model a situation where a traveller is waiting on the platform and a somewhat slower, but high frequency service is arriving. In that case the traveller is earlier at its destination when boarding this vehicle, instead of waiting for the next faster service.

$$p_{lum}^s = \frac{F_l e^{-\theta_4 \breve{c}_{lum}^s}}{\sum_{x \in L_{um}^s} F_x e^{-\theta_4 \breve{c}_{xum}^s}} \qquad (3.19)$$

From all the PT lines passing through a stop, only some provide an effective means of reaching the chosen destination. Lines may travel away from the destination, or at least not towards it. Others may travel towards the destination, but require too many transfers or travel too slowly to be effective. In order to reduce the number of options, the following criterion is used to determine the set of candidate lines L_{um}^s at stop u to reach destination s: given a PT line l which is immediately departing, if there is another PT line which has lower expected generalised cost even if it involves waiting its full average headway, then PT line l is deemed to be illogical, and is ruled out of L_{um}^s.

Step 4: transfers

When transferring between lines, the transfer can occur at the same stop, but it is also possible to transfer between lines with a (short) walk between 2 different stops. For every stop a set of possible transfer stops is identified based on a distance criterion. So all stops within a certain radius are considered as stops where a transfer may take place. These stops need to be connected by a pedestrian network. So for every possible stop u in the network a set of transfer stops is calculated, referred to as the transfer candidate set U_{um}^T.

Step 5: iterate to include transfers

From the set of stops reached during line choice, the process is repeated for routes that include a transfer (including the stops that can be reached by walk transfers) until a given maximum number of transfers $\hat{n}t_{\max}$. All trips that require more transfers are not considered.

Step 6: stop choice

Since a user may consider several lines at a stop, the waiting time at a stop is based on the combined frequency of lines in the candidate set L_{um}^s. The combined waiting time for stop u is defined in Eq. 3.20. By using this formula, the most attractive line fully contributes to the combined frequency, while less attractive lines only contribute to the combined frequency proportional to the attractiveness value $e^{-\theta_4 \breve{c}_{lum}^s}$ with respect to the attractiveness of the most attractive line.

$$tw_{um}^s = 1/(2 \sum_{l \in L_{um}^s} F_l \frac{e^{-\theta_4 \breve{c}_{lum}^s}}{\max_{x \in L_{um}^s} e^{-\theta_4 \breve{c}_{xum}^s}}) \tag{3.20}$$

$$\breve{c}_{um}^{rs} = \sum_{l \in L_{um}^s} p_{lum}^s \breve{c}_{lum}^s + tw_{um}^s + \breve{c}_{um}^r \tag{3.21}$$

The generalised costs to reach a destination s from origin r using stop u are calculated (Eq. 3.21), comprising of weighted average of the costs to reach the destination per line, of waiting time and of generalised costs of the access mode, between the origin r and each stop u. For each stop u in the access candidate set U_{rm}^A and for each origin, the stop choice fractions p_{um}^{rs} are determined using Eq. 3.22.

$$p_{um}^{rs} = \frac{e^{-\theta_3 \breve{c}_{um}^{rs}}}{\sum_{u \in U_{rm}^A} e^{-\theta_3 \breve{c}_{um}^{rs}}} \tag{3.22}$$

```
Step 1: Initialisation
For each stop u ∈ U
      Set čˢ_um to ∞
      For each line l ∈ Lˢ_um
            Set čˢ_lum to ∞
      End
End
For each stop u ∈ Uᴱ_sm
      Set čˢ_um to egress costs
End
Set initial stop set to be processed Uᴾ_n̂t to Uᴱ_sm
Set number of transfers counter n̂t to 0

Step 2: propagate backward along transit lines
For each u ∈ Uᴾ_n̂t
      For each line l ∈ L_u
            For each stop u′ prior to u on line l
                  If the currently recorded costs are higher,
                        Update costs čˢ_lu′m
                        Add u′ to Uᴾ_n̂t+1
                  End
            End
      End
End

Step 3: Update stop costs
For each u ∈ Uᴾ_n̂t+1
      Calculate line boarding fraction pˢ_lum by using Eq. 3.19
      Calculate the weighted average of line costs čˢ_lum (Eq. 3.21)
      and add waiting time tw_um (Eq. 3.20)
      If the currently recorded costs are higher
            Update costs čˢ_um
      End
End

Step 4: Include walk transfers
For each u ∈ Uᴾ_n̂t+1
      For each u′ ∈ Uᵀ_u
            If the currently recorded costs are higher than costs
            using a walk transfer
                  Update čˢ_um
                  Add u′ to U^Tused_n̂t
            End
      End
End

Step 5: Update and iterate
Set n̂t = n̂t + 1
Set Uᴾ_n̂t = Uᴾ_n̂t−1 + U^Tused_n̂t−1
If n̂t = n̂t_max
      Stop
Else
      Go to step 2
End
```

```
Step 6: Calculate access stop choice
For each  r ∈ R
     Calculate  p_{um}^{rs}  using Eq. 3.22 for each stop in  U_{rm}^{A} , based on
     generalised costs per stop  č_{um}^{rs} , including access costs
End
```

Algorithm 3.1: PT assignment algorithm

Summary

Altogether, the PT assignment algorithm can be summarised in the pseudo code in Algorithm 3.1. The result of the code is a path tree from all origins in the network to one destination *s*, using one specified mode chain *m*. By combining stop fractions p_{um}^{rs} and line fractions p_{lum}^{s}, the result of this algorithm is a set of PT route fraction ζ_{km}^{rs} per mode chain.

3.2.2.   Car assignment

For car the generalised costs of a route *k* consist of travel time and distance. Travel time on a link is dependent on the flow, give a freight pre-load. Distance represents fuel costs and other variable costs, for example maintenance costs (see Eq. 3.23).

$$c_k^{rs} = \alpha_1 di_k^{rs} + \alpha_2 tv_k^{rs} \qquad (3.23)$$

Freight trips are loaded to the empty network using all-or-nothing assignment. The car-only trips are assigned to the network assuming a capacity-dependent user equilibrium, taking the freight pre-load (expressed in Passenger Car Equivalents) into account. The car assignment problem is solved by the standard algorithm of Frank-Wolfe. This implies that all used routes have the same generalised costs, so the OD based generalised costs can be defined as follows: $c^{rs} = \min_k c_k^{rs}$.

The car travel times can be determined based on the flow following a BPR curve (see Eq. 3.24). The car travel time tv_a on link *a* depends on the car flow x_a, where tv_a^0 is the free flow travel time on link *a*, $\tilde{x}_a$ is the (fixed) freight flow on link *a*, x_a^{max} is the capacity of link *a*, and β_1 and β_2 are parameters. Travel time tv_k^{rs} and distance di_k^{rs} of a route is the sum of all link travel times / link lengths of the links of that route.

$$tv_a = tv_a^0(1 + \beta_1((x_a + \tilde{x}_a) / x_a^{max})^{\beta_2}) \qquad (3.24)$$

3.2.3.   Mode choice including multiple access and egress modes

Using the weighted average costs of the mode chains $\hat{c}_m^{rs}$ and the costs of car c_{rs}, a nested logit model is used as a mode choice model. The use of a logit model as a choice model implies variation in preferences and generalised costs perception among travellers. Depending on the costs per mode (chain), a distribution over the modes and mode chains is calculated using a nested logit model (Ben-Akiva and Bierlaire, 1999). This step splits the total OD matrix into several OD matrices: one for every mode chain. Within the nested logit model, two nests are used: one for the mode car and one for all PT mode chains. The upper nest is formulated in Eq. 3.3 and the lower nest in Eq. 3.12. The (generalised) costs of a mode chain are the weighted average of the used routes, as defined in Eq. 3.10.

The OD matrices per mode chain are iteratively assigned to the multimodal network (using the PT route choice algorithm described in Section 3.2.1). For each iteration of the modal split model updated costs are used for car and for PT chains that include car as access or egress (see Figure 3.5). The costs of other PT chains do not need to be updated, because these costs do not depend on flow. To prevent this iterative process from alternating, the new OD matrix is a weighted average of the matrix of iteration *n* (with weight *1/n*) and the old matrix of iteration *n-1* (weight *(n-1)/n*). A fixed number of iterations $n_{max} = 8$ is executed, which is a trade-off between computation time and convergence to user equilibrium (i.e. the OD matrices and costs of iteration 7 and iteration 8 only show little difference).
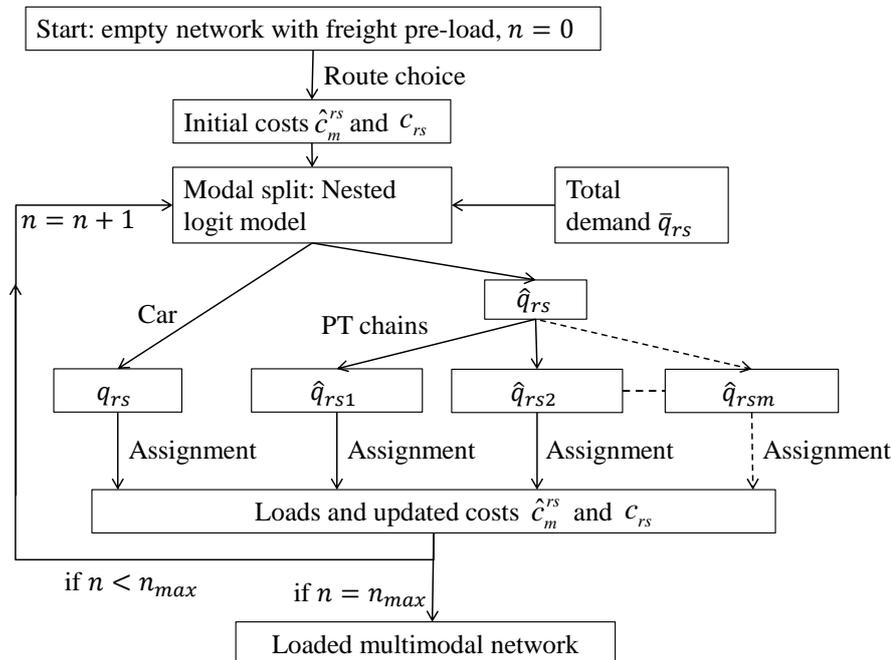
Figure 3.5: Flow chart of the lower-level model

3.3.   Objective functions

The three main aspects of sustainability (environmental, social and economic) can be made operational in objectives in various ways. The objectives to be used during this research were chosen during a workshop with policy officers and scientists from the SRMT project. A longlist of possible objectives (that can be assessed using the effect models available) was presented to the policy officers and scientists. Each policy maker and scientist chose his or her top three of most important objectives and top three of least important objectives. To be sure that all three aspects of sustainability are represented, the most popular objective to represent environmental sustainability was chosen, as well as the most popular objectives to represent social sustainability and to represent economic sustainability.

This resulted in total travel time to represent economic sustainability, as an easy to understand measure for accessibility, which is suitable because a fixed demand is assumed (and therefore is equivalent to average travel time). Urban space used by parking represents social sustainability, because space used for parking is privately used, instead of being available to

the society as a whole. CO_2 emissions represent environmental sustainability: the policy officers and scientists in the workshop chose CO_2 emissions (directly related to energy use) to be the most important problem. They expected that other problems like air quality and noise disturbance have easier technological solutions, making them less relevant in the future.

For the regional government (who are responsible for operating PT and therefore usually decide on the considered type of measures) operating costs are more relevant than investment costs in infrastructure, because these are annual costs. These annual costs are usually agreed upon for several years, because PT services are included in long-term contracts with PT operators that include investments in rolling stock. Furthermore, costs for investment and maintenance of rail infrastructure are irrelevant for the regional government, because these costs are usually paid by the national government. For the chosen measures investment costs are limited anyway: some measures (i.e. frequencies of bus lines) do not have investment costs and the other measures (i.e. new train stations) have limited costs (when compared to constructing a complete new rail line). Furthermore, investments in P&R facilities are included in their operating costs, regardless of the way it is financed, because to be attractive for users these operating costs should be paid by the regional government instead of the traveller. If the measures lead to more passenger-kilometres travelled the revenues of PT system increase, which is a compensation for the additional operating costs. Concluding, the costs of the measures are assessed by the resulting operating deficit: the difference between PT operating costs and operating revenues.

Another result from the workshop with policy makers is that costs are taken into account as an objective rather than as a budget constraint. First of all, this may result in finding solutions that have a better performance in terms of the other objectives, and also have lower costs (solutions that dominate other solutions). When it comes to trade-off between costs and benefits, the policy officers are interested in the marginal costs of reducing the other objectives: it is relevant to know whether a little more can be spend to reduce for example CO_2 emissions further. The other way around, it does not make sense to spend the entire budget if the final expenses only cause small improvements in policy objectives. These insights do not become available when costs are only taken into account as a (budget) constraint. Such a flexible budget better represents reality, because there is always interaction between the transportation department of a government and other departments of the government. Therefore operating deficit is defined as fourth objective function to represent cost-efficiency.

The objective functions are formulated mathematically in Table 3.1 and are all to be minimised. These values are calculated after applying the mode choice and traffic assignment models presented in Section 3.2. The values of the objective functions are calculated using loads and travel times in the network and using mode specific demand matrices that result from applying the lower-level model. This information is link based or OD based (matrices with size $|R| \times |S|$). The time scale of these objective values corresponds with the used lower-level model, e.g. when the model represents a one hour time period, the values of the objective values are also expressed for this one hour time period.

Total travel time in the network is calculated from travel times for car and PT, following from the OD based travel time components of the generalised cost functions (Eq. 3.17 for PT, using

the travel time and waiting time component and Eq. 3.23 for car, using the travel time component). Urban space used by parking is represented by the number of car trips to or from zones that are classified as highly urban. Such a trip requires a parking space that cannot be used for other urban land uses. Both origins (at residential areas) and destinations (at commercial areas) are considered to be relevant here, because in both situations the parking space needs to be present and double use (during the night for residents and during daytime for employees) is rare in practice. In less urban and rural areas this is considered to be irrelevant, because space is less scarce at those locations. Alternative urban land uses can for example be residential areas, parks or pedestrian zones. Although such alternative land uses can also be achieved by facilitating underground parking, this is often not done due to the large investments needed for such a facility. Implicit assumption here is that car ownership is gradually adapted as a result of less car trips, so that former parking spaces structurally come available for other purposes. CO_2 emissions are the main human cause of climate change (IPCC, 2007), so CO_2 emissions are a proper indicator for climate impact. The emissions for car E^{CO_2} are calculated for the year 2015 using the ARTEMIS traffic situation based emission model, where emission factors depend on link speed and volume / capacity ratio (INFRAS, 2007). The emission factors for PT $\hat{E}_b^{CO_2}$ are distance based with distinction per vehicle type and are taken from Dutch practice. Finally, operating deficit includes operating costs of PT (including depreciation of investments in vehicles) and P&R (including depreciation of investments), from which PT revenues are subtracted. The PT operating costs and fares are taken from Dutch practice and the costs for P&R are taken from Van Ommeren and Wentink (2012).

Table 3.1: Formulation of the four objective functions

Aspect of sustainability	Indicated by	Formulation	
Accessibility (economic)	Total travel time (TTT)	$Z_1 = \Omega \sum_{r \in R, s \in S} tv_{rs} q_{rs} + \sum_{r \in R, s \in S} \sum_{m \in M} \psi_m^{rs} \hat{q}_{rs} \sum_{k \in \hat{K}_m^{rs}} \zeta_{km}^{rs} (\hat{tv}_{km}^{rs} + \hat{tw}_{km}^{rs})$	(3.25)
Use of urban space by parking (social)	Number car trips to and from urban zones (USU)	$Z_2 = \sum_{r \in R_U} \sum_{s \in S} (q_{rs} + \sum_{m \in M_R} \psi_m^{rs} \hat{q}_{rs}) + \sum_{r \in R} \sum_{s \in S_U} (q_{rs} + \sum_{m \in M_S} \psi_m^{rs} \hat{q}_{rs})$	(3.26)
Climate impact (environmental)	CO_2 emissions (CE)	$Z_3 = \sum_{a \in A} x_a E^{CO_2} \left(v_a^0, \frac{x_a + \tilde{x}_a}{x_a^{max}} \right) d_a + \sum_{a \in A} \sum_{b \in B} \sum_{l \in L} \Delta_{bl} x_{al} \hat{E}_b^{CO_2} (v_a^0, \frac{x_a + \tilde{x}_a}{x_{max}}) d_a$	(3.27)
Cost-efficiency (economic)	PT operating deficit (operating costs – operating revenues) (OpD)	$Z_4 = \sum_{b \in B} \sum_{l \in L} [\Delta_{bl} O_b \sum_{a \in A_l} x_{al} t_{al}] + O_{PR} \sum_{r \in R} \sum_{s \in S} \sum_{m \in M_R \cup M_S} \psi_m^{rs} \hat{q}_{rs}$ $- \sum_{b \in B} \sum_{l \in L} [\Delta_{bl} F_b \sum_{a \in A_l} d_a \hat{x}_{al}]$	(3.28)

Explicitly taking into account the most important objective per main aspect of sustainability implies that other objectives related to sustainability are not taken into account. However, it is expected that also other sustainability objectives are improved, like noise hindrance, traffic safety and emissions of local pollutants, like NO_X, SO_2, CO, particulate matter and hydrocarbons. Just like CO_2 emissions, these objectives relate to the number of vehicle kilometres, yet with different emission factors. An important difference between CO_2

emissions and these other objectives is that the location of the emission / exposure is relevant: emission / exposure in populated areas is more harmful than in rural areas. The latter is related with the objective urban space used by parking, because the number of car trips to and from highly urban zones has a relation with the number of car kilometres driven in highly populated areas. Therefore it is expected that other sustainability objectives are at least partly in line with the objectives CO_2 emissions and urban space used, so that optimising CO_2 emissions and urban space used also leads to improving other sustainability objectives.

3.4. Case studies

Following from the context of this research (being a part of the Sustainable Accessibility of the Randstad research programme, as described in Chapter 1), the study area for the case studies is in the Randstad area in The Netherlands. In an ideal world, the study area would cover the whole Randstad area, but this is not feasible. First, the computation time to model the behavioural response of the travellers to network designs would be too large. Second, no existing transportation model is available for the whole Randstad area that contains both a detailed network of car and PT.

The Randstad consists of four city regions: Amsterdam, Rotterdam, Den Haag and Utrecht (see Figure 3.6). Each of these city regions has a transportation model that contains car, bicycle and PT, so in principle all these models could act as a starting point for a case study in this research. In the user group, policy officers from the city region of Amsterdam actively participated, and made the transportation model of the region available. Therefore, this model (with the name 'Venom' (Kieft, 2013)) is used as a basis for the case study.
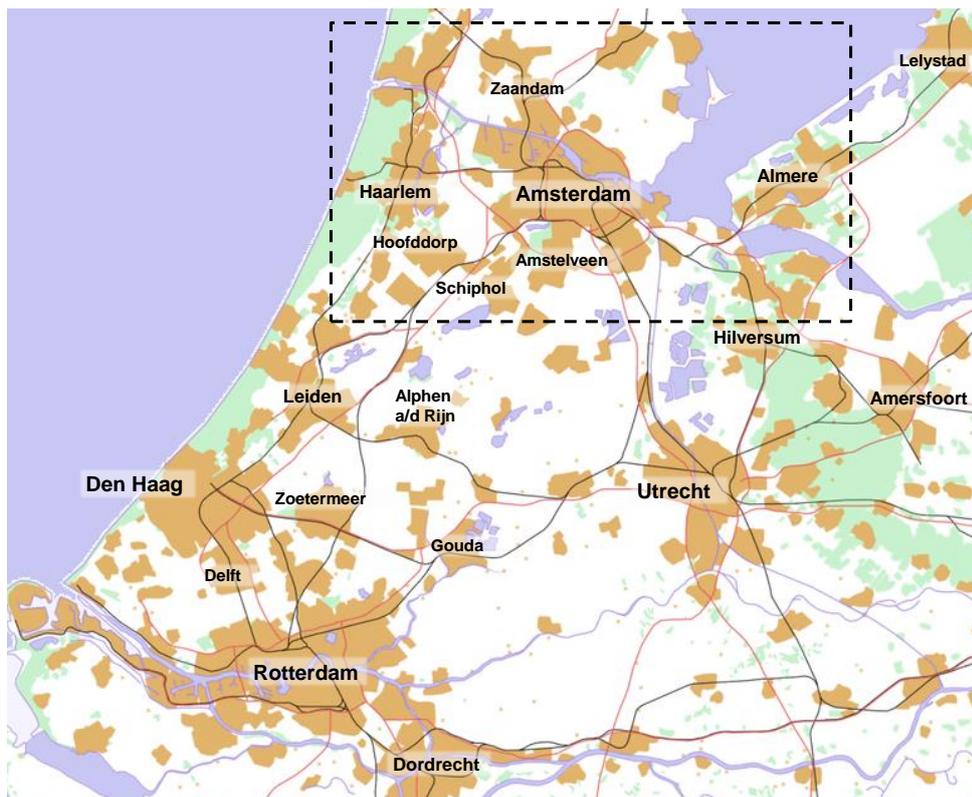
Figure 3.6: The Randstad area. The AMA is indicated by the dashed rectangle

Two case studies are used. Both case studies involve approximately the same study area in the Amsterdam Metropolitan Area (AMA) in the Netherlands (indicated by the dashed rectangle in Figure 3.6). The largest cities and towns in this area are Amsterdam, Haarlem, Almere, Zaandam, Hoofddorp and Amstelveen. It also contains the airport Schiphol. This area covers a large part of the Randstad and has an extensive multimodal network (car, train, metro, tram, bus, bicycle, walking).

The first case study (described in Section 3.4.2) has less OD pairs and decision variables, resulting in lower computation times. The second case study (described in Section 3.4.3) has more OD pairs and decision variables, resulting in more realistic modelling results, but also in a larger computation time. Both case studies are used to demonstrate the modelling framework by showing that the framework can be applied on a real-world scale. Furthermore, the first case study is used to test parameter settings of the solution algorithm as well as to investigate the influence of randomness in the algorithm. The second case study is used to test a faster optimisation algorithm, to test robustness of the results with respect to transportation demand input and to present the optimisation results such that they provide information that can be used as decision support information for policy makers in the study area.

3.4.1. Methods and parameter values used in both case studies

Data aggregation

In both case studies, the regional transportation model of the AMA (named Venom) is used as a basis (Kieft, 2013). This model contains a complete multimodal network (car, PT, bicycle and walking) and has 3722 transportation zones. The base year demand (for the year 2004) is available for car and for PT and is calibrated based on counts in both the car and the PT network. The demand forecasts for the year 2030 are based on the Dutch national scenarios for economical and spatial developments (WLO-scenarios) (CPB et al., 2006). The Venom model distinguishes car-only, bicycle-only and PT as separate modes in the matrix estimation process, but only for car and for PT a calibrated matrix is available. The demand of car and PT according to the Venom model is aggregated, resulting in a total demand matrix that is used as input for the modal split / route choice model. To reduce computation time, the demand matrix is aggregated into a smaller number of centroids (42 for case study 1 and 102 for case study 2), where zones are chosen such that it is still possible to distinguish between several access and egress modes. Especially zones outside of the study area are aggregated, which resulted in unused network elements that were removed (links, nodes, stops and PT lines). Internal traffic in the aggregated zones is removed from the matrix. All transportation demand within, to, from and through the study area is included in the OD matrix. Other demand is removed from the matrix (i.e. the corresponding cells are set to zero).

Parameters

Several parameters in the mode / route choice models need values when applied in the case studies. In both case studies the same values are used for these parameters. These values are derived from literature and from Dutch modelling practice and are used to calibrate the model, e.g. to achieve a realistic modal split.

The generalised costs function for car is expressed in euros, with $\alpha_1 = 0.15$ and $\alpha_2 = 6$. α_1 translates distance into euros and includes al distance based monetary costs, like fuel costs and other variable costs, for example maintenance costs. α_2 is the Value of Time (VoT), with travel time expressed in hours, and translates travel time into euros. In the generalised costs functions for PT the $\hat{\alpha}$ values are shown in Table 3.2. Monetary costs are derived from distance, so the $\hat{\alpha}_1$ value represents fare per kilometre. The travel time and waiting time are expressed in hours and translated into euros by using a VoT. Waiting time is valued twice as heavy as in-vehicle travel time (Wardman, 2004). Moreover, rail based PT vehicles are valued 20% more positively than bus (Bunschoten et al., 2013). A transfer is valued equal to 5 minutes of in-vehicle time (translated to euros this comes to 0.5 for bus and 0.42 for other PT modes).

The maximum number of transfers in the PT assignment algorithm $\hat{n}t_{max} = 4$. For walk and bicycle access and egress the minimum number of stops to find is set to 3. The search radius for candidate stops is set to 1.2 km for walking and to 5 km for bicycle. For walking, at least one train station is included in the candidate set and for bicycle at least one express train station. For car as access mode, a list of stops that have P&R facilities available predefines the candidate set. The logit parameter for stop choice (θ_3) is equal to 4 and the logit parameter for line choice (θ_4) is equal to 2.

Table 3.2: $\hat{\alpha}$ coefficients in the generalised costs function for PT in the case studies

Vehicle type b	Walking	Bicycle	Bus	Tram	Metro	Local train	Express train
$\hat{\alpha}_1$	0	0.01	0.07	0.07	0.07	0.07	0.07
$\hat{\alpha}_2$	6	6	6	5	5	5	5
$\hat{\alpha}_3$	-	-	12	10	10	10	10
$\hat{\alpha}_4$	-	-	0.5	0.42	0.42	0.42	0.42

The parameters in the BPR curve in the car model are $\beta_1 = 1$ and $\beta_2 = 4$. The logit parameter for mode choice (θ_1) is equal to 0.5 and the logit parameter for PT mode chain choice (θ_2) is equal to 1. Average car occupation Ω is set to 1.1 passengers per vehicle.

For the evaluation of the objective function values, the parameter values per vehicle type b for PT can be found in Table 3.3. Furthermore, costs to operate a P&R space O_{PR} are set to 2.8 euros per day. Emissions factors E^{CO_2} for car are dependent on speed and I/C ratio and can be found in detail in INFRAS (2007).

Table 3.3: Parameters to calculate objective function values in the case studies

Vehicle type b	Bus	Tram	Metro	Local train	Express train
Fare C_b (euros per km)	0.09	0.09	0.09	0.09	0.09
Operating costs O_b (euros per hour)	100	180	250	600	600
Emission factor $\hat{E}_b^{CO_2}$ (gram / km)	900	1800	2250	4500	4500

3.4.2. Description of case study 1

Study area

The first case study focusses on the Western part of the AMA (see Figure 3.7), covering Haarlem, Schiphol, Hoofddorp and the Western part of Amsterdam. It contains a detailed multimodal network, including bicycle links, car links and PT service lines. These lines comprise 404 bus lines, 44 tram / metro lines and 124 train lines (including distinction between local services and express services). Consequently, the PT vehicle type set B contains bus, tram, metro, local / stop trains and express trains, or shorter $B \in \{bu, tr, me, st, et\}$). This enables a detailed modelling of mode chains. On the other hand, the number of zones is limited, to ensure that the computation time is low enough. Important commercial areas are the city centres of Amsterdam and Haarlem (zones 17 and 1), the business district in the South of Amsterdam (zone 22), the harbour areas (zones 15 and 4) and the airport Schiphol (zone 26 and 27). The other zones are mainly residential, but still contain small or medium scale commercial activities.

✳ Zone and zone number
28

／ Railway with train numbers that are decision variables
3,4

6,9 Bus numbers that are decision variables

12 ↕ Candidate location for tram line extension

16 ● Candidate location for train station

20 ● Candidate location for express train status of train station

26 ● Candidate location for park and ride facility
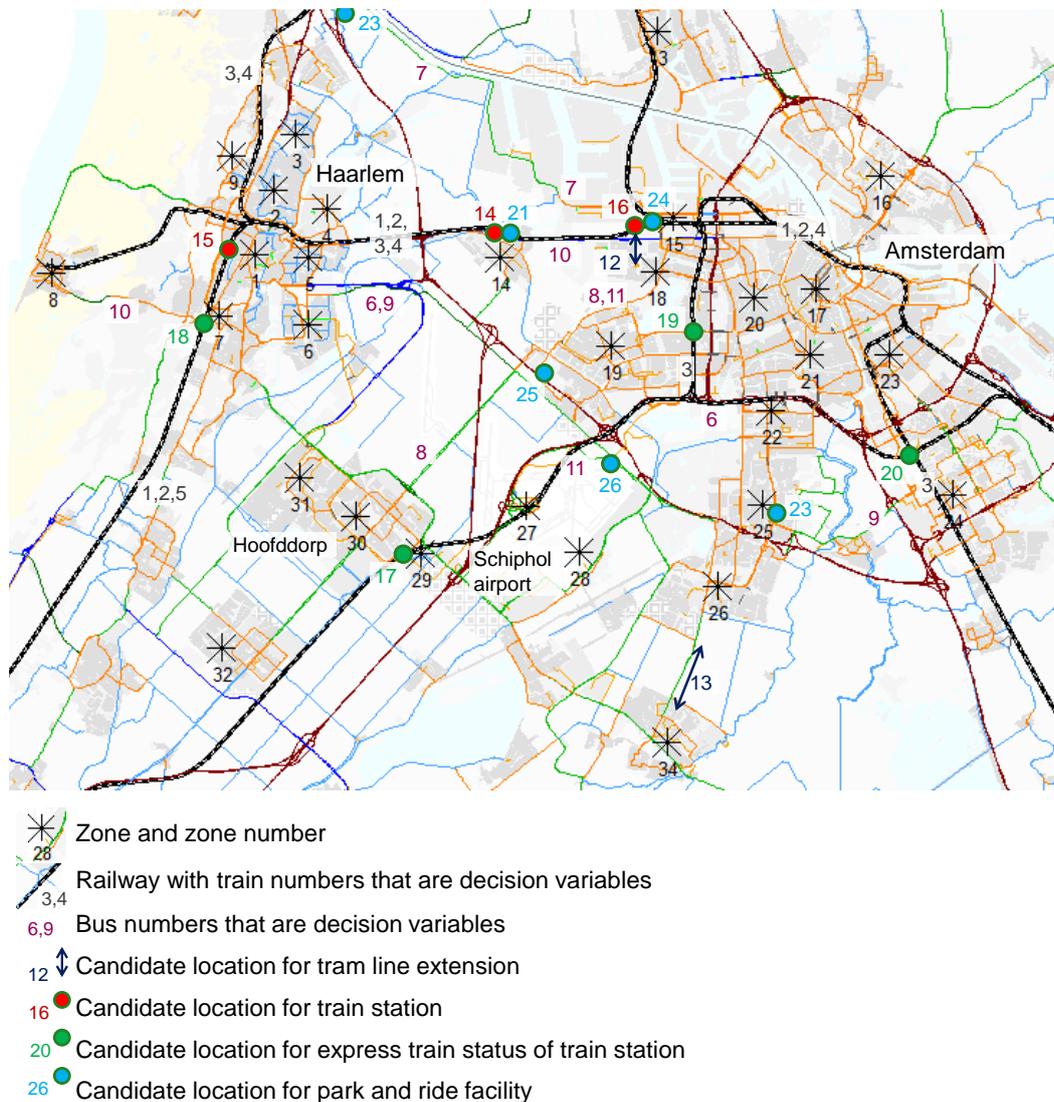
Figure 3.7: Study area of case study 1, including the geographical position of decision variables

Passenger transportation demand $\bar{q}_{rs}$ and freight demand $\tilde{q}_{rs}$ for a one hour period in the AM peak on an average working day in the base year 2004 are used. The total number of passenger trips is 248,866 and the total number of freight trips is 18,490. The resulting objective function values are also calculated for a one hour period in the AM peak for the whole network, so including the network outside of the study area, since the flows outside of the study area may change due to measures in the study area. Although the impact of a network solution also depends on the other periods of the day and week, the AM peak is studied because it is the busiest period: the transportation related problems are most relevant during this period. A plot with loads and I/C ratios is given for case study 2 in Figure 3.9, which has comparable loads as case study 1.

Decision variables

Case study 1 has 26 decision variables y_v in decision vector $\underline{y}$ (see Figure 3.7 and Table 3.4). These variables are of several types: frequencies of train lines, frequencies of bus lines, tram line extensions, existence of train stations, express train status of train stations and existence of P&R facilities. The list of variables is based on regional policy documents and on interviews with involved policy makers. For variable 1 to 11 that represent frequency, 2 or 4 different values are possible. For the elements of decision vector $\underline{y}$ values are chosen between 0 and 1, to make comparisons between decision variables possible. The decision variables in these cases correspond to network variable F_l for the relevant lines l (in all cases both directions of a PT service line are set to the same frequency). Opening / closure of train stations and express train status of train stations are binary variables and influence network property D_u^T. Opening / closure of P&R facilities is also a binary variable that influences network property D_u^P. Measures are only included as a candidate location if spatial and capacity constraints are met. For example, a P&R location is only potentially opened if the corresponding stop is served by PT. In total the feasible region Y contains approximately $5*10^9$ possible decision vectors. The characteristics of lines and stops that are not candidate locations are fixed at one value. Furthermore, the car and bicycle networks are assumed to be fixed. The evaluation of one solution takes approximately 3.25 minutes (using a computer with an Intel® Core™ i7 CPU 860 @ 2.8GHz and a 4 GB RAM).

Table 3.4: Definition of decision variables for the small case study

Possible values of y_v	Represents real value	Description of decision variable
$y_1 \in \{\frac{1}{3}, \frac{2}{3}\}$	$F_1(y_1) \in \{2,4\}$	Frequency of express train Amsterdam - Haarlem - Leiden
$y_2 \in \{0,1\}$	$F_2(y_2) \in \{0,2\}$	Frequency of local train Amsterdam - Haarlem – Leiden
$y_3 \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$	$F_3(y_3) \in \{0,2,4,6\}$	Frequency of new train line Uitgeest - Haarlem - Amsterdam Sloterdijk - Zuid – Bijlmer
$y_4 \in \{0, \frac{1}{3}\}$	$F_4(y_4) \in \{0,2\}$	Frequency of train line Uitgeest - Haarlem - Amsterdam Centraal
$y_5 \in \{0, \frac{1}{3}\}$	$F_5(y_5) \in \{0,2\}$	Frequency of train line Haarlem – Leiden
$y_6 \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$	$F_6(y_6) \in \{0,2,4,6\}$	Frequency of bus line 176: Amsterdam Zuid – Haarlem

$y_7 \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$	$F_7(y_7) \in \{0, 2, 4, 6\}$	Frequency of bus line 83: IJmuiden – Amsterdam Sloterdijk
$y_8 \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$	$F_8(y_8) \in \{0, 2, 4, 6\}$	Frequency of new bus line Hoofddorp - Amsterdam West - Sloterdijk
$y_9 \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$	$F_9(y_9) \in \{0, 2, 4, 6\}$	Frequency of bus line 175: Haarlem - Amstelveen - Amsterdam Zuidoost
$y_{10} \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$	$F_{10}(y_{10}) \in \{0, 2, 4, 6\}$	Frequency of bus line 80: Zandvoort - Haarlem - Amsterdam Marnixstraat
$y_{11} \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$	$F_{11}(y_{11}) \in \{0, 2, 4, 6\}$	Frequency of new bus line 'Westtangent': Amsterdam Sloterdijk – Amsterdam West – Schiphol
$y_{12} \in \{0, 1\}$	$F_{12}(y_{12}) \in \{0, 8\}$	Extension of tram 13 in Geuzenveld
$y_{13} \in \{0, 1\}$	$F_{13}(y_{13}) \in \{0, 8\}$	Extension of express tram Amstelveen to Uithoorn
$y_{14} \in \{0, 1\}$	$D_1^T(y_{14}) \in \{\{\}, \{st\}\}$	New train station Halfweg-Zwanenburg
$y_{15} \in \{0, 1\}$	$D_2^T(y_{15}) \in \{\{\}, \{st\}\}$	New train station Haarlem Zuid
$y_{16} \in \{0, 1\}$	$D_3^T(y_{16}) \in \{\{\}, \{st\}\}$	New train station Amsterdam Geuzenveld
$y_{17} \in \{0, 1\}$	$D_4^T(y_{17}) \in \{\{st\}, \{st, et\}\}$	Express train status of station Hoofddorp
$y_{18} \in \{0, 1\}$	$D_5^T(y_{18}) \in \{\{st\}, \{st, et\}\}$	Express train status of station Heemstede-Aardenhout
$y_{19} \in \{0, 1\}$	$D_6^T(y_{19}) \in \{\{st\}, \{st, et\}\}$	Express train status of station Amsterdam Lelylaan
$y_{20} \in \{0, 1\}$	$D_7^T(y_{20}) \in \{\{st\}, \{st, et\}\}$	Express train status of station Duivendrecht
$y_{21} \in \{0, 1\}$	$D_1^P(y_{21}) \in \{0, 1\}$	P&R at train station Halfweg-Zwanenburg
$y_{22} \in \{0, 1\}$	$D_2^P(y_{22}) \in \{0, 1\}$	P&R Velsen Zuid (at bus line 83)
$y_{23} \in \{0, 1\}$	$D_3^P(y_{23}) \in \{0, 1\}$	P&R at tram station Amstelveen Oranjebaan
$y_{24} \in \{0, 1\}$	$D_4^P(y_{24}) \in \{0, 1\}$	P&R at train station Geuzenveld and tram 13
$y_{25} \in \{0, 1\}$	$D_5^P(y_{25}) \in \{0, 1\}$	P&R Badhoevedorp at new bus line
$y_{26} \in \{0, 1\}$	$D_6^P(y_{26}) \in \{0, 1\}$	P&R Schiphol Noord at bus lines 'Westtangent' and 'Zuidtangent'

3.4.3. Description of case study 2

Case study 2 covers a larger area than case study 1, such that it covers the entire AMA in The Netherlands (Figure 3.8). This area has an extensive multimodal network with pedestrian, bicycle, car and PT infrastructure. PT consists of 586 bus lines, 42 tram and metro lines and 128 train lines, which include local / stop trains and express trains (just like in case study 1 the PT vehicle type set is defined as $B \in \{bu, tr, me, st, et\}$). Bicycles can be parked at most stops. A selection of PT stops facilitate P&R transfers. Origins and destinations are now aggregated into 102 transportation zones.

Passenger transportation demand $\bar{q}_{rs}$ and freight demand $\tilde{q}_{rs}$ for a one hour period in the AM peak on an average working day in the forecast year 2030 are used. The total number of passenger trips is 407,676 and the total number of freight trips is 35,839. The resulting objective function values are also calculated for a one hour period in the AM peak for the whole network. This includes the network outside of the study area, because the flows outside of the study area may change due to measures in the study area. Although the impact of a network solution also depends on the other periods of the day and week, the AM peak is studied because it is the busiest period: the transportation related problems are most relevant during this period.
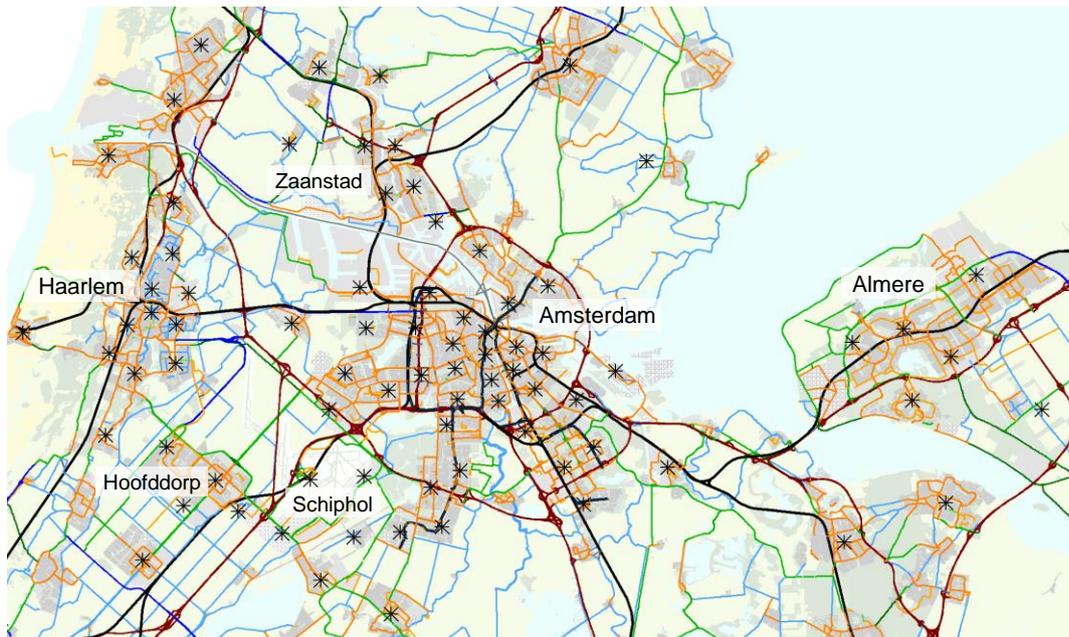
Figure 3.8: Map of the study area, showing origins / destinations, railways, roads

Network flows in the reference situation

In Figure 3.9 the flow through the car network is shown for the reference situation. It can be observed that many roads are congested, mainly in the direction of Amsterdam (since an AM peak is modelled). Bottlenecks are visible at motorways, but also on smaller roads.
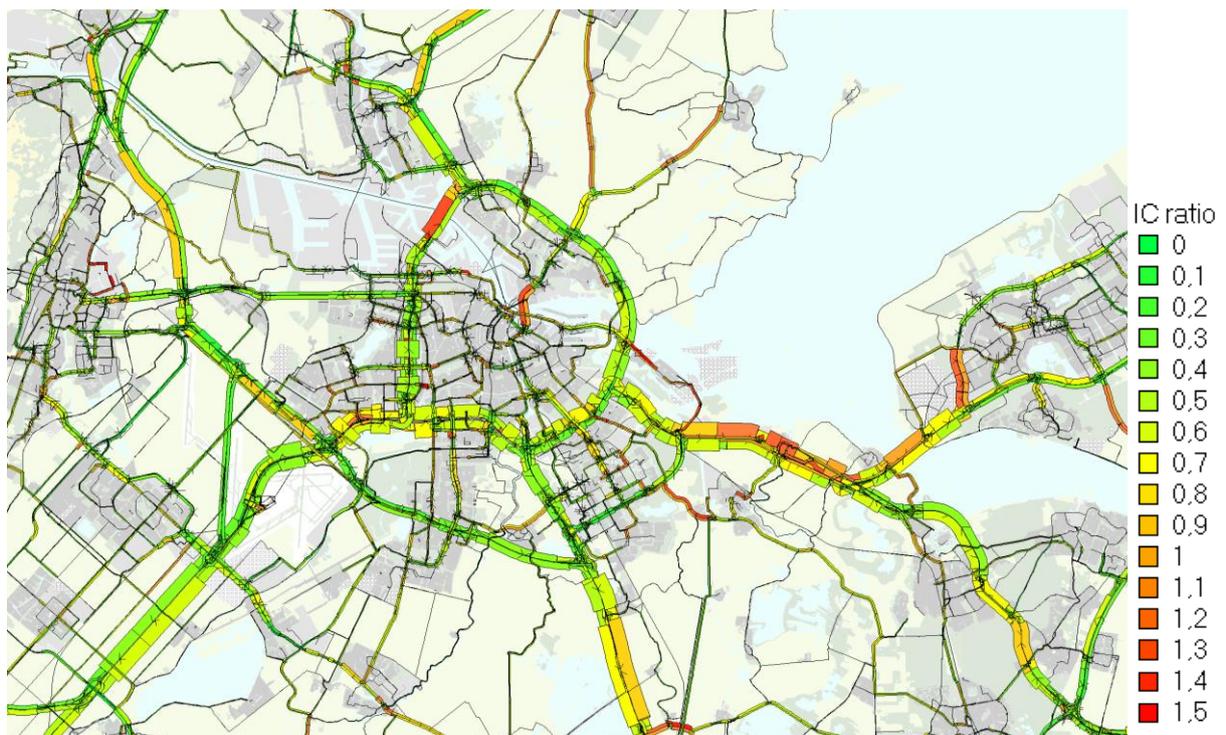
Figure 3.9: Network flows in the reference situation (width indicates flow, colour indicates I/C ratio)

Decision variables

In the network of the study area, 37 decision variables are defined related to transfer facilities or to PT facilities. The number of decision variables is larger than in case study 1, because the study area is larger in this case study. The list of variables is based on regional policy documents and on interviews with involved policy makers. For every potential measure in the network, a decision variable y_v is defined (see Figures 3.10-3.13 and Table 3.5). In this case, the feasible region Y contains approximately $7*10^{13}$ possible decision vectors. The evaluation of one solution takes approximately 6.5 minutes (using a computer with an Intel® Core™ i7 CPU 860 @ 2.8GHz and a 4 GB RAM).

Figure 3.10: Train stations that are decision variables in case study 2

Figure 3.11: P&R facilities that are decision variable in case study 2
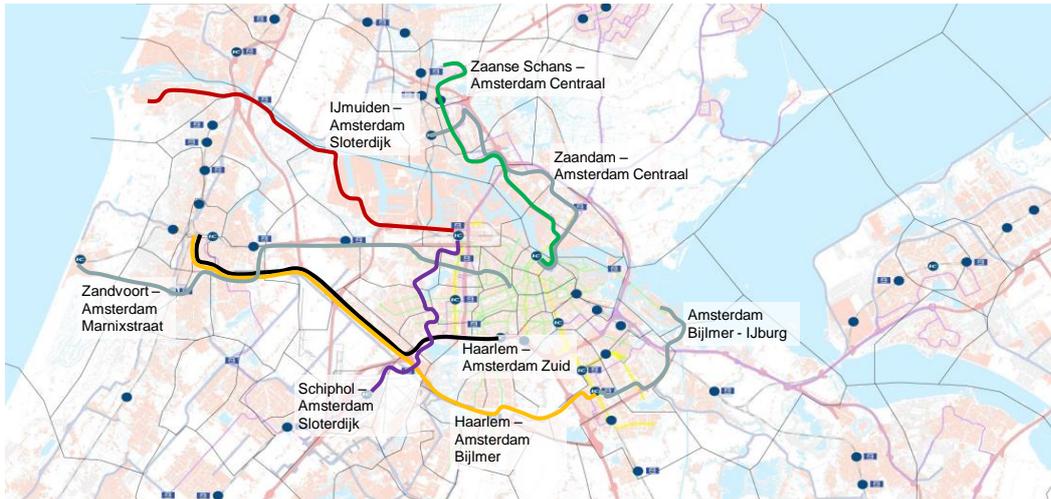
Figure 3.12: Bus lines with frequency as decision variable in case study 2
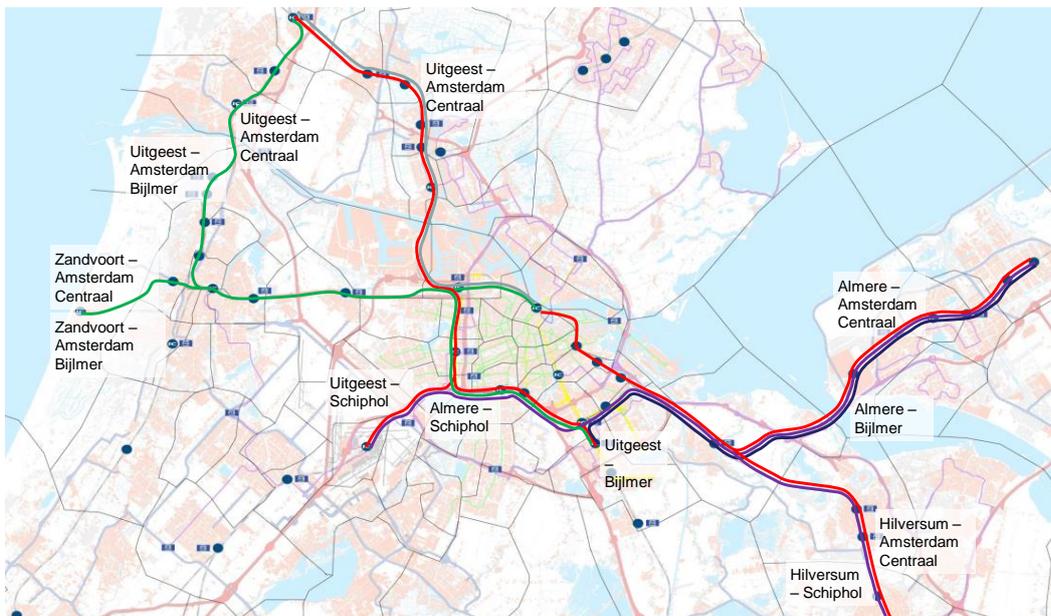
Figure 3.13: Local train lines with frequency as decision variable in the study 2

Table 3.5: Definition of decision variables for case study 2

Decision variable index v	Possible values of y_v	Represents real value	Description
1-6	$y_v \in \{0,1\}$	$D_u^T(y_v) \in \{\{\},\{st\}\}$	Opening / closure of train stations
7-9	$y_v \in \{0,1\}$	$D_u^T(y_v) \in \{\{st\},\{st,et\}\}$	Express train status of train stations
10-16	$y_v \in \{0,1\}$	$D_u^P(y_v) \in \{0,1\}$	Opening / closure of P&R facilities
17,21,22,24	$y_v \in \{0,\frac{1}{3},\frac{2}{3},1\}$	$F_l(y_v) \in \{0,4,8,12\}$	Frequency of bus lines
18,19,20,23	$y_v \in \{0,\frac{1}{3},\frac{2}{3},1\}$	$F_l(y_v) \in \{0,2,4,6\}$	Frequency of bus lines
25-32,34,36	$y_v \in \{0,\frac{1}{3}\}$	$F_l(y_v) \in \{0,2\}$	Frequency of local train lines
33,35	$y_v \in \{0,\frac{1}{3},\frac{2}{3},1\}$	$F_l(y_v) \in \{0,2,4,6\}$	Frequency of local train lines
37	$y_v \in \{0,1\}$	$F_l(y_v) \in \{0,8\}$	Extension of a tram line

3.5. Summary

In this chapter the mathematical optimisation problem is formulated. The main formulation is a bi-level mathematical problem, which is defined for the situation where infrastructural developments concerning multimodal trip making are involved. Therefore, a user equilibrium model is needed that involves mode choice over several trip chains, including car-only trips and PT trips with multiple access and egress modes. For PT assignment, multiple routing represents different preferences of travellers and variety over departure times. For car, travel times are calculated by taking congestion into account. Moreover, the problem is formulated as an MO problem, implying that the result of solving the problem is a set of Pareto-optimal solutions, rather than a singles optimal solution.

Furthermore, in this chapter the objective functions are defined. The values of objective functions are calculated based on network flows and corresponding speeds after mode choice and traffic assignment. Four objective functions are formulated, related to sustainability: travel time (accessibility) and operating deficit represent economic sustainability, urban space used by parking represents social sustainability and CO_2 emissions represent environmental sustainability.

Finally, two case studies are presented, for which the decision variables are defined in detail. These case studies are used throughout the thesis when applying the modelling framework. The solution approach is tested in Chapter 4, providing results on the influence of randomness in the genetic algorithm and providing an improvement of the optimisation results by applying an alternative MO genetic algorithm. The usefulness of MO optimisation in decision support

is illustrated in Chapter 5 and the robustness of the method for different transportation demand input is tested in Chapter 6.

Chapter 4: Solution approach

In Chapter 3 a mathematical formulation of the problem was given. This type of multi-objective optimisation problems is NP-hard and therefore cannot be solved analytically. Heuristic methods, and especially genetic algorithms, are generally used to solve these kinds of problems. More specifically, in Chapter 2 the algorithm NSGAII was identified as most suitable method to use.

In this chapter the more recent algorithm ε-NSGAII, which has had limited attention in literature and until now has no applications in the field of transportation, is applied and compared to NSGAII as a benchmark algorithm. The results show a considerable improvement in performance, especially in the early stage of running the algorithm. Based on these results, ε-NSGAII is further used as a solution algorithm for case study 2.

Furthermore, another important aspect is given attention in this chapter: the stochastic principles that are used in the genetic operators in this type of algorithms may cause variation in the resulting Pareto sets. The variation this causes, both in objective and in decision space, is relevant to take into consideration when the method is applied in practice to provide decision support information. An algorithm that produces optimisation results with small random variations can be seen as high performing. The level of variation gives an indication to what extent the decision support information provided to the policy maker is reliable. Therefore, the variation in the optimisation outcomes is analysed using the well-known algorithm NSGAII.

This chapter is partly based on:

Brands T., E. C. van Berkum (2014) Performance of a genetic algorithm for solving the multi-objective, multimodal transportation network design problem, *International Journal of Transportation* 2(1), pp. 1-20

Brands T., L. J. J. Wismans, E. C. v. Berkum (2014) Multi-objective Transportation Network Design: Accellerating Search by Applying ε-NSGAII. *IEEE Congress on Evolutionary Computation*, Beijing, China, IEEE Press

4.1.   Introduction

The problem defined in Chapter 3 is an NP-hard problem, since the bi-level linear programming problem is already an NP-hard problem (Gao et al., 2005). Therefore, the problem is computationally too expensive to be solved exactly for larger networks: a heuristic method is needed, accepting that the outcome is an approximation of the true Pareto set. As was stated in Chapter 2, when solving multi-objective (MO) problems, the class of genetic algorithms (GAs) is often used. These algorithms have a low risk of ending up in a local minimum, do not require the calculation of a gradient and are able to produce a diverse Pareto set (Deb, 2001).

More specifically, the algorithms NSGAII (Non-dominated sorting genetic algorithm II) and ε-NSGAII are used in this research. These algorithms were developed respectively by Deb et al. (2002) and by Kollat and Reed (2006). Both algorithms optimise multiple objectives simultaneously, searching for a set of non-dominated solutions, i.e. the Pareto set. NSGAII has been successfully applied by researchers to solve MO optimisation problems in transportation engineering and proved to be efficient for this type of problems with respect to other types of GAs (Sharma et al., 2009; Wismans et al., 2011). Kollat and Reed (2006) applied ε-NSGAII in the field of water resources, for long-term groundwater monitoring design. They found that on their 4-objective test problem ε-NSGAII outperformed benchmark algorithms NSGAII, ε-MOEA (ε-Multi-Objective Evolutionary Algorithm), see Deb et al. (2005), and SPEA2 (Strength Pareto Evolutionary Algorithm 2), see Zitzler et al. (2001), in terms of attained hypervolume and distance to the real Pareto set, especially in the early stage of running the algorithm. This excellent performance in limited number of iterations is very beneficial for computationally highly expensive problems, like the optimisation problem in this research. Albarello et al. (2014) provided a practical example of application of ε-NSGAII for the optimisation of a satellite system, where it performs comparable with ε-MOEA and better than NSGAII and random search. Given these promising results, ε-NSGAII is applied for the multimodal MO-NDP (network design problem), next to NSGAII. The latter algorithm serves as a benchmark algorithm because it is well-known and it has previously been applied successfully in transportation engineering. A comparison between the two algorithms provides new information concerning the performance of both algorithms, specifically when applied in the field of multimodal passenger transportation network design: it may perform differently in this different type of problem.

However, the stochastic processes involved in this type of algorithms may be a disadvantage. Due to these stochastic processes, the results from two runs of the same algorithm (with the same parameter settings) will probably not be exactly identical, because the result is an approximation of the true Pareto set and this true Pareto set might contain far more solutions than the algorithm will be able to assess and therefore find. Furthermore, the optimal values for the genetic parameters are case-specific and the time available to determine these optimal values is limited, so the genetic parameter values may also cause variation in the results. If a Pareto set is used as input for decision support information, the level of variation in the results is important for the decision maker. Note that this variation caused by the method is one of many aspects of the modelling framework that are subject to uncertainty, for example the behavioural rules in the lower-level model (both model structure, choice parameters and

general parameters like oil prices) and the transportation demand models. The influence of a different demand input will be addressed in chapter 6.

The primary goal would be to minimise variation in the resulting objective values, since this determines the extent to which policy goals are achieved. However, variation in the decision space is also important, because in the end the values of decision variables determine the measures that can best be taken. The level of variation in decision variable values determines if proper information is available to make the right decision. Therefore, this level of variation can be seen as a performance indicator as well. The relation between objective values and decision variable values is problem specific: in a steep, linear function a different decision variable value will imply a different objective value, while in a flat function with many local minima, several combinations of decision variables may result in (almost) the same objective values, making many solutions indifferent. In the first case, choosing a different network design results in different objective values, while in the latter case apparently many good network designs are possible, given the objectives that were defined. As long as a good set of designs is produced by the algorithm, it may not a problem if a different set of network designs is produced.

Attention is given to the influence of the stochastic process that takes place in GAs and of genetic parameter settings on the optimisation outcomes. This influence is tested by comparing multiple outcomes of the optimisation process of the well-known algorithm NSGAII, which is the most popular MOEA in current use (Zavala et al., 2014).

The algorithms are described in Section 4.2. In Section 4.3, methods are defined to compare Pareto sets. This enables an objective assessment of similarities and differences between Pareto sets, both concerning objective values as well as values of decision variables. The results of the comparison between ε- NSGAII and NSGAII are described in Section 4.4. The influence of the stochastic nature of GAs on the optimisation results is analysed in Section 4.5. The chapter ends with conclusions in Section 4.6.

4.2.   Description of algorithms used

Properties of GAs that apply to both NSGAII as ε-NSGAII are described in this section first, followed by a description of the genetic representation of the solutions. After that, both algorithms are described. NSGAII is a widely applied algorithm in MO optimisation. ε-NSGAII is less known and is a further development of the NSGAII algorithm. In fact, within ε-NSGAII, NSGAII is restarted multiple times when there is only few improvement in results.

4.2.1. Genetic algorithms

GAs belong to the class of evolutionary algorithms (EAs), which belong to the larger class of evolutionary computation. A GA is a population-based algorithm, inspired on the process of natural evolution, where well-adapted individuals within a species have a bigger chance to survive (survival of the fittest). Pairs of parents are selected based on their good performance compared to others, from which new individuals (children) are created in the population by reproduction. Furthermore, properties of individuals slightly change by mutation. Over generations, this leads to improvements in the population, since individuals with desired

properties have a better chance to survive. When GA is applied to the NDP, every population member is a specific network design. Properties of network designs may be mixed to form new solutions. A population-based algorithm is very suitable for MO optimisation, since the result of solving an MO optimisation problem is a Pareto set, instead of a single solution. Variation in population members is explicitly valued in MO optimisation, just like in its original counterpart in natural evolution. This is in contrast with other well-known metaheuristics like simulated annealing and tabu search where the solution space is explored by moving from solution to solution while searching for improvement, with one solution as result (usually a local optimum). To end up in different solutions, these algorithms should be run multiple times, without guarantee that a diverse Pareto set is found.

Figure 4.1 shows the general working principle of GAs. A GA begins its search with a set of solutions (Deb, 2001). Once a random population of starting solutions is created, each solution is evaluated in the context of the underlying design problem (in this research the multimodal MO-NDP): a value for each objective function is determined. Thereafter, a metric must be defined using these objective function values (shortly: objective values) to assign a relative merit to the solution (called fitness). A termination criterion is then checked. If the termination criterion is not satisfied, the population of solutions is modified by three main genetic operators (reproduction, crossover and mutation) and a new (and hopefully better) population is created. The generation counter h is incremented to indicate that one generation of the GA is complete.

Reproduction is also called (environmental) selection and has as primary objective to make duplicates of good solutions and eliminate bad solutions in a population, while keeping the population size constant (Deb, 2001). During reproduction, good solutions are given a larger chance to mate and to give their properties to the next generation than bad solutions. The result of this step is a set of parent solutions (also called mating pool) that will produce offspring (children) by the next operator: crossover.
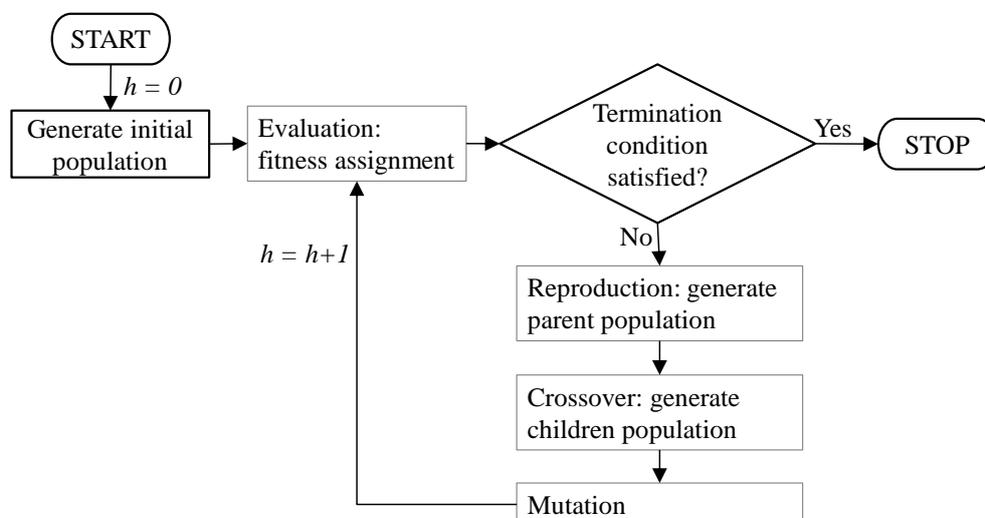
Figure 4.1: Flowchart of the working principle of a GA (Deb, 2001)

Crossover combines a good solution with another good solution to form new, hopefully even better solutions by combining good properties from both solutions. During crossover, two

parent solution strings are picked from the mating pool and some portion of the strings are exchanged between the strings to create two new children solution strings.

Finally, the mutation operator concludes the modification of the population. The bitwise mutation operator randomly changes the value of the bit with a defined mutation probability. Mutation promotes diversity in the population: it enables a wide exploration of the solution space.

4.2.2. Genetic representation

GAs like ε-NSGAII and NSGAII need a genetic representation of solutions. A binary representation is chosen, i.e. each solution is represented by a 0/1 string. In the case studies described in Chapter 3, the decision vector $\underline{y}$ contains both variables that are already binary as variables that can take 4 different discrete values (these variables represent the frequency of PT lines). The binary variables are directly put into the genetic string. The other variables are converted into two bits, to be able to represent the 4 different discrete values. This representation implies that the two-bit variables have a twice as large chance to be hit by the mutation operator than variables that are naturally binary. Another implication is that the first bit represents a larger difference in the value of the discrete variable (i.e. the frequency of the PT line) than the second bit. Although the effect of this representation on the optimisation process is not known, it is expected that this additional emphasis on frequency setting of PT lines is beneficial for the results: the PT lines that have frequency as decision variable are major lines, for which changing the frequency is expected to have a larger effect on the values of objective functions than changing the other binary decision variables (see chapter 5 for case-specific results).

The binary string is denoted by $\underline{y}'$. An invertible function exists between the two representations: $\underline{y} = f(\underline{y}')$ and $\underline{y}' = f^{-1}(\underline{y})$. In the GA the population consists of solutions coded in the genetic representation $\underline{y}'$, on which the genetic operators crossover and mutation are applied. $\underline{y}'$ is converted to decision vector $\underline{y}$ by function f to enable function evaluation by the lower-level model.

4.2.3. NSGAII

The steps of NSGAII are described in Algorithm 4.1. Step 1 is the initialisation: a random starting population is generated, using the parameter for population size.

Step 2 is the fitness assignment. Within NSGAII, the fitness value is calculated using a combination of non-dominated sorting and crowding distance. During non-dominated sorting, the solutions are ranked based on Pareto dominance. All solutions in the Pareto set receive rank 1. After that, these solutions are extracted from the set and all Pareto solutions in the remaining set receive rank 2, etc. Crowding distance is used to sort the solutions within these ranks, which is a measure for the location of a solution within objective space compared to the other surrounding solutions. The calculation is based on an average normalised distance to the closest solution per dimension (i.e. objective). Crowding distance calculation requires sorting of the population according to each objective value. The minimum values for each objective are assigned an infinite value, assuring that these values survive. All intermediate solutions are assigned a value equal to the difference in the normalised function values of two adjacent

solutions. The crowding distance value (and thus the fitness value) is higher if a solution is more isolated, promoting a more diverse Pareto set. To summarise, non-dominated sorting is the most important sorting criterion: crowding distance comes into play when solutions have the same rank in non-dominated sorting.

After checking the stop criterion (step 3), in step 4 the best ν_a solutions are selected from the union of the archive (containing the best solutions from the previous generation) and the offspring based on the fitness value. From these ν_a solutions, parents are selected to mate by binary tournament selection with replacement. This means that pairs of solutions are randomly selected, and the best of these two solutions (in terms of fitness value) is selected as a parent. Parents selected for the current tournament are again candidates for further tournaments. This implies that some solutions are selected multiple times, while other solutions are not selected as a parent at all. A recombination rate of 1 is applied, which means all selected parents are recombined.

```
Step 1: Initialisation
Set population size νp, which is equal to the archive size νa, the
maximum number of generations H
Generate an initial population φ0
Set h = 0 and π0 = ∅

Step 2: Fitness assignment
Combine archive πh and offspring φh, forming ϑh = πh ∪ φh
Calculate fitness values of solutions by dominance ranking and
crowding distance sorting

Step 3: Termination
If h ≥ H
      Determine Pareto set P from the set of all evaluated
      solutions φ = {φ0 ∪ ··· ∪ φH} and STOP
Else
      Go to step 4
End

Step 4: Reproduction / environmental selection
Determine new archive πh+1 by selecting the νa best solutions out of
ϑh based on their fitness
Perform binary tournament selection with replacement on πh+1 to
determine the mating pool of parents σh+1

Step 5: Variation
Apply recombination (uniform crossover) to the mating pool σh+1 to
create offspring φh+1
Apply mutation (assign a random value to decision variables with
probability φ) to offspring φh+1
Set h = h+1 and go to step 2
```

Algorithm 4.1: NSGAII

In step 5 the next generation of offspring is created by uniform crossover: two offspring solutions are constructed by choosing every bit from one of both parents with probability 0.5. Consequently, every bit of both parents will be passed on to one of the two children. In

addition to this mating process, in every generation a stochastic mutation operator is applied to promote the exploration of unexplored regions in the decision space. This mutation operator assigns a random value to each 0/1 variable in $\underline{y}'$ with probability φ. Because of this, there is a probability of $\frac{1}{2}$ that the newly assigned value is equal to the old value, implying that the probability the mutation operator actually changes the variable is equal to $\varphi / 2$. In case an infeasible solution is generated during either recombination or mutation, that solution is discarded and a new solution is generated by repeating the recombination or mutation operator until a feasible solution is generated.

Since NSGAII is a heuristic method, the result of a run of NSGAII is an approximation of the true Pareto set. This result may differ among optimisation runs (depending on the random seed that was used), because the genetic operators are stochastic processes. More details of the algorithm can be found in Deb et al. (2002).

4.2.4. ε-NSGAII

ε-NSGAII has NSGAII as a basis, but adds two main elements: ε-dominance and restarts with adaptive population sizing (see Figure 4.2).
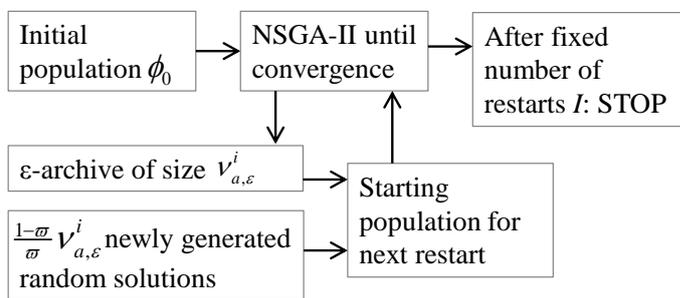
Figure 4.2: The main loop of the ε-NSGAII algorithm

For every objective z, a value ε_z is set, resulting in an ε-grid in the objective space. In this way, every Pareto solution is placed within one ε-box (see Figure 4.3). This allows the user to specify the precision of the algorithm for each objective. The concept of ε-dominance is defined in two steps (step A and step B).

In step A only one Pareto solution is chosen to represent an ε-box. If that box contains more than one Pareto solution, the solution that is closest to the vertex of the hyperbox with minimum objective value for all objectives is chosen. When determining this distance, the Euclidian distance is calculated, using the specified ε-values to normalise the objective values.

In step B boxes that are dominated by another box (i.e., have a worse or equal objective value for every objective) are eliminated. The objective value for a box is chosen to be the minimum value of all objectives in the box (similar to the vertex of the hyperbox used to find the best Pareto solution per ε-box in step 1). The solutions that are contained by every remaining ε-box are denoted as the ε-Pareto set.
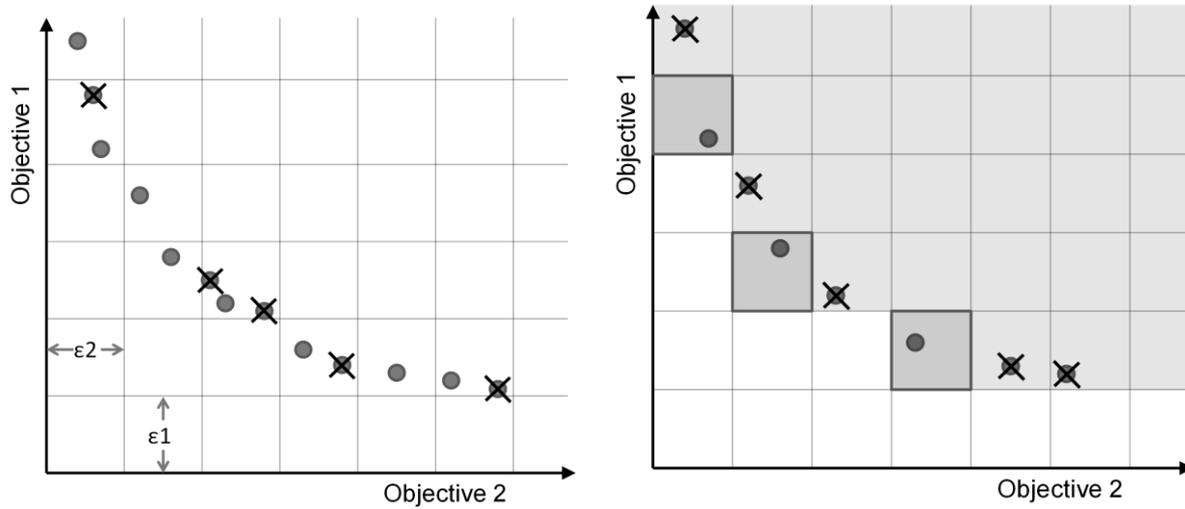
Figure 4.3: Illustration of ε-dominance for the 2 objective case. Left: step A, choosing 1 solution per ε-box. Right: step B, applying dominance to the ε-boxes: the light grey boxes are dominated by the darker boxes

In Algorithm 4.2 ε-NSGAII is described in steps. Step 1 is the initialisation: a random starting population is generated, using the parameter for initial population size.

In step 2, the NSGAII algorithm is run multiple times. During every generation of the NSGAII algorithm in which the Pareto set is determined by non-domination sorting, the ε-Pareto set is also determined and saved in an ε-archive π_h^ε. This ε-archive is updated every generation by applying ε-dominance over the union of the ε-archive π_h^ε of generation h and the new archive π_{h+1} of Pareto solutions in NSGAII in generation h. When no progress is made any more in the ε-archive for a specified number of generations ς the NSGAII algorithm is terminated.

After checking the termination criterion in step 3, a restart is activated in step 4, where the starting population after the restart consists of the ε-archive, replenished with new randomly generated solutions (the number of solutions to add depends on injection scheme parameter ϖ). The generation size after the restart depends on the size of the ε-archive: the larger the archive, the larger the generation size. The ε_z values remain constant over the total run, i.e. do not change during a restart. The setting of the ε_z values is very important, since these values determine the trade-off between precision and speed of the algorithm.

```
Step 1: Initialisation
Set the number of restarts I, the initial population size v_p^0, the
injection scheme parameter ϖ, the number of generations without
change to trigger a restart ς.
Randomly generate an initial population φ_0
Set h=0, i=0 and π_0=∅

Step 2: NSGAII
Run the NSGAII algorithm (step 2 to 6 in Algorithm 4.1) with
generation size v_p^i and initial population φ_h. As termination
criterion use: no ε-progress is made for ς generations (i.e. π_h^ε has
not changed for ς generations). Note that h increases with every
generation within NSGAII

Step 3: Termination
If i≥I
      Determine Pareto set P from the set of all evaluated
      solutions φ={φ_0∪φ_1∪⋯∪φ_h∪⋯} and STOP
Else
      Go to step 4
End

Step 4: Restart
Set v_p^i=v_{a,ε}^i/ϖ, randomly generate (1−ϖ)v_p^i new solutions as starting
solutions for generation h+1 and put them in φ_{h+1}, add ε-archive to
these solutions: φ_{h+1}=π_h^ε∪φ_{h+1}
Set i=i+1 and h=h+1
Go to step 2
```

Algorithm 4.2: ε-NSGAII

4.3. Methods to compare Pareto sets

When comparing Pareto sets the quality of and similarity between the Pareto sets can be objectively assessed by indicators. These indicators are usually based on the objective values of the solutions in the Pareto set. They either measure the extent to which the set is close to the real Pareto set (or if the real Pareto set is unknown, the best-known Pareto set) or the extent to which the set is diverse (Deb, 2001).

However, another aspect that is relevant when comparing outcomes resulting from different MO optimisation processes, is the extent to which two Pareto sets are similar in the decision space. Since it is very unlikely that different runs of GAs (with different random seeds) produce exactly the same results, it is desirable that these results are as similar as possible: if a certain algorithm has a small variation among optimisation results, this algorithm can be seen as high performing. In this section comparison indicators are defined, enabling a formal assessment of this similarity for the outcomes of the case study experiments.

Firstly, the notation is formalised for optimisation results from various runs of the optimisation process. After that, a selection of indicators from literature is presented for comparison in the objective space and new indicators are defined for comparison in the decision space.

4.3.1. Output definitions

The set ϕ^j is defined as all decision vectors (or solutions) that are evaluated during one optimisation process j. The set of N^j solutions $P^j = \left\{ \underline{y}_1^j, \underline{y}_2^j, \cdots, \underline{y}_{N^j}^j \right\}$ is defined as the Pareto set resulting from process j, which includes all non-dominated solutions with respect to all solutions in ϕ^j, or mathematically, there is no $\underline{y}_i^j \in P^j \mid \exists \underline{y}_{i'}^j \in \phi^j : \underline{y}_{i'}^j \preceq \underline{y}_i^j$. As defined in Chapter 3 every element or solution $\underline{y}_i^j$ (solution i in Pareto set j) represents one transportation network design, which is defined by the values of V decision variables in decision vector $\underline{y}_i^j = \left[y_{i1}^j, y_{i2}^j, \cdots, y_{iv}^j, \cdots, y_{iV}^j \right]$. For every element $\underline{y}_i^j$ a vector of objective functions $\underline{Z}(\underline{y}_i^j)$ is evaluated, which consists of W objective functions, denoted by index w: $\underline{Z}(\underline{y}_i^j) = \left[Z_1(\underline{y}_i^j), \cdots Z_w(\underline{y}_i^j), \cdots, Z_W(\underline{y}_i^j) \right]$.

P^j is the j-th outcome of the optimisation problem defined in Chapter 3: one approximation of the Pareto-optimal set. The j-th run of the algorithm is also denoted as run j. Pareto set P^Σ with respect to all evaluated solutions $\phi^\Sigma = \{\phi^1 \cup \cdots \cup \phi^J\}$ is denoted as the superset and is the best known Pareto set based on J runs.

As a result of this definition, the Pareto sets used for analysis are based on all solutions that are evaluated during one optimisation process (see also step 3 of Algorithm 4.2). An alternative would be to analyse the Pareto set resulting from the final generation of the GA. Determining the Pareto set from a set of known solutions takes considerable computation time when the set of known solutions grows larger. Therefore, a trade-off exists between computation time to determine the Pareto set from all known solutions and to calculate the values of the performance indicators and computation time to evaluate the objective values of a solution using the lower-level model. For the case studies the lower-level model is expensive in terms of computation time, so the number of evaluated solutions in one run is never too large to be infeasible any more to determine the Pareto set from all known solutions. Therefore, the trade-off is clearly in favour of determining the Pareto set and calculating the indicator values after optimisation. Furthermore, the generation sizes (and thus size of the Pareto sets resulting from the final generations) in NSGAII and ε-NSGAII are considerably different, making comparing these sets more difficult.

As a consequence of taking all evaluated solutions within one run into account to determine the Pareto set for that run, the indicators based on objective values (that are used to assess the algorithms) focus on attainment of objective values rather than on diversity / spread.

4.3.2. Performance indicators

A distinction between indicators is made based on two main characteristics of the methods.

- The first is a distinction between the number of Pareto sets to be compared (Zitzler et al., 2008):
    - A unary indicator is defined for one Pareto set on its own.
    - A binary indicator is defined to compare two Pareto sets.
    - A distribution-based indicator is defined to compare J Pareto sets.
- The second distinction is whether the indicator compares the Pareto sets in the objective space or in the decision space:

      o   Indicators based on the objective values $\underline{Z}$ of the solutions in the set(s).
      o   Indicators based on the decision values $\underline{y}$ of the solutions in the set(s).

An overview of the indicators used is given in Table 4.1. Some of these indicators use a distance function, which can be any distance function defined for two solutions as $d(\underline{y}_i^j, \underline{y}_{i'}^{j'})$. For the analyses the distance function is defined as in Eq. 4.12. This formulation can be interpreted as the number of decision variables that have a different value for binary variables and the extent to which decision variables are different for discrete variables.

$$d(\underline{y}_i^j, \underline{y}_{i'}^{j'}) = \sum_{v=1}^{V} \left| y_{iv}^j - y_{i'v}^{j'} \right|. \tag{4.12}$$

Table 4.1: Overview of indicators to compare Pareto sets

Indicator	Type	Description
N^j (4.1)	Unary	Cardinality of Pareto set j. The bigger a Pareto set, the more complex it is to analyse for a decision maker. Furthermore, with more solutions in the set it is easier to have a higher score on other criteria, like hypervolume.
$MIN_w(P^j) = \min\limits_{\underline{y}^j \in P^j} Z_w(\underline{y}^j), \forall w$ (4.2) $MAX_w(P^j) = \max\limits_{\underline{y}^j \in P^j} Z_w(\underline{y}^j), \forall w$ (4.3) $RMI_w(P^j) = \dfrac{MIN_w(P^j) - MIN_w(P^\Sigma)}{MAX_w(P^\Sigma) - MIN_w(P^\Sigma)}, \forall w$ (4.4)	Unary, objective space	The minimum objective value attained by Pareto set j for every objective function Z_w, with respect to the range covered by the superset. For that, the minimum and maximum values over all Pareto solutions in the superset are determined, as well as the minimum value over all solutions in Pareto set j. This indicates the extent to which set j covers the entire Pareto front for every objective dimension.
$SSC(P^k)$ **Figure 4.4: Hypervolume 2-D visualisation**	Unary, objective space	Hypervolume, implemented as in While et al. (2006), also known as S-metric or space coverage. In the 2-dimensional case it determines the area that is covered by the Pareto set with respect to a reference point (the star in Figure 4.4). The reference point represents the upper boundary of all objectives: the reference point is defined in a way that it is dominated by all solutions in the Pareto set. Because the true maximum values of the objective functions are not known a conservative point is chosen, based on the evaluated solutions. In the 3 dimensional case area is replaced with volume, and in the more dimensional case with hypervolume.
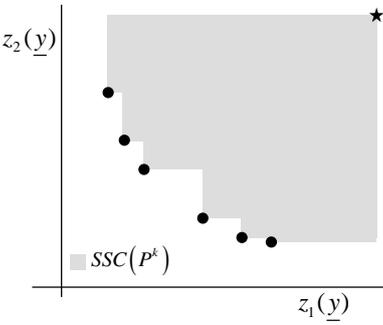
$CTS(P^j, P^{j'}) =$ $$\frac{\left\| \left\{ \underline{y}^{j'} \in P^{j'} \mid \exists \underline{y}^j \in P^j : \underline{y}^j \preceq \underline{y}^{j'} \right\} \right\|}{N^{j'}}$$ (4.5)	Binary, objective space	The set coverage or C-metric, see Zitzler and Thiele (1999). The level in which the solutions in $P^{j'}$ are weakly dominated by at least one solution in set P^j so a higher value indicates a better set coverage of P^j over $P^{j'}$.
$$AD(P^j, P^{j'}) = \frac{1}{N_j} \frac{1}{N_{j'}} \sum_{i=1}^{N_j} \sum_{i'=1}^{N_{j'}} d(\underline{y}_i^j, \underline{y}_{i'}^{j'})$$ (4.6)	Binary, decision space	The average distance (in the decision space) between the solutions in two Pareto sets P^j and $P^{j'}$ (where j may also be equal to j'). This indicator is used to check whether the variation within a Pareto set is different from the variation between Pareto sets.
$$AND(P^j, P^{j'}) = \frac{1}{N^j} \sum_{i=1}^{N^j} \min_{i' \in \{1..N^{j'}\}} d(\underline{y}_i^j, \underline{y}_{i'}^{j'})$$ (4.7)	Binary, decision space	The distance of an element in P^j to the closest element (i.e. most similar in decision space) in Pareto set $P^{j'}$, averaged over all elements P^j. A low value indicates higher similarity (the extreme case $AND(P^j, P^j) = 0$ holds).
$$AFD(P^j, P^{j'}) = \frac{1}{V} \sum_{i=1}^{V} \left| \rho_v^j - \rho_v^{j'} \right|,$$ (4.8) with $$\rho_v^j = \frac{1}{N^j} \sum_{i=1}^{N^j} \kappa(y_{iv}^j)$$ (4.9) with $$\kappa(y_{iv}^j) = \begin{cases} 0 & \text{if } y_{iv}^j = 0 \\ 1 & \text{otherwise} \end{cases}$$ (4.10)	Binary, decision space	Comparison of fractions of nonzero decision variables. The fraction ρ_v^j of solutions in Pareto set j that have a positive value for decision variable y_v characterises the decision space of the set: decision variable y_v represents the existence of a measure in the transportation network, so a higher fraction for variable y_v implies a larger ability to attain Pareto solutions when that variable has a positive value. The κ function defines the nonzero relation. To indicate the difference between Pareto set j and j', the differences between these fractions are calculated, averaged over all V decision variables.
$$\alpha(\underline{Z}) = \frac{1}{J} \sum_{j=1}^{J} \Gamma(\exists \underline{y}^j \in P^j : \underline{y}^j \preceq \underline{Z})$$ (4.11)	Compare J Pareto sets, objective space	Attainment function, as defined in Zitzler et al. (2008). For each objective vector $\underline{Z}$ this function returns the probability that $\underline{Z}$ is attained by a fraction α of J Pareto sets. It is approximated on the basis of the approximation set sample. Function Γ returns 1 if its argument is true and 0 if it is false.

4.4. Performance comparison between NSGAII and ε-NSGAII

In this section the results of experiments that compare the performance of NSGAII and ε-NSGAII are presented. First, the experiments are described. Then, the results are presented for various well-known, objective value-based performance indicators from literature, as defined in Section 4.3.2.

4.4.1. Experimental set-up

The optimisation algorithm is run four times: two runs using NSGAII and two runs using ε-NSGAII. Differences between the outcomes of each pair of runs occur, because both runs have a different random seed in the Monte Carlo simulation that is used for the recombination and mutation operators in the GAs. Case study 2 is used, because it has more decision variables than case study 1 and is therefore a more complex problem, which needs fast progress in a limited number of function evaluations. Another advantage is that it has a more detailed network in the lower-level model and therefore evaluates objective values more accurately. The objective function evaluation for one solution in case study 2 takes approximately 6.5 minutes. Two weeks is seen as a computation time that is still reasonable, so the number of function evaluations should be kept below 3000.

For ε-NSGAII, the parameters from Kollat and Reed (2006) could not be used, because the number of function evaluations would be much larger than 3000. A test problem is used to come to appropriate settings (see Equations 4.13-4.16). This test problem is comparable with case study 2 (equal number of objectives, objectives that are partly opposed, equal decision space and multiple local minima), but has small computation time. The function can be evaluated for all values for the number of decision variables V.

$$Z_1(\underline{y}') = \sum_{v'=1}^{V'} y'_{v'} + \sin(\sum_{v'=1}^{V'} y'_{v'}) . \tag{4.13}$$

$$Z_2(\underline{y}') = \lceil V'/2 \rceil + \sum_{v'=1}^{\lfloor V'/2 \rfloor - 1} y'_{v'} - \sum_{v'=\lceil V'/2 \rceil}^{V'} y'_{v'} + \sin(\sum_{v'=\lceil V'/2 \rceil}^{V'} y'_{v'}) . \tag{4.14}$$

$$Z_3(\underline{y}') = V' - |y'_{v'} : \mathrm{mod}(v',2) = y'_{v'}| . \tag{4.15}$$

$$Z_4(\underline{y}') = \sum_{v':y'_{v'}=0} (v' - \lfloor V'/2 \rfloor)^2 . \tag{4.16}$$

The settings in Table 4.2 resulted in a number of function evaluations just below 3000, using an initial population size v_p^0 of 10 (just like in Kollat and Reed (2006)). A 1/3 injection scheme is used ($\varpi = 1/3$), implying that at every restart, 1/3 of the new starting population will consist of the ε-archive and 2/3 of new randomly generated solutions. Next, the parameter for convergence that triggers a restart ς was set in a way that after 3 generations without a change in the ε-archive, a restart is activated.

Table 4.2. Parameter values of ε-NSGAII

Parameter	*Parameter value*
Initial population size v_p^0	10
Mutation rate φ	0.03
Injection scheme parameter ϖ	1/3
Number of generations ς without change to trigger a restart	3
Epsilon values ε_z	5 to 8 ε-intervals per objective

The epsilon values ε_z are set for each objective in a way that the range known for each objective (bounded by the minimum and maximum known values) is divided into 5 to 8 ε-intervals. In the situation of a 4-objective problem this results in approximately 2000 ε-boxes in the objective space. However, since the Pareto set contains a certain structure by definition (i.e. in the 2-dimensional case the Pareto set is a line and in the multidimensional case a hyperplane), a majority of the boxes is empty. Step B in Figure 4.3 further reduces the ε-Pareto set, resulting in a typical size of the ε-archive of approximately 20 solutions in case study 2.

With these parameter settings for ε-NSGAII, run 1 results in 2631 evaluated solutions and run 2 in 2330 evaluated solutions, both after 5 restarts. This difference occurred because the 5[th] restart is activated at a different moment in time for the two runs with different random seeds. For NSGAII, the number of generations is fixed, so the number of function evaluations is comparable, but never shorter than the computation time of ε-NSGAII. Using a generation size of 80, this results in 33 generations to come to 2640 evaluated solutions.

4.4.2. Minimum per objective

In Figure 4.5 the values per objective function $MIN_w(P^j)$ that are achieved in each run are shown (normalised by the range known for each objective from all four Pareto sets and rescaled such that NSGAII run 1 has index 100). Both algorithms perform comparably: on average NSGAII performs better for two objectives and for the other two objectives ε-NSGAII performs better, although the differences are small. However, NSGAII shows larger variation in results, which may indicate a larger influence of the stochastic processes in the GA. Note that the absolute difference per objective varies depending on the specific case studied: when the measures have a strong influence on an objective, a larger absolute reduction for that objective is achieved than when an objective can hardly be influenced by the decision variable values.
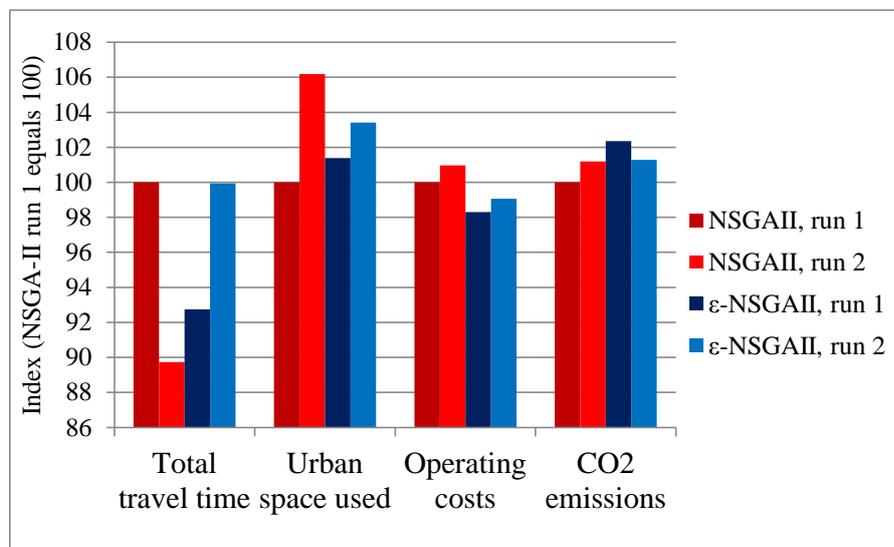
Figure 4.5. Normalised values per objective for each of the 4 runs

A relatively large variation can be observed for total travel time. This may be explained by the network dynamics in the transportation NDP, which cause travel time to have a less clear relation with the decision variables than the other three objectives. For example, a measure may cause a traveller to choose a different route, having influence on the travel time of many other travellers, but having less influence on CO_2 emissions and urban space used, because the travellers still use the same mode to reach the destination.

4.4.3. Hypervolume

The next indicator that is considered is the hypervolume indicator $SSC(P^k)$ (While et al., 2006). In the chosen definition (using a reference point) a larger hypervolume means a better solution. In Table 4.3 the normalised hypervolume values for the Pareto sets resulting from the 4 optimisation runs are shown. On average ε-NSGAII outperforms NSGAII, although NSGAII run 1 has a slightly better performance than ε-NSGAII run 2. The better average performance of ε-NSGAII is not due to more solutions in the final Pareto set: both algorithms produce a comparable number of Pareto solutions.

Table 4.3: The normalised hypervolume values for four optimisation runs

	Number of Pareto solutions N^j	Normalised hypervolume
NSGAII, run 1	241	0.990
NSGAII, run 2	202	0.987
ε-NSGAII, run 1	210	1.000
ε-NSGAII, run 2	211	0.988

4.4.4. Convergence

In Figure 4.6 the hypervolume is shown as a function of the number of evaluated solutions until that moment in time. The hypervolume is calculated based on the Pareto set that is known with respect to all solutions that had been evaluated until that moment in time. In Figure 4.7 the size of the corresponding Pareto sets is plotted. Note that the hypervolume value never decreases, since each new solution either improves the Pareto set or the Pareto set remains the same. It can be observed that ε-NSGAII converges faster than NSGAII in terms of hypervolume. The number of found Pareto solutions is larger for ε-NSGAII in the early stage of running the algorithm, while the number of Pareto solutions for NSGAII is larger in the later stage. These results show that ε-NSGAII may be stopped after 1800 solutions, almost without losing any hypervolume.
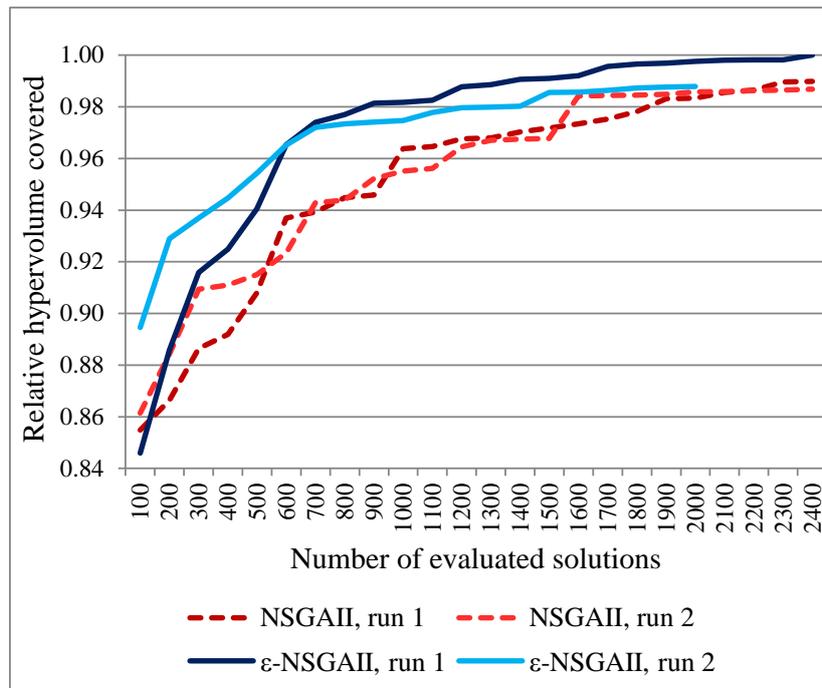
Figure 4.6: Hypervolume covered after increasing number of evaluated solutions
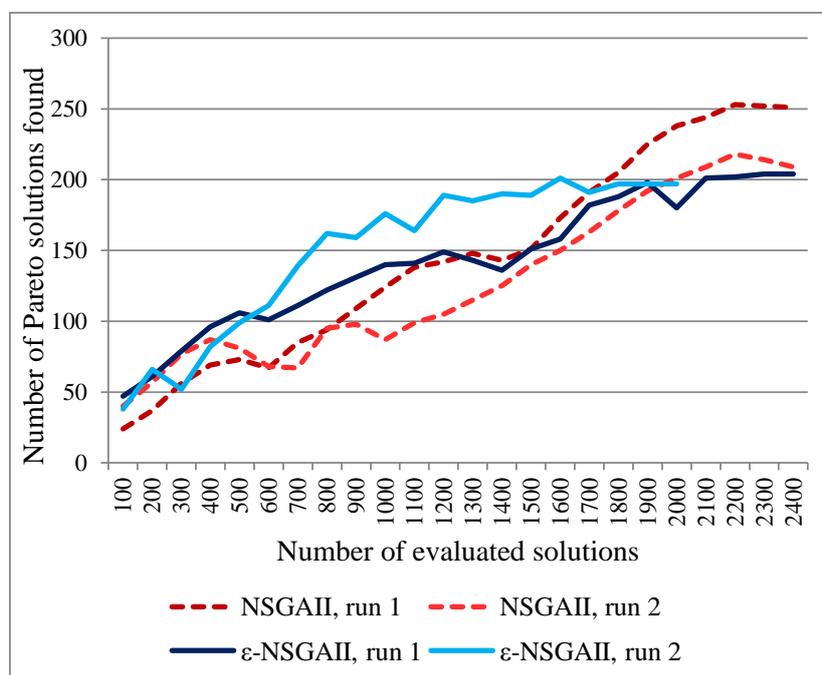
Figure 4.7: The size of the found Pareto set with respect to the number of evaluated solutions

4.4.5. Set coverage

Finally the Pareto sets are compared using the set coverage indicator $CTS(P^{j}, P^{j'})$ (Zitzler and Thiele, 1999). The set coverage is a pairwise indicator that shows the fraction of solutions in another set $P^{j'}$ that is dominated by at least one solution in a set P^{j}. This definition

implies that a higher value for a Pareto set P^j indicates a better score over the set $P^{j'}$ that it is compared with.

In Table 4.4 the values for set coverage are shown for all pairs to be formed out of the 4 optimisation runs. For example, a fraction 0.20 of Pareto solutions produced by NSGAII, run 1 dominates at least one Pareto solution produced by ε-NSGAII, run 1. The other way around, a fraction 0.62 of the solutions from ε-NSGAII, run 1 dominates at least one solution from NSGAII, run 1. It can be noted that comparing Pareto sets produced by ε-NSGAII with Pareto sets produced by NSGAII results in considerably higher (and therefore better) values than the other way around (an average value of 0.54 in the four lower left cells compared to an average value of 0.18 in the four upper right cells). Note that the set coverage of run 1 and run 2 with the same algorithm shows the influence of the random seeds: also the different scores for set coverage reveal that the results of both runs do not have an equal quality.

Table 4.4: Coverage $CTS(P^j, P^{j'})$, shown for each pair of runs

P^j \ $P^{j'}$	NSGAII		ε-NSGAII	
	run 1	run 2	run 1	run 2
NSGAII, run 1		0.26	0.20	0.17
NSGAII, run 2	0.38		0.14	0.22
ε-NSGAII, run 1	0.62	0.57		0.45
ε-NSGAII, run 2	0.46	0.50	0.20	

4.4.6. Conclusions

Putting these results together, ε-NSGAII shows better results for the experiments than NSGAII with respect to the indicators set coverage and hypervolume. For 2 out of 4 objectives, the best solutions for each individual objective are found by ε-NSGAII and for the other 2 by NSGAII, therefore both algorithms have a similar score on this third indicator. The high performance of ε-NSGAII may be explained by the focus on large gains by the ε-dominance relation: if no large progress is made any more, no computation time is wasted to achieve little improvements. Instead, a restart is triggered, stimulating exploration of new areas of the Pareto front, but using the properties of the high quality solutions found earlier (the ε-archive). Another explanation may be the dynamic population size: by starting with a small population, the algorithm focuses on dominance rather than on producing a diverse Pareto set. In limited computation time large progress is made in an early stage of the algorithm. Later, as the population size grows, it allows for more diverse solutions. The population size is directly dependent on the size of the ε-archive, so no computation time is wasted on too much detail.

The results also confirm that different runs of the same algorithm may produce results that show quite some differences. Therefore, the next section includes analyses concerning differences between runs.

4.5. Influence of stochastic principles and parameter settings in GAs

Firstly it is important that an algorithm produces high quality solutions in a limitedly available computation time, as was investigated in the previous section. Secondly, it is desirable that the algorithms are not too sensitive to the stochastic principles and parameter settings in GAs. This section investigates the influence of stochastic processes that are inherent in GAs on the optimisation outcomes. Furthermore, it investigates the sensitivity of the optimisation outcomes to the genetic parameter settings.

The influence of stochastic processes is relevant, because when applied in practice, there is no time available to do multiple runs of the algorithm. Therefore, it is desirable that the outcomes of multiple optimisation runs are similar, both in terms of objective values attained and in terms of decision variable values of solutions. For this purpose mainly binary (pairwise) indicators (defined in section 4.3.2) are used.

The sensitivity to genetic parameter settings is relevant, because when applied in practice, there is no time available to test several parameter settings to find the best settings for a specific problem. For the method to be useful in practice, the optimisation outcomes (both in terms of objective values attained and in terms of decision variable values of solutions) should therefore not be too sensitive for the genetic parameter settings.

This is done for the well-known GA NSGAII, which is the most popular MOEAs in current use (Zavala et al., 2014). It is assumed that the findings in this section for NSGAII are approximately valid for ε-NSGAII, since ε-NSGAII has NSGAII as a basis.

First, the experiments are described. After that, the results are presented: first for unary indicators, then for binary indicators and finally for the attainment function that compares multiple Pareto sets.

4.5.1. Experimental set-up

The NSGAII algorithm is run 13 times in total (see Table 4.5) using case study 1. Case study 1 is chosen, because it has smaller computation time, enabling a larger number of experiments. In 5 runs with identical parameter settings the impact of the stochastic principles in the algorithm is investigated. In two more groups of 3 runs, 2 different combinations of parameter settings are chosen to investigate the influence of generation size. Finally, in 2 runs the influence of a different mutation rate is investigated. A set of runs with the same parameter settings is further denoted as a group of runs. The groups are highlighted in Table 4.5.

All parameter settings are chosen in such a way that computation times are still reasonable (one week per run, without distributed computing), while convergence is almost reached in terms of attained hypervolume. Since the evaluation of one solution takes approximately 3.25 minutes and since the evaluation of solutions takes most computation time by far, the maximum number of evaluated solutions is around 3000 per run. For a fair comparison, the total computation time (indicated by the total number of evaluated solutions $v_a H$) is approximately the same for all runs, implying that a lower value for v_a is combined with a higher value for H and vice versa.

4.5.2. Unary performance indicators

All 13 runs attain considerably larger hypervolumes (see Table 4.5) than the random starting populations (the average value of hypervolume covered by the starting populations ϕ_0^1 to ϕ_0^J is $7.31*10^{21}$). Furthermore, the individual runs are not far from the best known Pareto set, the superset P^Σ, taking into account that the superset contains a larger number of solutions than the individual runs.

Table 4.5: Parameter settings and unary indicators of the 13 runs of the optimisation process

Set number j	archive size v_a	Number of generations H	Mutation rate φ	Number of solutions in set N^j	Hypervolume $SSC(P^j)$ $(*10^{21})$
1	80	36	0.03	339	9.57
2	80	36	0.03	282	9.78
3	80	36	0.03	301	9.84
4	80	36	0.03	364	9.76
5	80	36	0.03	291	9.82
6	50	60	0.03	287	9.80
7	50	60	0.03	329	9.88
8	50	60	0.03	340	9.76
9	100	30	0.03	363	9.86
10	100	30	0.03	255	9.79
11	100	30	0.03	312	9.85
12	80	36	0.1	159	9.28
13	80	36	0.01	337	9.80
Σ	-	-	-	1019	10.12

When the numbers of attained Pareto solutions are compared, only run 12 shows a large difference with the other runs. The variation within the number of Pareto solutions is similar within groups of runs and between groups of runs. The same holds for the hypervolume that is covered by the Pareto sets. In addition, run 1 has a relatively low value for hypervolume, which seems to be an outlier resulting from the stochastic optimisation process.

Next, the minimum values attained by the different runs $RMI_w(P^j)$ for all 4 objective functions are compared. Some observations from Figure 4.8 are:

- Travel time has the most (relative) variation from all objectives (a similar result was found earlier in section 4.4.2).
- Run 12 does not provide (near) optimal solutions for any objective.
- Run 7 performs excellently, achieving 3 minimal values (i.e. contains the best scoring solutions for 3 individual objectives) and 1 average value.
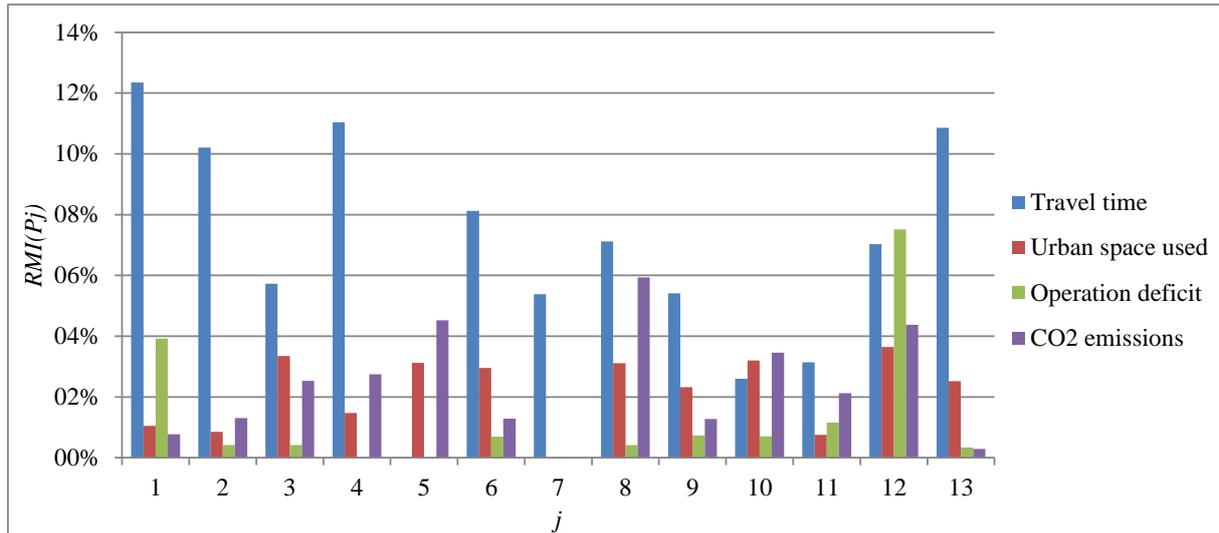
Figure 4.8: The relative values $RMI_w(P^j)$**, for 13 runs and for 4 objectives**

Although no hard conclusions can be drawn based on one optimisation outcome, the unary indicators show that P^{12} does not perform as well as the other Pareto sets. This is also shown by additional information on the set coverage concerning P^{12}: the average value over $j, j \neq 12$ of $CTS(P^j, P^{12})$ is 0.94, while the average value over $j, j \neq 12$ of $CTS(P^{12}, P^j)$ is 0.43. This indicates that a mutation rate of 0.1 (as was applied in run 12) is too high for this type of problem. Apparently this is too much randomness for this problem: good properties of solutions are no longer well-preserved by the GA, coming too close to random search.

4.5.3. Comparing pairs of Pareto sets

In this section a pairwise comparison between Pareto sets is made. Firstly, the Pareto sets resulting from runs 2 and 3 (that have comparable hypervolume values) are plotted as an example in Figure 4.9. This shows that the sets have comparable shapes, but the extremes (i.e. minimum values per individual objective) are different, and one Pareto set contains more solutions in certain areas than the other set. In the remainder of this section the indicators defined in Section 4.3.2 are used to compare the Pareto sets.

Table 4.6 shows the values of $AD(P^j, P^{j'})$ for each pair of Pareto sets. The diagonal of the table contains the values for the average distance within the same Pareto set. When looking at the table, the only clear difference that can be observed is between run 12 and the other runs. As was already seen in Table 4.5, this run is also different from the other runs for the unary indicators. It has a smaller number of Pareto solutions, which may explain the larger average distance to other solutions in the set. More generally speaking, there is a small difference between cases where $j = j'$ and cases where $j \neq j'$: overall, these average values are 6.5 and 6.8, respectively. Furthermore, there is a small difference between the average value within groups of runs (6.7) and between groups of runs (6.9). These differences seem too small to be substantial, although this result seems to confirm the hypothesis that a pair of Pareto sets with the same parameter settings has more similarities than a pair of Pareto sets with different parameter settings.
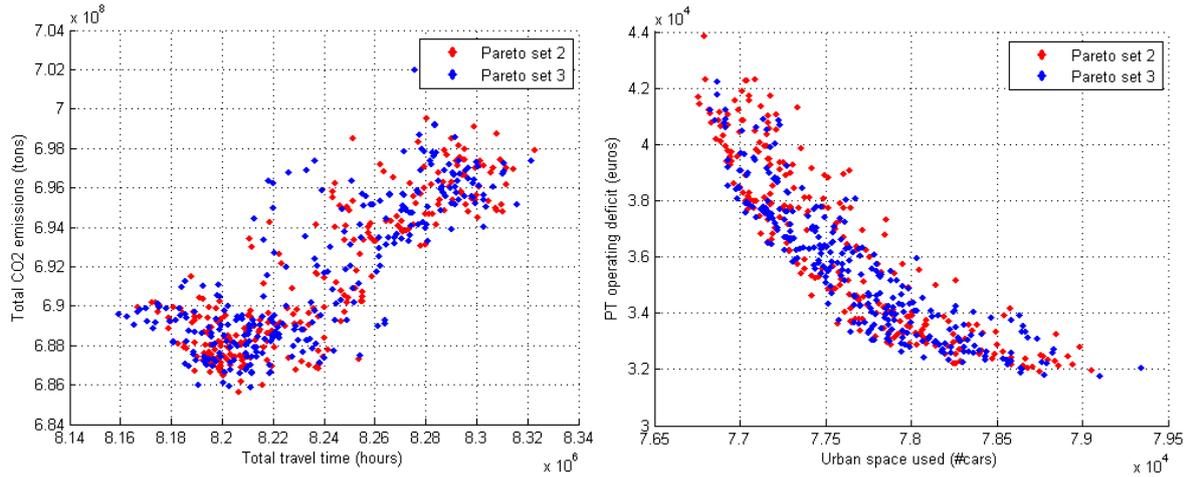
Figure 4.9: 2-dimensional plots of two different outcomes of the optimisation process, for objectives total travel time and CO2 emissions (left) and for objectives urban space used by parking and PT operating deficit (right)

Table 4.6: The average distance between Pareto solutions $AD(P^j, P^{j'})$

	1	2	3	4	5	6	7	8	9	10	11	12	13
1	6.4	6.9	6.9	6.9	7.1	6.7	7.2	6.9	7.1	7.0	7.0	7.6	7.0
2		6.5	6.5	6.6	6.7	6.6	7.0	6.8	6.7	6.8	6.7	7.5	6.8
3			5.8	6.4	6.3	6.4	6.6	6.6	6.4	6.2	6.4	7.3	6.4
4				6.4	6.7	6.7	6.9	6.9	6.8	6.7	6.7	7.6	6.8
5					6.4	6.6	6.8	6.7	6.5	6.7	6.6	7.6	6.8
6						6.1	6.7	6.5	6.6	6.6	6.7	7.4	6.6
7							6.6	6.9	6.8	6.9	6.8	7.7	7.0
8								6.4	6.6	6.8	6.7	7.5	6.8
9									6.2	6.7	6.5	7.4	6.8
10										6.5	6.7	7.4	6.7
11											6.5	7.6	6.8
12												7.8	7.5
13													6.6

Table 4.7: The average distance to the nearest Pareto solution $AND(P^j, P^{j'})$

	1	2	3	4	5	6	7	8	9	10	11	12	13
1	0.0	2.3	2.7	2.2	2.6	2.4	2.6	2.5	2.5	2.6	2.5	3.1	2.4
2	2.3	0.0	2.3	2.1	2.2	2.2	2.4	2.2	2.2	2.3	2.1	2.9	2.1
3	2.5	2.0	0.0	1.9	1.8	2.0	2.0	2.0	2.0	1.9	1.9	2.7	1.8
4	2.3	2.2	2.2	0.0	2.2	2.4	2.4	2.4	2.3	2.2	2.1	2.9	2.2
5	2.4	2.0	2.1	2.0	0.0	2.2	2.0	2.0	1.9	2.3	1.8	3.0	2.1
6	2.3	2.1	2.1	2.2	2.2	0.0	2.1	2.0	2.1	2.2	2.1	2.7	2.0
7	2.4	2.3	2.3	2.1	2.1	2.2	0.0	2.3	2.2	2.4	2.1	2.9	2.3
8	2.4	2.1	2.2	2.3	2.1	2.0	2.2	0.0	1.8	2.4	2.0	2.8	1.9
9	2.5	2.1	2.1	2.1	1.9	2.1	2.1	1.8	0.0	2.2	1.8	2.7	2.0
10	2.5	2.3	2.1	2.1	2.2	2.3	2.4	2.2	2.1	0.0	2.2	2.7	2.0
11	2.5	2.1	2.1	2.0	2.0	2.2	2.3	2.2	1.9	2.2	0.0	2.8	2.1
12	3.2	3.2	3.2	3.0	3.2	3.1	3.1	3.0	2.9	3.0	3.0	0.0	3.0
13	2.3	2.1	2.1	2.2	2.2	2.1	2.3	2.0	2.1	2.0	2.1	2.7	0.0

Table 4.7 shows the values of $AND(P^j, P^{j'})$ for each pair of Pareto sets. Note that the diagonal only contains zeroes: in the same Pareto set an identical solutions with $d(\underline{y}_i^j, \underline{y}_{i'}^{j'}) = 0$ can always be found. Another observation is that the indicator is not symmetrical, i.e. $AND(P^j, P^{j'}) \neq AND(P^{j'}, P^j)$: a solution $\underline{y}_i^j$ in set j that acts as nearest solution for solution $\underline{y}_{i'}^{j'}$ may not have solution $\underline{y}_{i'}^{j'}$ as nearest solution in set j'. Again, only minor differences can be observed in the table, with run 12 as an exception, probably because run 12 resulted in a smaller number of Pareto solutions. Furthermore, also for this indicator the average value within groups of runs (2.2) is slightly smaller than between groups of runs (2.3), which was expected, although these differences seem too small to be substantial.

The absolute values of $AND(P^j, P^{j'})$ have an implication as well: on average, the closest solution in another Pareto-optimal set can be found at a distance of 2.3. This implies that on average more than 2 of the 26 decision variables are different if a solution from another random outcome of the optimisation process is used, which may be a substantial difference when measures are implemented in practice. Note that in the definition of the indicator, all decision variables are weighted equally, while in reality some decision variables may represent much more comprehensive measures than others (i.e. a new rail line vs. a new park-and-ride facility at a small station). Consequently, the sensitivity of decision variables with respect to objective values will vary.

Table 4.8: The difference between fractions of nonzero decision variables in the Pareto sets $AFD(P^j, P^{j'})$

	1	2	3	4	5	6	7	8	9	10	11	12	13
1		9.6%	12.5%	8.7%	12.0%	9.3%	13.7%	11.7%	14.6%	9.4%	11.7%	9.8%	9.8%
2			8.3%	5.6%	7.7%	7.5%	8.3%	9.9%	9.7%	7.8%	5.9%	9.5%	6.7%
3				7.0%	6.4%	8.9%	8.7%	9.5%	9.4%	5.4%	6.1%	11.7%	6.9%
4					6.8%	8.3%	9.6%	10.5%	10.0%	5.9%	6.5%	10.8%	6.2%
5						7.5%	7.6%	8.1%	7.4%	7.6%	4.5%	11.7%	8.5%
6							8.0%	7.6%	9.7%	7.1%	7.8%	9.1%	7.4%
7								9.6%	9.0%	9.3%	6.7%	10.8%	9.2%
8									7.1%	10.0%	8.4%	9.7%	8.3%
9										10.2%	7.1%	11.3%	9.9%
10											6.8%	7.7%	5.8%
11												10.5%	7.1%
12													9.4%

Table 4.8 contains the values for $AFD(P^j, P^{j'})$. Relatively large differences between pairs of Pareto sets can be observed. The absolute value of this indicator has an implication as well: on average, the fraction of Pareto solutions in which a certain decision variable is Pareto-optimal, differs 8.7%. This can be considered substantial, taking into account that some decision variables are (almost) always equal to 0, or (almost) always nonzero. Furthermore, also for this indicator the average value within groups of runs (8.4%) is slightly smaller than between groups of runs (8.8%) as expected, although these differences seem too small to be substantial.

4.5.4. Comparing multiple Pareto sets

Finally, Figure 4.10 shows the attainment plots for the 5 identical parameter settings, for the objective combinations of total travel time - CO_2 emissions and urban space used - operating benefit. The scale of the axes contains the range covered by the entire Pareto set. The dark red area in the plot is the area that is dominated by all 5 Pareto sets, while the dark blue area is not dominated by any of the Pareto sets. The intermediate area is dominated by some of the Pareto sets and visualises the variation between the sets. For the first combination of objectives, some clear variation in results can be observed, mainly in the extreme ends of the Pareto set. This could already be observed by the best values found for individual objectives (Figure 4.8). The second combination of objectives only shows a small boundary of variation, also because the trade-off region of the Pareto set is much bigger with these objectives.
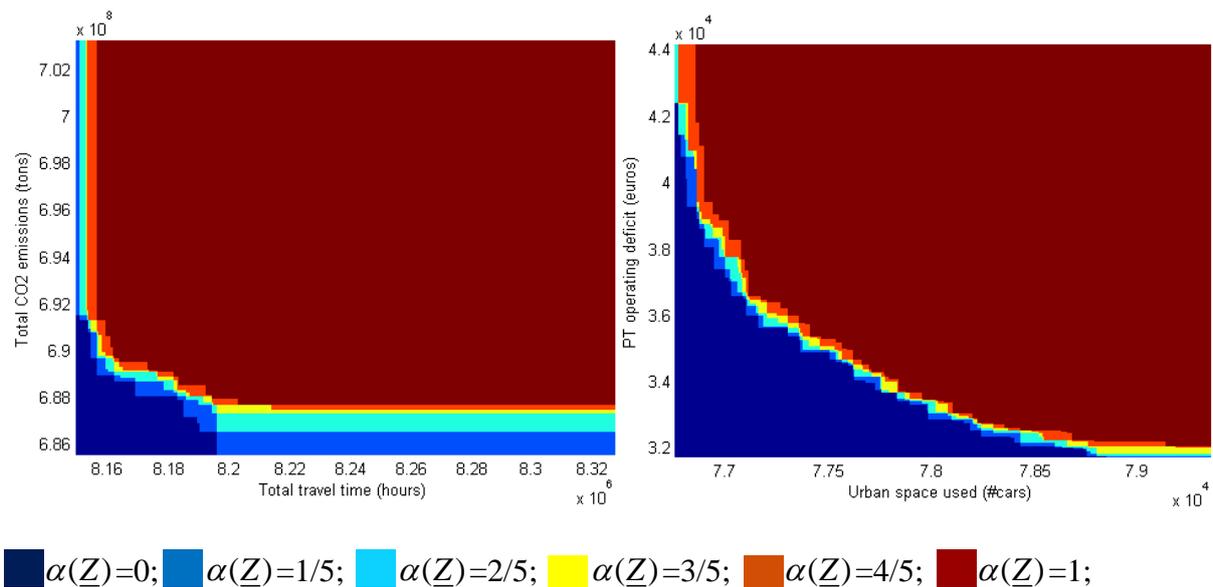
Figure 4.10: The attainment plots of the outcome of optimisation processes 1-5. The blue surface is attained by none of the Pareto sets, while the dark red surface is attained by all Pareto sets

4.5.5. Conclusions

The first conclusion from this analysis is that the primary goal to minimise variation in the resulting objective values is met by the algorithm, because the differences in the attained objective values are small among the runs with the same parameters: the influence of stochastic principles on the attained objective values is small. Furthermore, the differences in the attained objective values are small when comparing runs with different values for generation size and number of generations. However, the result is sensitive to the mutation rate, because a different mutation rate showed differences in attained objective values. At least a few different values for this parameter should be tried when setting up a new case study.

Secondly, it can be concluded that the decision variable values that occur in the final Pareto sets from runs with the same parameters are substantially different in all analysed pairs of sets. This implies that the advised decisions to be taken are influenced by the stochastic

principles used in the algorithm. Apparently many good network designs are possible in this case study, given the objectives that were defined. As long as a good design is chosen by the algorithm, it may not a problem if a different network design is chosen. When taking decisions based on Pareto sets that are constructed by such an algorithm it is important to be aware of this. Sound methods to analyse the Pareto set are useful in this matter.

When comparing runs with different values for generation size and number of generations, variations in decision variable values are only slightly higher than the variations caused by the stochastic principles of the algorithm. Because this difference in variation is too small to be substantial, the influence of generation size and number of generations in the algorithm is not large, implying that the risk for running the algorithm with bad parameter values is small. However, optimisation results are again more sensitive to the value of the mutation rate.

4.6.   Conclusion

Considering the high computation times in both case studies, a solution method is needed that produces high-quality results, while using a limited number of function evaluations. To this end, the MO evolutionary algorithms NSGAII and ε-NSGAII are applied to the multimodal passenger transportation NDP. The performance of ε-NSGAII is compared with the performance of NSGAII. Furthermore, the variation among optimisation outcomes using NSGAII is assessed, both with the same parameter settings and with different parameter settings. Various indicators and methods are used to assess both the differences between the attained objective values as well as the differences between decision variables of the Pareto sets.

The experiments confirm that earlier findings (Kollat and Reed, 2006; Albarello et al., 2014) of ε-NSGAII outperforming NSGAII on most indicators are also true for the multimodal MO-NDP. Especially in the early stage of running the algorithm, larger progress is observed. Furthermore, ε-NSGAII found comparable or better values for the objective functions individually. So ε-NSGAII is more efficient to find high quality trade-off solutions in the multimodal MO-NDP, especially for objective functions that are highly computationally expensive to evaluate, like in the case studies in this research. Concluding, ε-NSGAII is a suitable optimisation algorithm to find a high quality approximation of the Pareto set that may serve as input for decision support information. However, since results of evolutionary algorithms are always case-specific, it would be good to do more experiments on other real-life NDPs, e.g. a different study area or different decision variables.

Further experiments showed that the variation in the attained objective function values among multiple random outcomes of the optimisation process is small, both for runs with the same parameters and for runs with different parameter settings. This implies that in each run high quality solutions are found, leading to the conclusion that the method is suitable for the purpose it serves in this research: providing input for decision support information to give insight into the measures to be taken to reach objectives, depending on trade-offs between objectives.

However, the decisions to be taken are substantially different in different outcomes of the optimisation algorithm, so the advised decisions are influenced by the stochastic principles of

the algorithm. Combining this observation with the small differences observed in objective function values leads to the conclusion that the same can be achieved by taking different measures. Evidently the objective functions have flat shapes, with several combinations of decision variables that result in (almost) the same objective values, making many solutions indifferent. Concluding, this issue is more related to the definition of the optimisation problem than to the solution algorithm. This may be caused by decision variables that only have little influence on the objective function values. It may also be caused by interdependencies between decision variables. Therefore, in Chapter 5 the sensitivity of decision variables with respect to the objective values is tested, to identify important and less important decision variables.

Chapter 5: Optimisation results from the Randstad case study to support decision making

In chapter 4 an optimisation method was selected that is able to produce a high quality Pareto set within a reasonable computation time. This Pareto set is the (approximation of) the solution of the multi-objective, multimodal passenger network design problem. However, the concept of a Pareto set may be too complicated for decision makers to be directly interpreted. Following the needs of decision makers to choose one final solution for implementation, the objective of this chapter is twofold. The first objective is to analyse and understand the problem using information in the Pareto set. This is done for the case study in the Northern part of the Randstad area in the Netherlands. The second objective is to enhance existing methods for decision support based on results from multi-objective optimisation, which help decision makers choose a final solution.

Several methods are demonstrated and developed further which is useful to finally come to one best compromise solution for implementation. These methods involve deriving problem knowledge and reducing the number of solutions in the set. The outcomes of the methods provide insight into the extent to which measures that enable multimodal trip making can contribute to the sustainability objectives in the northern part of the Randstad area in the Netherlands (either simultaneously or by making choices among objectives) and which (types of) measures can contribute to these objectives. These measures involve park-and-ride facilities, new train stations (that can also be reached by bicycle) and frequency setting of PT lines.

5.1. Introduction: needs of decision makers

The Pareto-optimal set is the outcome of a multi-objective (MO) optimisation procedure, which contains all network solutions that might be optimal for the decision maker dependent on the compensation principle used to combine the objectives. In this section the information needs of decision makers are identified, so they are able to use the information in the Pareto set. These needs are derived from literature and from three interviews with Dutch policy officers, who prepare decision making at three different local governments in the Netherlands (municipality of Amsterdam, city region of Amsterdam and province of Overijssel). An overview of the resulting needs is presented in Table 5.1. Each of these needs corresponds with a certain type of information that supports the decision making process. The ultimate goal of this decision making process is to choose the final solution for implementation that best fits the policy of the region that is studied.

Compared to the current practice in Dutch policy making in the field of transportation, using optimisation results in the decision making process has several advantages. First, all possible solutions are considered simultaneously based on the policy objectives and the defined decision variables, without knowing the exact preferences of the decision maker. The search process is not limited to only a few (often expert-judgment based) solutions, for which it is not known whether they are Pareto-optimal. Second, after optimisation the Pareto set can be used to interactively reveal consequences of choosing certain decisions or choosing preferred objectives. Third, suboptimal solutions are excluded beforehand, so no valuable time is used to discuss them.

Disadvantages of using MO optimisation in policy making practice are the complexity and the work load. The methods and results are difficult to explain to decision makers, require long computation times and require a suitable transportation model for the region under study (that is fast enough and still has enough quality to be able to calculate objective values with enough accuracy). In addition to that, the process of investigating all possible measures among the involved stakeholders may be labour intensive. These possible measures have to be manually coded in the network of the transportation model (i.e. to define solution space).

Visualisation techniques can help present the Pareto set in such a way that it provides insight to decision makers. Deriving these insights from Pareto set visualisations, also called manual innovization by Deb (2003), is a method where both common and different properties of all obtained Pareto solutions are identified, as well as where and how these similarities and differences occur. During the interviews with policy officers it became clear that visualisations of the Pareto set are useful, but in addition background information is needed for a correct interpretation.

A next step is to mathematically analyse the Pareto set to derive useful, more general problem knowledge. Trade-off information between objectives is a common result that is derived from Pareto sets, but is not straightforward when more than 2 objectives are considered (Wismans et al., 2013). In the interviews this information was valued as well, because this is related to social cost-benefit analysis: a decision support methodology that is often used in the Dutch infrastructural planning practice (Mouter et al., 2013). For the case study in the Randstad area, such trade-offs are determined in the current chapter, as well as minimum values per objective, range covered per objective and correlation between objectives.

Design rules for decision variables in relation with objectives are more complicated to obtain. Atashkari et al. (2005) identified such rules for a problem with a limited number of continuous decision variables. The mathematical structure of relationships between objectives, decision variables and constraints is analysed to discover useful design principles. Deb et al. (2014) constructed such rules by fitting functions to represent these relationships, which they call automated innovization. This method only works for optimisation problems with few continuous decision variables, so it cannot easily be applied to the case study in multimodal passenger transportation network design, which has many discrete variables. In this chapter mathematical techniques are adjusted in a way that they can be applied to Pareto sets that result from a problem with many discrete decision variables. The sensitivity of the decision variables with respect to the objectives is tested. Furthermore, during the interviews it became clear that policy officers are interested in so called 'no regret' and 'always regret' measures, i.e. measures that should always or never be taken, no matter which priorities are given to the (in this case four) objectives. Therefore for each decision variable the fraction of solutions in the Pareto set where that decision variable is active is determined, to indicate the effectiveness of the related measure (in relation to all defined objectives simultaneously).

Although such analyses provide more insight into the structure of the problem, in the end one solution that reflects the decision makers' preferences best has to be chosen for implementation. A reasonable number of alternatives for a decision maker to consider (during an explanatory phase in the planning process) is somewhere between 10 and 20 solutions (based on the interviews). In case of competing objectives the Pareto set is likely to be larger: the Pareto set can easily contain a few hundred solutions. The complete set then provides too much or too complex information for the decision maker, as is also argued by Kulturel-Konak et al. (2008) and Chaudhari et al. (2010).

The problem of choosing the best compromise solution to implement is rarely addressed in relation to the MO-NDP. During the interviews it was discovered that policy officers like to cope with a large Pareto set by pruning it step by step in dialogue with a decision maker. This can be done by setting additional bounds on objective values, as is done by Kasprzyk et al. (2013), and / or by including / excluding certain values for decision variables that are or are not politically desirable. If there are still multiple solutions left after this process, traditional methods can be used to choose the best compromise solution from the remaining solutions. Another method is to define priorities among objectives, where a decision maker receives direct feedback on the consequences of these priorities. Such a method has the advantage that it is relatively easy to understand and it provides the decision maker with direct feedback on his choices. On the other hand, early choices may exclude large areas of the Pareto set, so there is a risk that the final decision is different from the decision that would be made based on the complete overview of possible solutions.

To provide a concise overview of possible solutions, the Pareto set can be systematically pruned to a small number of solutions to choose from. This was done by Taboada et al. (2007) and Ulrich (2013), whose pruning methods showed to be useful to provide a comprehensive overview of the main choices for the decision makers: clusters of similar solutions are identified, where the pruned set contains one representative solution per cluster. When pruning the Pareto set, the challenge is to select a representative subset of the original set, i.e. a subset that maintains the main characteristics of the original set. The pruned Pareto set

should therefore contain solutions that show significant differences, the extreme solutions and solutions that are divers, i.e. equally spread along the entire Pareto front. Taboada et al. (2007) only paid attention to the objective space during the pruning process. Ulrich (2013) made clusters based on both objective space and decision space. As good clusters in decision space do not necessarily correspond with good clusters in objective space, she formulates the clustering problem as a bi-objective optimisation problem and proposes an MO evolutionary algorithm to generate promising trade-off partitionings, which contain similar designs and are located in compact regions in objective space. During the interviews, the policy officers agreed that reducing the number of solutions to choose from helps making their task easier.

In this chapter several methods that systematically reduce the number of solutions are applied. Their aim is to help the decision maker in choosing the best compromise solution. These pruning methods include two conceptual improvements for one specific method, the PIT filter, which focuses on regions of the Pareto set with significant trade-off (Mattson et al., 2004). The enhanced method, the original filter and two other possible pruning methods (convex hull and K-means clustering) are compared.

Table 5.1: Overview of needs of decision makers

Need of decision maker	Made operational by	Originates from	Type of results	Section
Visualisation of Pareto-optimal solutions	Several types of plots	Interviews, Hettenhausen et al. (2009) and Deb (2003)	Case-specific	5.2
Trade-off information between objectives	Range covered per objective; trade-off values between objectives	Interviews, Wismans et al. (2013)	Case-specific	5.3.1. 5.3.2.
Derive useful, more general knowledge of the problem	Sensitivity of the decision variables with respect to the objectives; relation between types of decision variables with objective values; correlation between objectives; correlation between types of decision variables	Deb et al. (2014)	Case-specific; Method-logical	5.3.3.
Identification of 'no regret' and 'always regret' decision variables	Percentage of active decision variables in Pareto set	Interviews	Case-specific	5.3.3.
Interactively set priorities to objectives	Set a bound per objective step-by-step; until only one solution remains	Interviews	Case-specific; Method-logical	5.4.1.

Search for compromise solutions	Choose the solution that does not have a bad value for any objective	Interviews	Case-specific; Methodo-logical	5.4.1.
Take political preferences for certain measures into account	Including / excluding certain decision variables	Interviews	Case-specific; Methodo-logical	5.4.2.
Reduce the number of choice options	Systematically reduce the number of solutions by pruning	Interviews; Taboada et al. (2007); Mattson et al. (2004)	Methodo-logical	5.5.

The elements that contribute to the goal of choosing one final solution for implementation from the Pareto set are summarised in Table 5.1. All elements are made operational by a mathematical analysis of the Pareto set. Results of these analyses are case-specific, but the methods may be applied to other Pareto sets as well. The case-specific results provide general problem knowledge for the case under study, which is the design of a multimodal passenger network in the Northern part of the Randstad area in the Netherlands. Next to these case-specific results, methodological results provide insights that may be generalised and methods that may be applied to other MO design problems. It is indicated in Table 5.1 whether an element mainly provides case-specific knowledge, mainly contributes to methodology development, or both.

These elements are made operational and applied to one Pareto set that was constructed by ε-NSGAII for case study 2. This algorithm is used, because in chapter 4 it resulted in the highest quality in terms of attained objective values. Case study 2 is used, because it is the larger case study, including more decision variables and a more detailed network in the lower-level model. The objective values are therefore calculated more accurately, which is relevant when relations between these objectives and between objectives and decision variables are interpreted. In total 2384 solutions were calculated during running the algorithm. From these solutions, 210 were Pareto-optimal (i.e. non-dominated). This Pareto set is an approximation of the real Pareto set, since it would take too much computation time to calculate all solutions and thus the real Pareto set is not known.

In section 5.2 the Pareto set is visualised in several ways. In section 5.3 more general rules are derived from the Pareto set that are valid for the study area in the Northern part of the Randstad area in the Netherlands. In section 5.4 the Pareto set is pruned step-by-step to one solution by applying a manual technique that uses visualisation of the Pareto set. In section 5.5 pruning techniques are applied to systematically select a number of solutions to provide a comprehensive overview of the complete Pareto. Section 5.6. contains the conclusions, both on methodology development and on case-specific problem knowledge.

5.2.   Visualisation of optimisation results

The first step for a Pareto set to be useful information to support decision making is visualising the Pareto set. The Pareto set contains information that is presented in various

ways in this section. This information helps to understand the network design problem (NDP) in a multimodal context better and to enable a decision maker to choose one solution from the Pareto set for implementation. When visualising the Pareto set, a distinction can be made between visualising objective values of the Pareto set and visualising values of decision variables in the Pareto set. According to Deb (2003), visualisation could help obtain a holistic view of the data, but visualisation of the whole Pareto-optimal set in more than three dimensions is a problem which may be too complicated to solve. According to Hettenhausen et al. (2009), visualisation of the Pareto frontier can be carried out up to four or five objectives, but it requires more cognitive effort on the viewer's part if the number of objectives increases. Although visualising four dimensional Pareto sets is difficult various ways of visualising the objective values are presented. Therefore, in this section the results of the multimodal NDP are presented using various ways to visualise the Pareto set: scatter plots, a parallel coordinate plot, decision maps and a heatmap.

5.2.1. Traditional visualisation

The scatter plot shown in Figure 5.1 is a common way to visualise a Pareto set, especially to show trade-offs between objectives. Since only two objectives can be shown per objective, several plots are needed to show all interactions between objectives in a scatterplot matrix (Hettenhausen et al., 2009). Here, only 2 pairs of 2 objectives are shown. The plots also show solutions which are not Pareto-optimal when only the two objectives in the plot are considered, but these solutions are Pareto-optimal as a result of the four considered objectives during optimisation. The first plot for urban space used (USU) and PT operating deficit (OpD) shows a clear trade-off. The first improvements in USU are relatively cost-efficient, but further improvements become more expensive in terms of OpD. The second pair of objectives is total travel time (TTT) and CO_2 emissions (CE). These objectives are rather in line with each other: two clusters can be observed, one with lower TTT and lower CE and one with higher values for both objectives. On the other hand, within the lower left cluster a trade-off between these two objectives can be observed as well.

A parallel coordinate plot (see Figure 5.2) captures all 4 objectives in one plot (as used earlier by Kasprzyk et al. (2013)). Normalisation per objective is required for this plot, because the four objectives have different orders of magnitude and different units. The normalised values range from 0 to 1, corresponding to the minimum per objective and the maximum per objective. When interpreting these normalised value, it should be noted that there is a difference per objective in terms of absolute difference (and that this absolute difference also depends on the unit). Furthermore, the impact of this absolute difference is different depending on the objective. This is especially relevant when weight factors are used to combine normalised objectives.
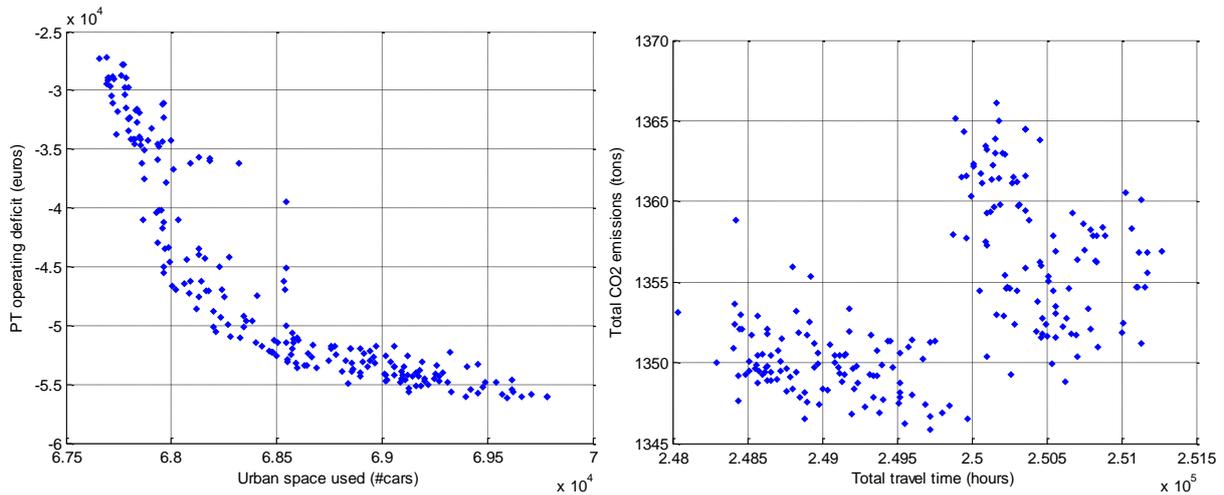
Figure 5.1: Two scatter plots of the Pareto set: one dot represents one solution in the set, where the corresponding scores for the two objective values can be read on the two axes
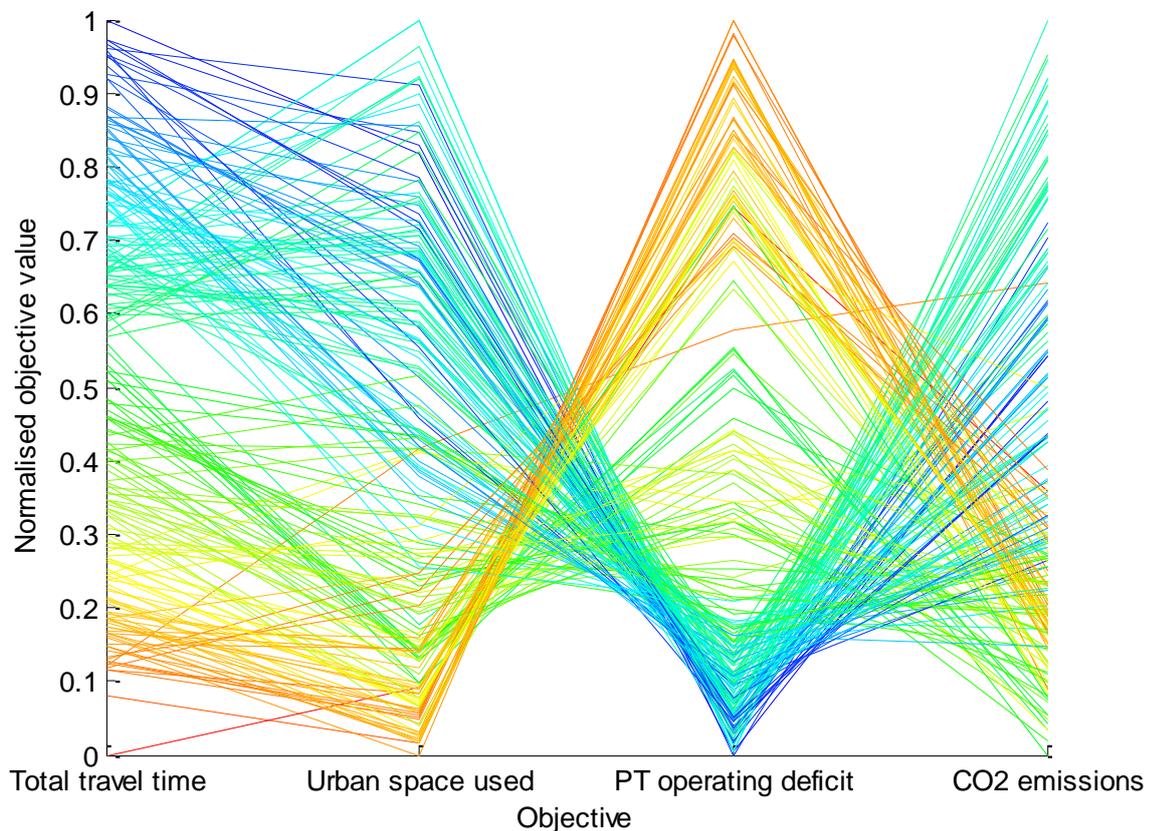
Figure 5.2: Parallel coordinate plot of the Pareto set: one line represents one solution, where the normalised values of the 4 objective values are plotted in the 4 columns. The lines are coloured using their value for total travel time

In a parallel coordinate plot trade-offs between objectives cannot be observed as directly as in a scatter plot. Using a colour scale to represent one specific objective value improves this. When looking at the colour distribution, it can be observed that high TTT implies low OpD, but usually also high USU. The relation with CE is less clear, but roughly TTT and CE are in

line (as was earlier observed in the scatter plot). The ordering and colouring of the objectives can be changed, to emphasise different relations and therefore providing different insights.

5.2.2. More dimensional scatter plots

By using different colours for the dots in a scatter plot, one additional objective is included in the scatter plot (see Figure 5.3). Hettenhausen et al. (2009) earlier referred to this way of visualisation as decision map. A limited number of categories is defined for the 3rd objective, each represented by a separate colour. This type of visualisation can for instance be used to visualise the effect of introducing a constraint for the 3rd objective on the Pareto front of the two objectives at the axes, i.e. its effect on which solutions remain Pareto-optimal and their related outcome concerning the two other objectives.
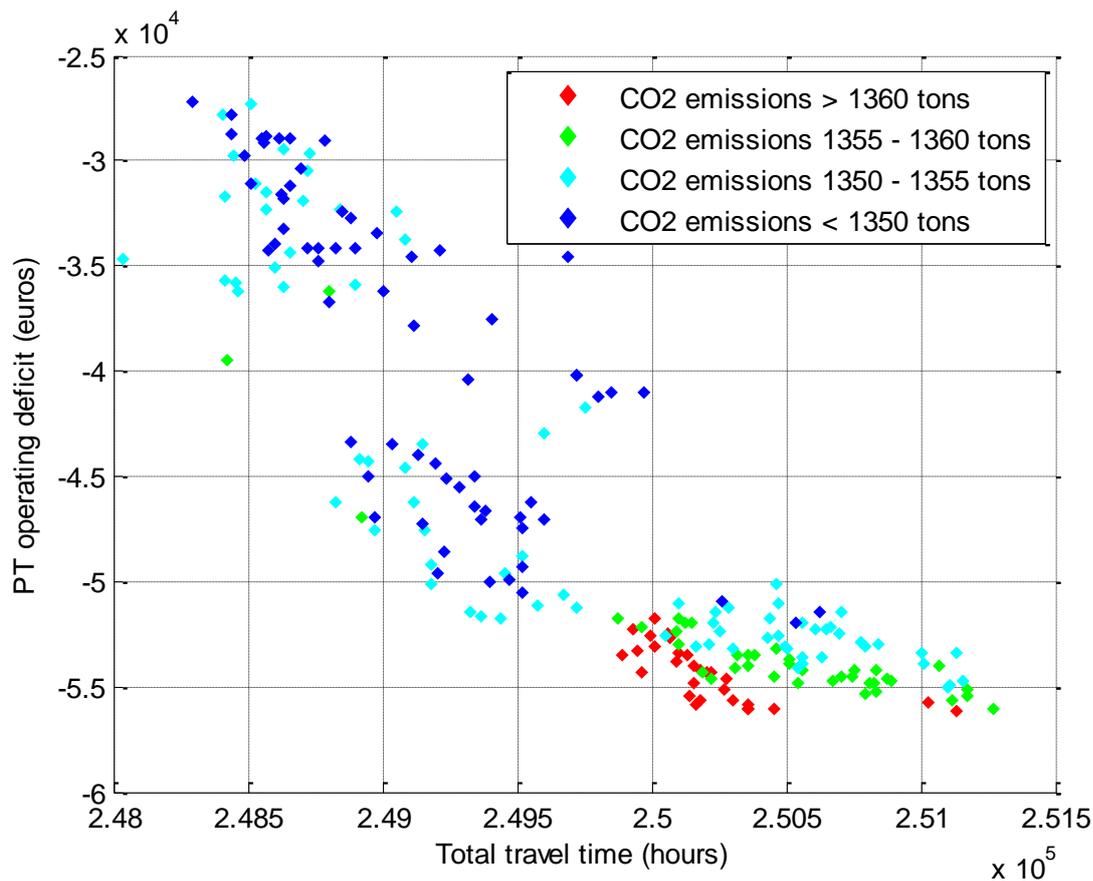
Figure 5.3: A 2D scatterplot for total travel time and operating deficit, with a distinction between CE categories (also called decision map)

It can be observed that solutions with high CE only occur in an area of the plot with low OpD and medium to high TTT. On the other hand, a very interesting observation is that when only solutions with low CE are considered (i.e. less than 1350 tons, or similarly, less than 1351.9 tons, the CE in the base situation), still a large variation exists in scores for TTT and OpD. In other words, when an additional constraint is set for CE, there is still a choice between solutions with low TTT, with low OpD or trade-off solutions with intermediate values for both objectives.

In Figure 5.4 the 3^{rd} objective is represented by a continuous colour scale. This could be referred to as a continuous decision map. It contains more detailed information than the decision map based on categories, but it may be more difficult to read. Which of the two is to be preferred depends on the desired accuracy of the 3^{rd} objective. Especially along the Pareto front for the two objectives TTT and OpD interesting trade-off information is found. The solutions that belong to this front are the best when only TTT and OpD are considered, but moving a little away from the front (in upper right direction) results in a large gain in CE at little costs in terms of TTT and / or OpD.
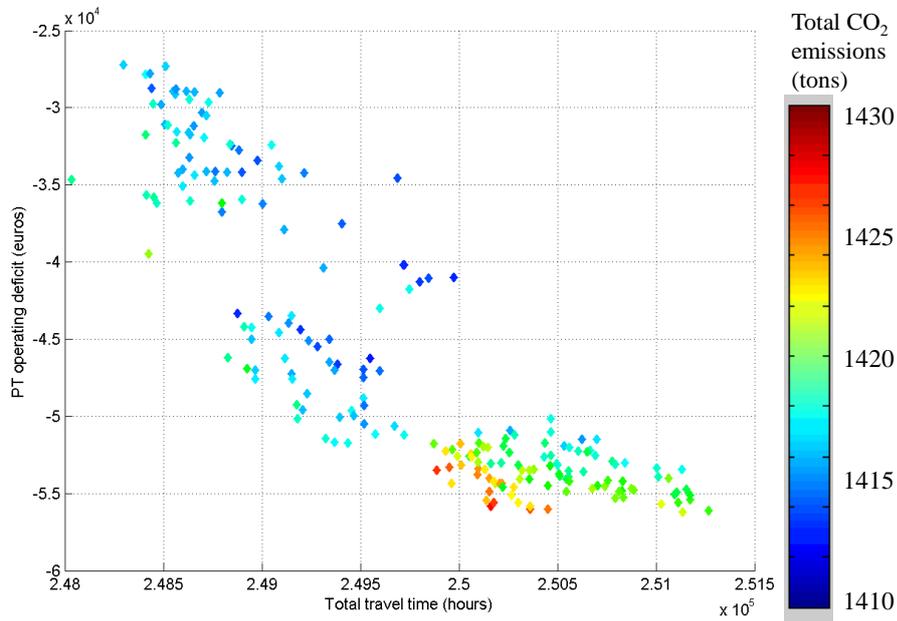
Figure 5.4: Continuous decision map of the Pareto set

It is even possible to include information on more objectives in a figure like this, by using arrows instead of dots in a scatter plot, where the orientation of the arrow represents a fourth objective and the size of the arrow represents a fifth objective, see for example Kasprzyk et al. (2013). It is questionable whether all this information is interpreted correctly by the viewer when presented together in one figure.

5.2.3. Decision variable values in specific solutions

In contrast to earlier visualisations, the heatmap in Figure 5.5 shows the values of decision variables of all solutions in the Pareto set, see also Pryke et al. (2007). Each row corresponds to one solution i and each column corresponds to one decision variable v. The grayscale in each cell represents the value of the corresponding decision variable y_{iv} in solution i. Recall that the values of decision variables y_v range from 0 to 1, where each value of a decision variable represents a real value, for example frequency of PT line l $F_l(y_v)$. The solutions (i.e. the rows) are sorted based on decreasing values for CO_2 emissions (CE).

This plot has the advantage that all values of decision variables in the entire Pareto are given in one overview. However, this amount of data is close to the limitations on a viewer's ability to interpret such a large heatmap. Sorting the solutions (in this case based on CE) helps to

derive information from it. The first clear observation is that solutions with higher CE (in the upper half of the plot) have lower values for train frequencies (variables 25-36). Also for bus frequencies (variables 17-24) this is the case, but this relation is less clear. For the other variable types no relations with CE are identified. Specifically for variable 12 and variable 16, park-and-ride (P&R) facilities at Zuiderzeeweg in Amsterdam and at Schiphol North respectively, it can be observed that these measures are only included in solutions with CE values that are relatively low. Finally, the heatmap shows that some variables (e.g. variables 6 and 11) are present almost in all solutions in the Pareto set, while other variables (e.g. variables 5 and 10) are present in (almost) no solutions in the Pareto set. This is analysed in more detail in Section 5.3.3).
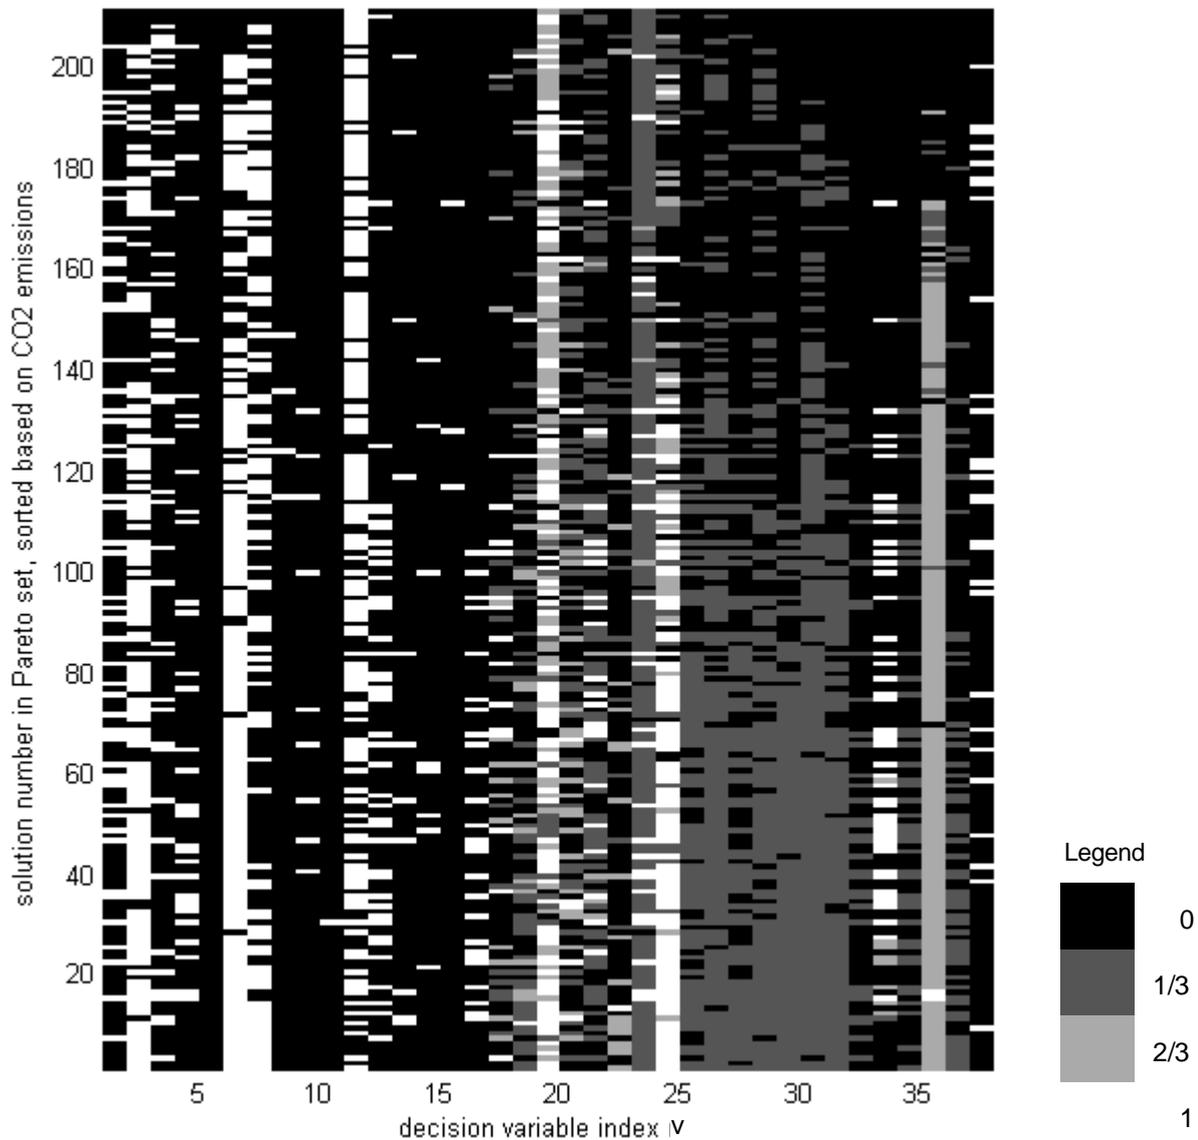
Figure 5.5: Heatmap of all Pareto solutions

5.2.4. Discussion

Various ways to visualise the information in the Pareto set are presented. A parallel coordinate plot is useful to get a general overview of the data, because all objectives can be

included in one plot. It is also suitable to visualise the position of one solution or a selection of solutions in the entire set. Therefore, in Section 5.3.1 this type of plot is used to illustrate the position of the best solutions per objective in the entire Pareto set and in Section 5.4 to visualise the results during a step-by-step pruning process to come to one final solution for implementation. Scatter plots provide more detailed information on trade-offs between 2 objectives (or 3 objectives in a decision map) and are therefore more suitable to choose an exact solution once it is known which of the objectives are to be preferred, because it is complicated to include more than 3 objectives in a scatter plot. In Section 5.3.2 scatter plots are used to illustrate trade-offs between pairs of objectives.

Some case-specific conclusions can be drawn concerning trade-offs between objectives. When only solutions with low CE are selected, it is still possible to cover largely diverse scores for OpD and TTT. Furthermore, when searching for a trade-off solutions between OpD, TTT and CE, small losses in OpD and TTT result in a large gain in CE. In general, higher frequencies for train lines result in lower CE.

5.3.   Analysis of the Pareto set for decision support

In this section more general information is derived from the Pareto set. The ideal result from such an analysis would be design rules, stating that certain (combinations of) objectives require certain types of measures to be taken. Unfortunately, the design problem to find an optimal multimodal transportation network contains many interdependencies between decision variables and objectives, so it is not straightforward to set up such design rules. Furthermore, the results are case-specific, which means that the effectiveness of measures may depend on the geography of the study area. As a result, general rules cannot be given with certainty. Still, in this section insights are given into the effects of the proposed measures at a more aggregate level and interdependencies between objectives and decision variables, for the study area in the Amsterdam metropolitan area. First, the range of objective values covered by the Pareto set is given and put in perspective by comparing these values with the base solution and with measures outside of the scope of the defined multimodal network optimisation problem. Second, as a result of the MO approach, information is derived on the relations between multiple objectives. Finally, relations are established between decision variable values and objective values.

5.3.1.   Scores reached per objective

The ultimate goal of the multimodal NDP, as formulated in the context of the 'Sustainable Accessibility of the Randstad' research programme, is to make the transportation network in the Randstad more sustainable. This general objective is put into operation by defining four objective functions. The extent to which the formulated objective values can be improved (i.e. reduced compared to the base solution) indicates to what extent this goal can be reached by the selected measures. In this section the minimum and maximum values per objective are compared with the objective values in the base solution. The base solution is defined as the situation where the current plans for 2030 are realised. In other words, this is the most likely transportation option network that will be developed for 2030 when all known plans are realised. More details on the base network are given in Section 5.3.3.

In Figure 5.6. the optimal solutions per objective are highlighted in a parallel coordinate plot in blue and the base solution is shown in black. It can be observed that for three out of four optimal solutions per objective the value of at least one other objectives is worse than the value of the base solution. Although in the end this depends on the preferences of the decision maker, it sounds reasonable to search for a solution that has improvement with respect to the base solution for all objectives. The optimal solution for CE comes with improvements on all four objectives, but the improvements in TTT and OpD are limited. It appears to be possible to find a solution that has a relative improvement of at least 0.40% for all four objectives (plotted in red in the figure). In total, 47 out of 210 Pareto-optimal solutions have better values for all objectives than the base solution (i.e. dominate the base solution).
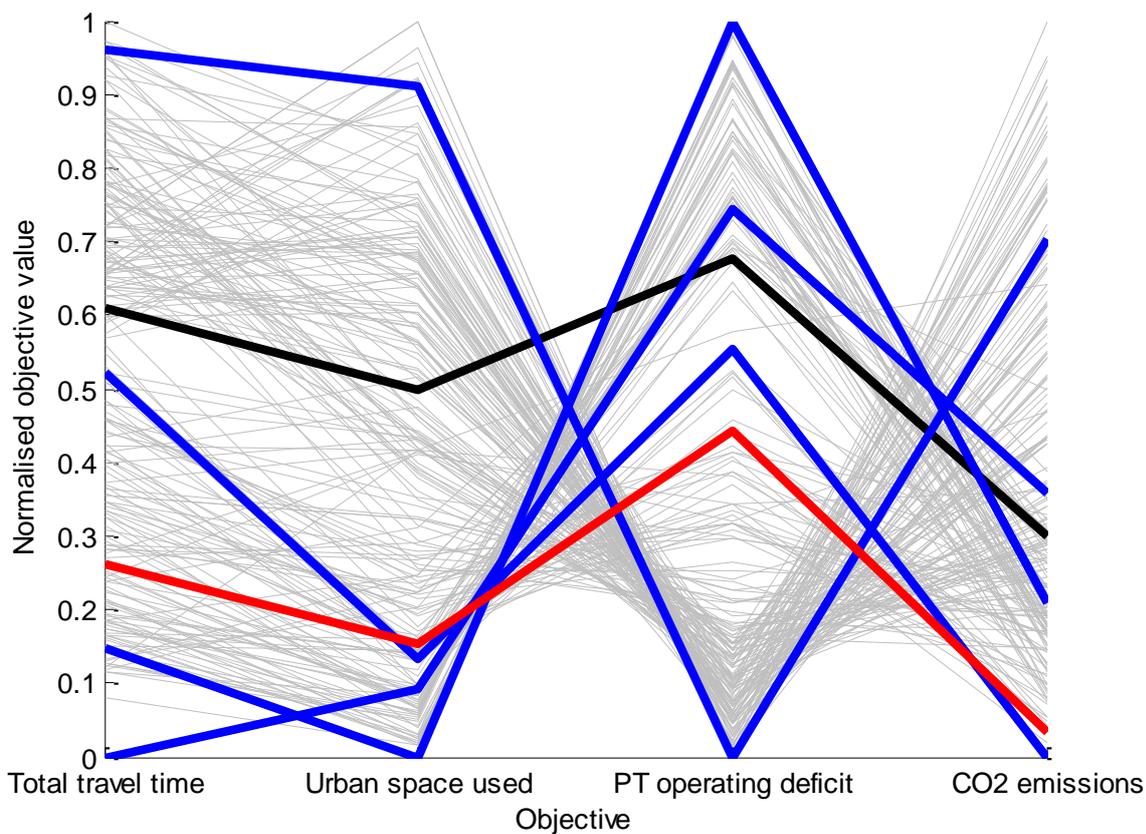
Figure 5.6: Parallel coordinate plot for the optimal solutions per objective (blue), for the base solution (black) and for one solution that gives considerable improvement with respect to the base solution in all four objectives (red)

In Table 5.2 the absolute minimum and maximum values for each individual objective are shown, as well as the relative change with respect to the base solution. Since OpD is the balance resulting from costs minus revenues, the relative values for OpD are determined by dividing the difference in OpD (compared to the base solution) by the total operating costs in the base solution (instead of the value for OpD itself, which can be both positive and negative). Improvements with respect to the base solution are possible for all objectives, but there is a large difference among the objectives. CE can only be reduced by 0.5%, while OpD can be reduced by more than 6%. Note that OpD turns out to be negative in all solutions in the Pareto set, so in fact there is an operating surplus. Still the definition from Section 3.3 is

maintained, so that the objective is to be minimised, consistent with the other three objectives. The large reduction potential for OpD is a direct result of the measures in the case study: putting additional PT into service directly involves costs (and taking PT out of service saves costs). USU can be reduced by making PT and P&R alternatives more attractive to attract the car users that have their origin or destination in an urban area, resulting in the possibility to reduce USU by more than 1.5%. TTT is more difficult to reduce, because travel time of cars is only affected by the measures to a small extent. Travel time of PT travellers can be reduced more strongly by the measures, but this is only a part of all trips made in the study area. Finally CE are difficult to influence by the measures, because reducing CE is only possible by attracting car drivers to the PT network, which is only possible by putting more PT into service. These additional PT services increase CE, partly nullifying the positive effect of attracting car users.

The contribution of the measures is small in relative terms, because the objective functions are defined as the sum over the entire multimodal transportation network in the study area, implying large aggregate values. The measures at transfer locations and in the PT network typically only influence a limited part of the network (the car network is hardly influenced). On the parts of the network that are influences, the relative gains in objective values are much larger. Furthermore, in absolute terms these potential gains per objective are considerable, e.g. the maximum reduction in CE is still 12 tons per AM peak, equivalent to the daily CE of more than 500 Dutch households.

Table 5.2: Minimum and maximum values found per objective, compared to the base situation. All objectives values are for a one hour period in the AM peak

	TTT (hours)	USU (# of cars)	OpD (euros)	CE (tons)
Base situation	250007	68720	-36594	1351.9
Minimum	248034	67659	-56203	1345.8
Maximum	251267	69786	-27218	1366.1
Relative range covered	1.29%	3.09%	9.14%	1.50%
Best possible individual improvement w.r.t. base	-0.79% -1973	-1.54% -1061	-6.19% -19609	-0.45% -6.1
Best possible simultaneous improvement w.r.t. base	-0.45% -1127	-1.07% -733	-2.12% -6730	-0.40% -5.4

Decision variable values of selected solutions

To visualise the values of decision variables for the selection of solutions in Figure 5.6 a heatmap containing these specific solutions is used (see Figure 5.7). In the first row the base solution is plotted for comparison reasons. Next, the four solutions that are optimal for the individual objectives are shown. Also in decision space, the optimal solutions for CE, TTT and USU are quite similar and different from the optimal solution for OpD. The three optimal solutions for CE, TTT and USU have many trains with a nonzero frequency (variable 25-36), while the optimal solution for OpD has many trains with frequency of zero. For train stations

and P&R facilities (variables 1 to 16), depending on the objective different values occur in the optimal solutions, but in general only few decision variables have nonzero values in these solutions. The next shown solution is the solution with considerable improvements with respect to the base solution (shown in red in Figure 5.6). This solution is similar as the optimal solution for CE: both in objective space and in decision space (see Figure 5.7: only two variables have a different value). In the parallel coordinate plot in Figure 5.2, one solution that is different from the rest of the solutions stands out: it has a low value for TTT, but at the same time a fairly high value for CE and a not so high value for OpD. This solution is shown in the 7[th] row of the heatmap. This solution has relatively low values for train and bus frequencies, and has a relatively large number of P&R facilities. This combination causes low costs (OpD), relatively large car use (high CE), but lower TTT, because car and P&R alternatives can have low travel times. Finally, the min-max solution (i.e. a compromise solution, see Section 5.4.1) typically looks like an intermediate solution, with all types of variables included to some extent, but it is remarkable that the local trains to Almere (variables 32, 33 and 34) all have zero frequency, i.e. they would not run at all. This would suggest that some objectives (low costs) are achieved in the Almere region, while other objectives (TTT, CE, USU) are achieved in the rest of the region. Equity (distribution of benefits of the measures throughout the region) may be an important issue here.
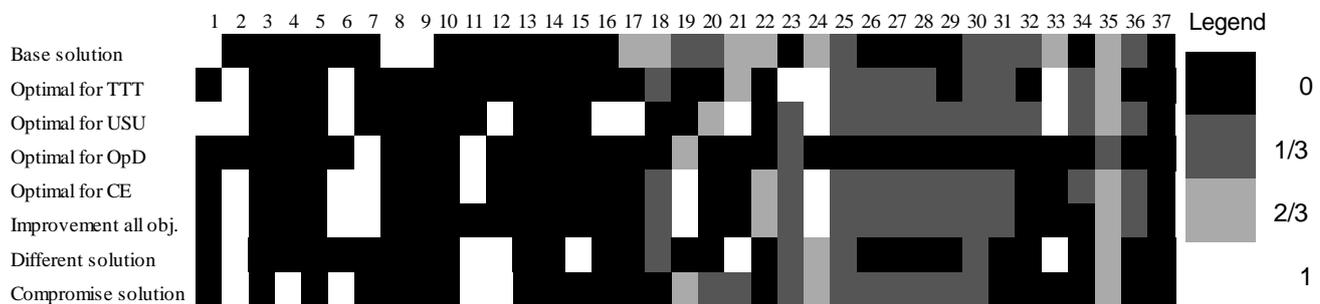
Figure 5.7: Heatmap of a selection of solutions

Scores of alternative measures

The selected decision variables facilitate transfers in the multimodal passenger transportation network by providing transfer locations and by making the PT network more attractive. This has resulted in simultaneous improvements in all 4 objectives. To put them in a wider perspective, these results are benchmarked with two measures that lie outside the scope of the defined optimisation problem. The first measure is to realise transit-oriented development (TOD) and the second measure is a time-independent, kilometre-based road user charge. These measures are first described and then the results follow.

The alternative measures TOD and the road user charge are applied to the entire study area and compared with the base situation. These measures are assessed using the same lower-level model that was used during optimisation (see Section 3.2). This model only contains route choice and mode choice as behavioural response, with fixed total transportation demand as input. Although effects on trip distribution and trip generation are expected for road pricing too, this model is used to enable the researcher to compare the results in this section with the earlier results. For TOD the effect on transportation demand is taken as input for the model.

These strong assumptions are in line with the purpose of this section to get an impression of the order of magnitude of the effects, rather than to obtain a highly reliable forecast.

In the TOD scenario (described in more detail in Section 6.2.2) the same level of economic growth is assumed as in the base demand forecast for 2030. All growth (like new jobs and new houses) takes place in the vicinity of train stations, instead of being spread out over the whole region, resulting in a demand pattern that better fits the PT network. This is typically a measure that can only be realised on the long term, unlike to redesigning the multimodal network, which can be implemented on an earlier term. For the government this policy does not necessarily involve expenses. It does involve costs for society, because residents and companies have to deviate from their preferred locations due to regulations. It is assumed that each zone will generate / attract the same number of trips per resident / job. The result is a different demand pattern that fits the PT network better: more trips are generated to and from zones near a train station and fewer trips are generated to and from zones far away from train stations. The distribution of travellers from each origin over the destinations is assumed to be constant. Similarly, the distribution of travellers from origins to each destination is assumed to be constant. As a result, the trip length distribution and average trip length stay approximately the same. The total number of trips slightly changes, from 407,676 in the base situation to 408,554 in the situation with TOD. Therefore in this section the objective TTT is replaced by total delay time (TDT, equal to TTT minus the travel time in a situation with no congestion or road pricing) for a fair comparison.

Next, road pricing is taken as an example of a 'push' measure, instead of only taking 'pull' measures. Road pricing is introduced in the model by increasing the distance based costs of car by 4 cents per kilometre (from €0.15 to €0.19 per km). This represents the most simple road pricing scheme that is not time-dependent or location-dependent: the effects on specific objectives can be larger when more differentiation is allowed in space and time. Determination of such a differentiated toll scheme can also be seen as an optimisation problem, as was for example done by Brands et al. (2009). The introduction of such a road pricing scheme generates revenues for the government, but it also involves implementation costs, i.e. to introduce and operate a registration system. These costs are not considered in this benchmark.

In Table 5.3 the results of the described alternative measures on the objective values are presented. The relative effects with respect to the base situation are shown. The effects on TDT are relative to TTT in the base situation, to be able to compare with the simultaneous improvements that can be achieved by solving the NDP, presented earlier in Table 5.2. The first observation is that the effects of road pricing are much larger than the effects of redesigning the multimodal transportation network. The main reason for this is that road pricing is a network-wide measure (applied to all roads in the study area), while redesigning the multimodal network only involves several local measures. As a consequence, road pricing influences the costs of the car alternative for all OD pairs, while redesigning the multimodal network only influences a selection of travel alternatives for a selection of OD pairs. The effect of road pricing on USU and CE is very positive, indicating that environmental and social sustainability can largely benefit from road pricing (or more generally speaking: from making the car alternative less attractive). However, the shift from car to PT on average results in an increased travel time, because travellers avoid the more expensive car alternative

that they would have preferred in the situation without road pricing. This harms economic sustainability. Although in reality a part of the former car users will not travel at all, PT use will increase, leading to increasing revenues and therefore in a large decrease of OpD. It is questionable whether the resulting PT loads realistically fit into the PT vehicles, because capacity of PT is not checked in the lower-level model. If this is the case, the capacity of PT service lines needs to be increased accordingly, implying that operating costs and PT-related CO_2 emissions will also increase, making the gains in OpD and CE smaller. In case the capacity of PT is not increased accordingly, a part of the new PT users will shift back to car due to overcrowded PT vehicles, despite the road pricing costs, also leading to more CO_2 emissions and a lower gain in OpD.

Compared to road pricing, the effect of TOD is relatively small, which can be explained as follows. Although more demand is concentrated near train stations, the total demand for transportation is still the same and travelling by car is still an option. Also for trips with origin and / or destination near a train station the car alternative may be faster than the PT alternative or multimodal alternative. So without any additional policy (for example a restrictive parking policy) TOD is not very effective in terms of objective values. Nevertheless a positive effect is observed for CE and for OpD, indicating that car trips are replaced by PT trips or multimodal trips. This is possible with a concurrent improvement in TDT. USU increases, because due to the higher densities in station areas, which are mainly highly urban, the total number of trips to these zones increases. Consequently, the number of car trips to these zones increases too, which is a negative effect in the current definition of USU.

Concluding, the introduction of a network-wide road user charge is more effective to achieve larger improvements in environmental and social sustainability than redesigning the multimodal transportation network. The main reason for that is that the car alternative is made less attractive. Economic sustainability is worse off with road pricing, probably because it is introduced as a single measure and not as a part of an optimisation approach. TOD policy is also more effective than redesigning the multimodal network when it comes to reducing environmental sustainability. TOD does not harm economic sustainability in terms of delay time, but increases the number of cars in highly urban areas. Redesigning the multimodal network is the only type of measure that enables simultaneous improvement for all considered aspects of sustainability, due to the adopted MO optimisation approach. It should also be noted that the costs of road pricing and TOD are not considered in this analysis, for example implementation expenses, societal costs of trips not being made any more and societal costs to deviate from preferred locations.

Table 5.3: Effects of the alternative measures TOD policy and fixed road pricing compared to the earlier results from solving the multimodal NDP

	TDT (hours)	TTT (hours)	USU (# of cars)	OpD (euros)	CE (tons)
Base situation (over 1 hour in the AM peak)	56222	250007	68720	-36594	1351.9
Effect of TOD policy	-2042	-0.8%	0.7%	-4.3%	-1.8%
Effect of fixed road pricing	9818	3.9%	-9.3%	-42%	-22%

5.3.2. Interdependencies between objectives

Correlation between objective values

In Table 5.4 the correlation matrix is shown for the 4 objective values of the solutions in the Pareto set (note that the matrix is symmetrical and the diagonal is filled with ones). OpD has a negative correlation with all other three objectives: it is opposed to all these objectives. The three remaining pairs of objectives have a positive correlation, so these objectives are more or less in line with each other. Especially USU is in line with CE, which can be explained because both objectives benefit from a reduction of car traffic. However, the correlation is not equal to 1, so still a trade-off exists between these objectives, as will be further elaborated on in the next section. The relation between CE and TTT is less clear: these objectives have a lower, but still positive correlation.

Table 5.4: Correlation between pairs of objective values

	TTT	USU	OpD	CE
Total travel time (TTT)	1.00	0.79	-0.87	0.53
Urban space used (USU)	0.79	1.00	-0.87	0.84
Operating deficit (OpD)	-0.87	-0.87	1.00	-0.62
CO_2 emissions (CE)	0.53	0.84	-0.62	1.00

Optimal solutions per objective: scores for other objectives

As was shown earlier in Figure 5.6, each of the four objectives has its own optimal solution. No solution exists that is optimal for two or more objectives at the same time, so it remains useful to take all four objectives into account during optimisation. In line with the observations concerning correlation between objectives, the objectives TTT, USU and CE are aligned to some extent and opposed to OpD: the optimal solutions for TTT, USU and CE are close to optimal for the other objectives in that group of three objectives. On the other hand, the optimal solutions for TTT, USU and CE are far from optimal for OpD and vice versa.

In Table 5.5 the objective values for these optimal solutions are given. The values in the table are index values, where for TTT, USU and CE the minimum known value is scaled to 100 (see the diagonal of the table). Since OpD is the balance resulting from costs and revenues (that can be both positive and negative), the value for OpD is determined by dividing the additional OpD (compared to the optimal solution for OpD) by the total operating costs in the base solution (see Section 5.3.3).

When looking at the table, for example the solution optimised for USU has a 0.2% worse score for TTT than the optimal solution for TTT, a 9.1% worse score for OpD than the optimal solution for OpD and a 0.3% worse score for CE than the optimal solution for CE. Note that each cell in the diagonal contains the best possible value 100.0, because the solutions were optimised for that objective. An interesting observation is that the relative costs in terms of TTT, USU and CE when choosing for the optimal solution for OpD is smaller than the relative costs of OpD when optimised for the other three objectives. This

emphasises the relevance of trade-offs between the absolute values of the objectives, as will be discussed in the next subsection.

Table 5.5: Relative scores of optimal solutions per objective for the other objectives

		Has relative score for			
		TTT	USU	OpD	CE
Solution optimised for	TTT	100.0	100.3	106.8	100.5
	USU	100.2	100.0	109.1	100.3
	OpD	101.3	102.9	100.0	101.1
	CE	100.7	100.4	105.1	100.0

Trade-off values between pairs of objectives

A trade-off value indicates the extent to which an increase for an objective has to be accepted, if a decision maker wants to decrease another objective by moving to another solution (e.g. it is possible to reduce CE with an amount x, but this means an increase in TTT of y) (Wismans, 2012). For each pair of solutions in the Pareto set such a trade-off exists for at least one pair of objectives, because each Pareto solution is non-dominated. In Eq. 5.1 trade-off is formally defined: the difference when moving from solution $\underline{y}$ to solution $\underline{y}'$ (or vice versa) for one objective w is divided by the difference in another objective w'.

$$TO_{w,w'}(\underline{y}, \underline{y}') = \frac{Z_w(\underline{y}) - Z_w(\underline{y}')}{Z_{w'}(\underline{y}) - Z_{w'}(\underline{y}')} \qquad (5.1)$$
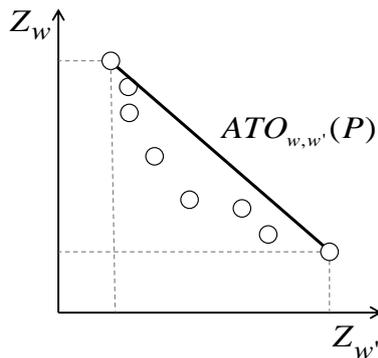
Figure 5.8: Illustration of the average trade-off value, based on the extreme solutions for a pair of objectives

Although this definition of trade-off is straightforward, using it is still complicated for two reasons. Firstly, this trade-off value can be calculated for each pair of solutions and there are many pairs possible: the Pareto set analysed in this chapter contains 210 solutions, resulting in 210*209/2 = 21945 pairs. Secondly, when the trade-off is determined for one pair of objectives, other objective values will change, too (positively or negatively). To be complete for all objectives, the change in all objectives should be assessed simultaneously, which makes the task more complicated. Therefore the average trade-off value is introduced,

providing one value (see Figure 5.8) for trade-off between a pair of objectives, along the entire (known) Pareto front. In Eq. 5.2 average trade-off is formally defined.

$$ATO_{w,w'}(P) = \frac{Z_w\left(\underset{\underline{y} \in P}{\arg\min} \, Z_{w'}(\underline{y})\right) - \underset{\underline{y} \in P}{\min} Z_w(\underline{y})}{Z_{w'}\left(\underset{\underline{y} \in P}{\arg\min} \, Z_w(\underline{y})\right) - \underset{\underline{y} \in P}{\min} Z_{w'}(\underline{y})} \qquad (5.2)$$

Recall that for the case study all objective values are calculated for a one hour period in the AM peak. The resulting trade-off values are shown in Table 5.6. For example, on average it is possible to reduce TTT on a daily basis by one hour at a daily additional expense of €6.94. It is likely that more efficient investments than average are possible by selecting specific solutions. Note that the other two objectives also may change values and result in additional benefits or costs, which are disregarded in this table.

Table 5.6: Average trade-off values between the four objectives ($ATO_{w,w'}(P)$)

$w \setminus w'$	TTT (hours)	USU (# of cars)	OpD (euros)	CE (kilos)
TTT (hours)		-2.4	-0.14	-0.23
USU (# of cars)	-0.41		-0.067	-0.066
OpD (euros)	-6.9	-15		-1.1
CE (kilos)	-4.3	-15	-0.89	

It is checked whether spending additional money on the multimodal transportation network is cost-efficient, i.e. whether the benefits are larger than the costs. The average trade-off values found in the optimisation problem are benchmarked with trade-off values derived from values to monetise traffic externalities in literature (see Table 5.7). The value of time (VoT) in the Netherlands (for all purposes) ranges from €6.75 for bus, tram and metro to €9 for car and €9.25 for train (Significance et al., 2013). In the benchmark an average value of €8.50 is used. Schroten et al. (2014) identify €78 per ton CO_2 as median value for the Dutch situation, which is based on Kuik et al. (2009). Last of all, one parking space on average has a yearly rent of €758 with a standard deviation of €504 (Van Ommeren and Wentink, 2012). In the benchmark, a value of €1000 per year is used (a little above average), because it is expected that the average parking rent in highly urban areas is higher than average. When this parking space is not needed any more, the space becomes available for other uses that will have a comparable economic or societal value. Using 250 working days a year, preventing one commuting trip in the AM peak saves €4.- per day on parking costs.

Given the current insights into monetary values for TTT, USU and CE, the values in Table 5.7 are the maximum acceptable deterioration of one objective to gain one unit of another objective. When comparing the values in both tables, the following observations can be made. Spending additional OpD to reduce TTT is cost-efficient on average, since the cost is lower than the known VoT. On the other hand, spending OpD to reduce USU and CE is not cost-efficient on average: the trade-off values are higher than the current monetisation values. Note that this trade-off value is an average: in some areas of the Pareto front cost-efficient

measures may still exist for USU and CE. When comparing CE and TTT and following the current monetary values, 109 kilos of additional CE may be accepted for a one-hour gain in TTT, compared to only 4.31 kilos that are needed on average for a one-hour gain. This implies that in the multimodal NDP additional CE are to be accepted to a large extent to gain TTT. When comparing USU and TTT the situation is similar.

Table 5.7: Trade-off values derived from monetisation values from literature

$w \setminus w'$	TTT (hours)	USU (# of cars)	OpD (euros)	CE (kilos)
TTT (hours)		-0.47	-0.12	-0.0092
USU (# of cars)	-2.1		-0.25	-0.020
OpD (euros)	-8.5	-4.0		-0.078
CE (kilos)	-109	-51	-13	

Table 5.8: Costs to reduce (transportation related) CE when using other types of measures and monetary values for CE

Means to reduce CE	**Cost to reduce CE**	**Reference**
Evaluation of policy measures in the Netherlands		
Tighter CO_2 emission standard for cars after 2020: from 95 to 70 g/km.	110 euro per ton CO_2	Daniëls et al. (2014)
Stimulate electric cars (zero-emission vehicles) by for example tax benefits and CO_2 emission standards	350 euro per ton CO_2	Daniëls et al. (2014)
World-wide analysis for long-term CE reduction		
The BLUE Map scenario of the International Energy Agency: halve 2005 emission levels by 2050 by exploiting technology options	USD 175 per ton CO_2	IEA (2010).
As a part of the BLUE map scenario: large-scale introduction of Plug in hybrid electric vehicles	USD 140-210 per ton CO_2	IEA (2010)
Currently known prices for CE		
European Union Emissions Trading System, price of CO_2 emission right	€4.50 / ton	Hekkenberg and Verdonk (2014)
Price of external effects of CE	€78 / ton	Schroten et al. (2014)

Finally, the value for costs to reduce CE found by taking measures in the multimodal transportation network is compared to costs of other measures taken to reduce CE (Table 5.8). All these measures are more cost-efficient in reducing CE than the value found (which is more than 1,100 euros per ton CO_2). However, these other types of measures also have higher costs than the current price of external effects of CE used in cost-benefit analysis. The price of CO_2 in the European emissions trading system is even much lower than that. To conclude, it should be noted that although measures in the multimodal transportation network are not cost-efficient when only looking at CE, they may still be so when (positive) effects on the other objectives USU and TTT are taken into account. Furthermore, in some areas of the Pareto front measures related to multimodal trip making may be more cost-efficient than other policy measures that were taken recently in the Netherlands to reduce CE. Recall that it is even possible to reduce all four objectives concurrently compared with the base solution, achieving both cost reductions and reductions in CE.

Trade-offs between more than two objectives

As a next step information of all four objectives is taken into account to determine trade-off information. First, the three objectives that are rather in line with each other (TTT, USU and CE) are combined to one value TUC, using weights. As values for these weights, the average trade-off values $ATO_{w,w'}(P)$ are used (implying that the two extreme solutions for a pair of objectives are equally valued when combining those two objectives). The combined value for the three objectives is expressed in hours as generalised travel time (see Eq. 5.3).

$$TUC = TTT + ATO_{TTT,USU}(P)*USU + ATO_{TTT,CE}(P)*CE \qquad\qquad (5.3)$$

For each solution the value of TUC is determined. Together with OpD, only 2 objectives remain. The values of TUC and OpD for all 210 solutions in the Pareto set are plotted in Figure 5.9. These two objectives are clearly opposed: a clear negative relation can be observed in the figure. When only TUC and OpD are considered, 36 solutions are Pareto-optimal and most other solutions are near-optimal (i.e. close to a solution in the Pareto set).

Using the new objective TUC the average trade-off value $ATO_{TUC,OpD}(P)$ is equal to -0.37 hours / euro, meaning that on average for one euro of additional OpD, 0.37 hours of generalised travel time is gained. When comparing the optimal solution for TUC with the optimal solution for OpD, it appears that the scores for all three original objectives (TTT, USU and CE) are improved at the cost of additional OpD. This means that on average for one euro of additional OpD, concurrently approximately 0.1 hours of TTT, 0.07 parking spaces of USU and 0.5 kilos of CE are gained.
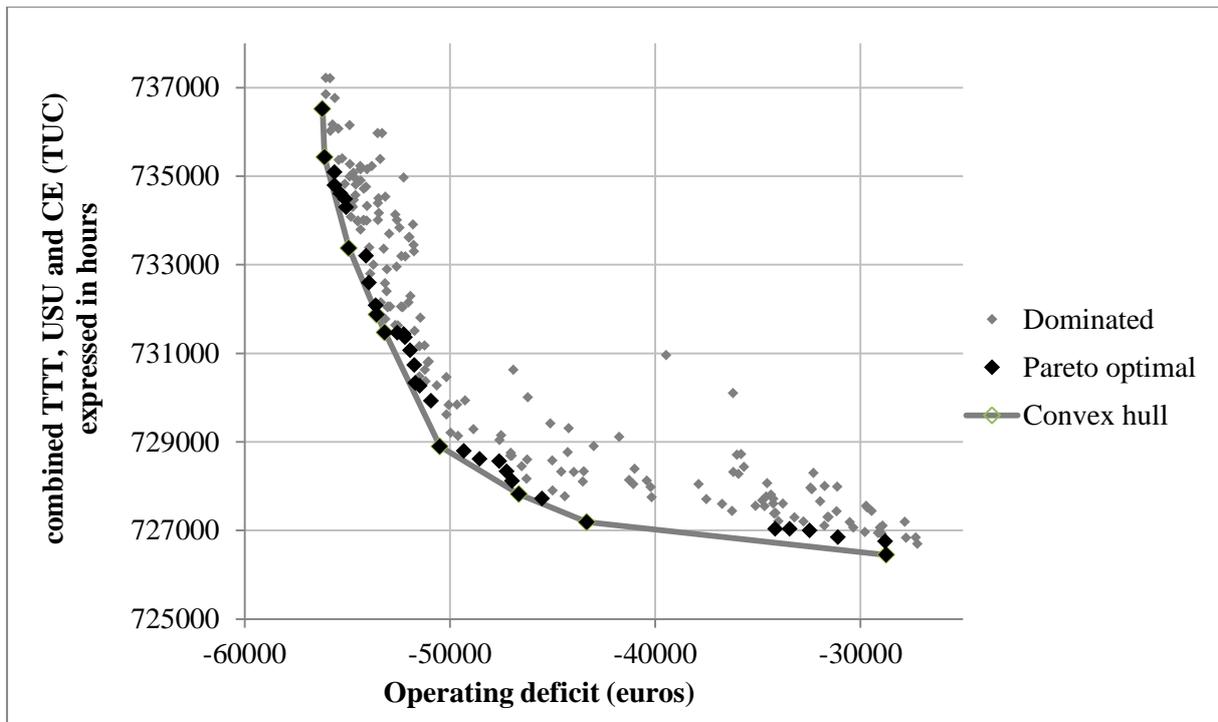
Figure 5.9: Trade-off between the combined objective TUC and OpD

A next step is to determine trade-off values at specific areas of the Pareto front instead of an average value based on the extreme ends of the set. The 36 Pareto-optimal solutions are numbered starting from the lower right end of the Pareto front, i.e. starting with the optimal solution for TUC. When assuming that within the decision making process the objectives are eventually linearly weighted (e.g. using the weighted sum method), some of the Pareto-optimal solutions are irrelevant and never chosen as the final design. The only relevant solutions are vertices of the convex hull (see Figure 5.9). The convex hull has 9 solutions as its vertices. For each pair of neighbouring vertices a trade-off value is identified, resulting in 8 different values (see Table 5.9). Starting from solution 36 (that is optimal for OpD), the first improvements for TUC are very cost-efficient compared to the average trade-off value: more than 10 hours of generalised travel time can be saved by investing 1 additional euro. Because TUC is composed of three other objectives using weights, this does not necessarily mean that all three objectives improve. In the first step this is the case: TUC consists of a large gain in USU and CE and a loss in TTT. When moving from Pareto solution 35 to 29, again the largest gains are in CE and USU, but now also TTT improves. In the next step from Pareto solution 29 to 25 the share of TTT in the gain in TUC is much larger, but still the other objectives improve as well. Moving further along the front, the gain in TUC that is achieved by an additional euro of OpD becomes smaller and smaller. The contribution of TTT, USU and CE to the gain in TUC strongly varies. In the final and most expensive step (in terms of OpD) from Pareto solution 7 to solution 1 CE becomes worse, but including the gain in TTT and USU, there is still a little gain in TUC.

Table 5.9: Trade-off values for pairs of vertices of the convex hull

Pareto solution indices i,i'	$TO_{TUC,OpD}(\underline{y}_i, \underline{y}_{i'})$ (generalised hours / euro)	$TO_{TTT,OpD}(\underline{y}_i, \underline{y}_{i'})$ (hours/euro)	$TO_{USU,OpD}(\underline{y}_i, \underline{y}_{i'})$ (#/euro)	$TO_{CE,OpD}(\underline{y}_i, \underline{y}_{i'})$ (kilos/euro)
1-7	-0.05	-0.030	-0.016	0.08
7-9	-0.19	-0.153	-0.005	-0.11
9-15	-0.28	-0.035	-0.054	-0.48
15-24	-0.96	-0.367	-0.141	-1.08
24-25	-1.02	-0.142	-0.019	-3.56
25-29	-1.10	-0.402	-0.177	-1.16
29-35	-1.73	-0.135	-0.472	-1.92
35-36	-11.34	1.359	-2.066	-33.03
Average (1-36)	-0.37	-0.098	-0.067	-0.45

5.3.3. Relation between decision variable values and objective values

When improving the transportation network, measures have to be chosen for implementation. To motivate these decisions, insight into the influence of each measure on the objective values is needed. Certain measures will be beneficial for one objective, other measures for another objective. This includes possible interdependencies between measures, for example when measures strengthen or weaken each other for an objective. Therefore in this section first a sensitivity analysis is conducted, providing insight into the effect of each measure on its own. Thereafter, correlations between values of decision variables and objective values and between pairs of decision variable types are analysed. Finally, for each measure the effectiveness in attaining Pareto-optimal solutions is shown.

Sensitivity analysis

The isolated effect of each of the 37 variables on the objective values is determined by a sensitivity analysis. To this end, the partial derivative is used. The exact partial derivatives cannot be determined, because the decision variables are discrete and because the objective functions are not known analytically (the objective values are calculated numerically based on the outcome of the lower-level model). Therefore, the slope is used instead, as defined in Eq. 5.4. SL_{w,v,∇_v}, also known as average rate of change, is the amount objective w changes per unit of change in decision variable v, using step size ∇_v and starting from a base solution $\underline{y}^0$ (see Figure 5.10).

Table 5.10: Value of decision variables in base solution

Decision variable type	Decision variable	Decision variable name	Real value in base solution
Opening / closure of train station	y_1	Halfweg-Zwanenburg	1
	y_2	Haarlem Zuid	0
	y_3	Amsterdam Geuzenveld	0
	y_4	Amsterdam Nieuw Sloten direction East	0
	y_5	Amsterdam Nieuw Sloten direction North	0
	y_6	Amsterdam Westerpark	0
Express train status of train station	y_7	Hoofddorp	0
	y_8	Duivendrecht	1
	y_9	Heemstede-Aardenhout	1
Opening / closure of P&R facility	y_{10}	Train station Halfweg-Zwanenburg	0
	y_{11}	Velsen South	0
	y_{12}	Amsterdam, Zuiderzeeweg	0
	y_{13}	Train station Geuzenveld	0
	y_{14}	Amsterdam, Buikslotermeerplein	0
	y_{15}	Amstelveen, Oranjebaan	0
	y_{16}	Schiphol North	0
Frequency of bus line	y_{17}	Haarlem - Amsterdam South	8
	y_{18}	Haarlem - Amsterdam Bijlmer	4
	y_{19}	IJmuiden - Amsterdam Sloterdijk	2
	y_{20}	Zandvoort - Amsterdam Marnixstraat	2
	y_{21}	Zaanse Schans - Amsterdam Centraal	8
	y_{22}	Zaandam - Amsterdam Centraal	8
	y_{23}	Schiphol - Amsterdam Sloterdijk	0
	y_{24}	Amsterdam Bijlmer – Ijburg	8
Frequency of local train line	y_{25}	Uitgeest - Amsterdam Centraal	2
	y_{26}	Uitgeest - Schiphol - Hoofddorp	0
	y_{27}	Uitgeest - Zaandam - Zuid - Bijlmer	0
	y_{28}	Zandvoort - Haarlem - Zuid - Bijlmer	0
	y_{29}	Uitgeest - Haarlem - Zuid - Bijlmer	0
	y_{30}	Zandvoort - Haarlem - Centraal	2
	y_{31}	Uitgeest - Haarlem – Centraal	2
	y_{32}	Amsterdam Centraal - Almere Oostvaarders	2
	y_{33}	Hoofddorp - Almere Oostvaarders	4
	y_{34}	Bijlmer - Almere Oostvaarders	0
	y_{35}	Utrecht - Hilversum - Amsterdam Centraal	4
	y_{36}	Utrecht - Hilversum - Hoofddorp	2
Extension of a tram line	y_{37}	Amsterdam Geuzenveld	0

$$SL_{w,v,\nabla_v} = \frac{Z_w(\underline{y}^0 + [0,\cdots,0,\nabla_v,0,\cdots,0]) - Z_w(\underline{y}^0)}{\nabla_v} \qquad (5.4)$$
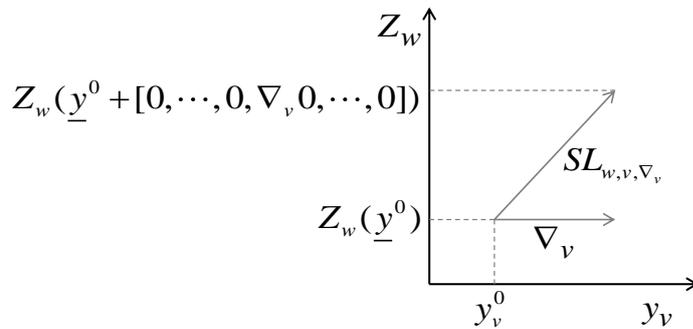
Figure 5.10: Slope SL_{w,v,∇_v}

Although interdependencies between decision variables are neglected in this way, this gives an impression of which variables have a strong effect on the objectives and which have a weak influence, and whether this influence is positive or negative. The same base solution $\underline{y}^0$ is used as in Section 5.3.1: the situation where the current plans for 2030 are realised. As can be seen in Table 5.10, most decision variables have value 0 in this solution, but some of them also have a nonzero value, implying that these measures are included in the current plans.

In the sensitivity analysis each decision variable is changed compared to the base solution, where for most decision variables this means an additional measure is realised (i.e. changing the value of the decision variable from 0 to a nonzero value by a positive value ∇_v). However, some variables are nonzero in the base situation and are therefore changed to 0 (using a negative value ∇_v). This means that a measure is reversed: a PT service line is cancelled or in one case an existing train station is not used any more. A correct interpretation of the effect of these variables is assured by the definition ∇_v, which has a negative value in that case. 48 solutions are evaluated: the base solution and 47 solutions with one decision variable changed. If more than two values are allowed for a non-binary discrete variable, this variable is changed twice, using two different step sizes ∇_v: a small change ∇_v^S and a large change ∇_v^L.

Since the objectives have different units, the relative change is presented here (see Eq. 5.5). In the left half of Table 5.11 the minimum and maximum values for relative change $\min_{v,\nabla_v} RC_{w,v,\nabla_v}$ and $\max_{v,\nabla_v} RC_{w,v,\nabla_v}$ are shown. Just like in Section 5.3.1, it can be observed that operating deficit has the largest range of relative change, because operating costs are influenced directly by the decisions. The measures mainly cause an increase of the deficit because they involve additional costs, but in some cases the deficit is reduced, implying that the revenues of PT increase more than they cost. Another observation is that TTT and USU can be reduced to a larger extent than they can be increased: the types of measures are mainly improvements in the PT network, with lower travel times and less car use as a result. Note that the measures that deteriorate the PT network are analysed by redirecting the effect, i.e. the effect is divided by a negative step size ∇_v. Finally CE are the most difficult to influence by these types of measures, because attracting car users to the PT system (and therefore reducing CE) can only be achieved by putting additional PT vehicles into service (and therefore increasing CE). The 4 objectives largely vary in their range of relative changes that occur, so to visualise the 4 objectives in one figure, the relative changes in objective values are normalised in such a way

that each range is rescaled to 100% using Eq. 5.6. (see right half of Table 5.11; the rescaled values are used in Figures 5.11 and 5.12).

$$RC_{w,v,\nabla_v} = \frac{SL_{w,v,\nabla_v}}{Z_w(\underline{y}^0)} \tag{5.5}$$

$$NRC_{w,v,\nabla_v} = \frac{RC_{w,v,\nabla_v}}{\max\limits_{v,\nabla_v} RC_{w,v,\nabla_v} - \min\limits_{v,\nabla_v} RC_{w,v,\nabla_v}} \tag{5.6}$$

Table 5.11: Minimum and maximum relative change that occurred for TTT, USU, OpD and CE, after changing only one decision variable during the sensitivity analysis

	Relative change RC_{w,v,∇_v}				Normalised relative change NRC_{w,v,∇_v}			
	TTT	USU	OpD	CE	TTT	USU	OpD	CE
Minimum	-0.50%	-0.51%	-0.28%	-0.14%	-77.3%	-81.2%	-14.5%	-54.8%
Maximum	0.15%	0.12%	1.67%	0.12%	22.7%	18.8%	85.5%	45.2%
Range	0.65%	0.63%	1.95%	0.26%	100.0%	100.0%	100.0%	100.0%
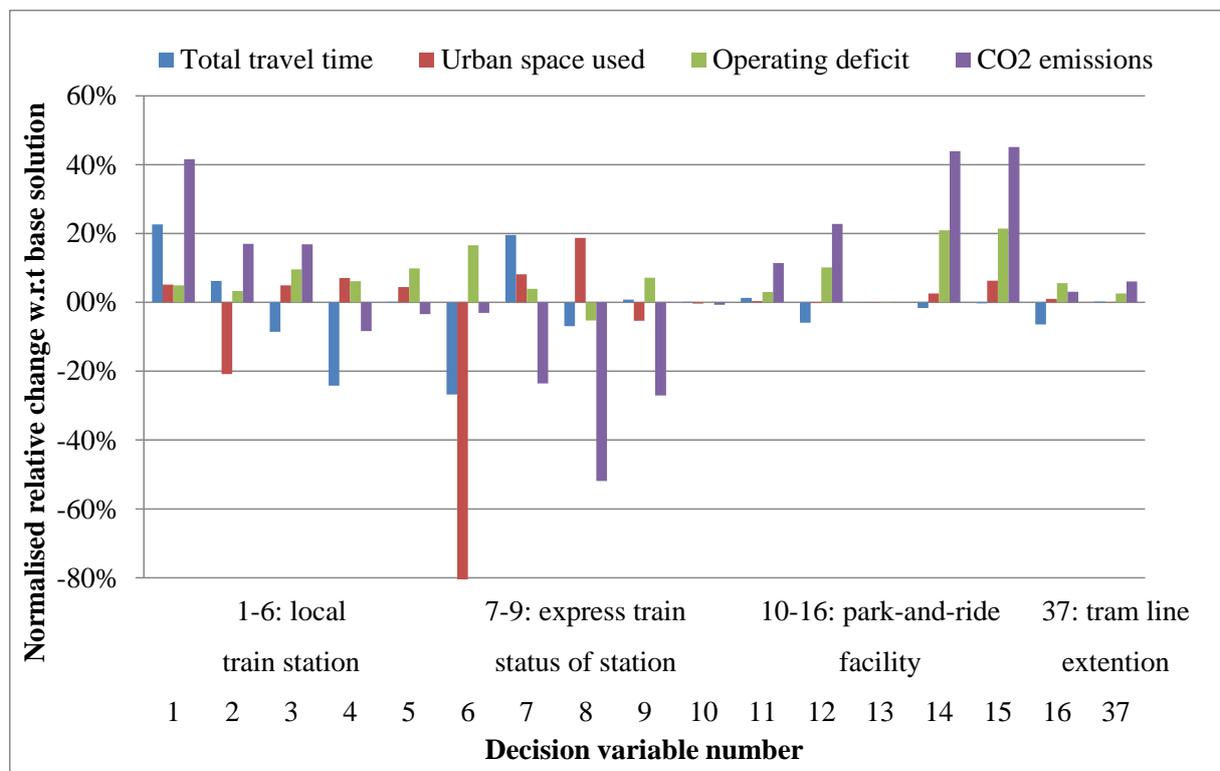
Figure 5.11: NRC_{w,v,∇_v} **for binary decision variables**
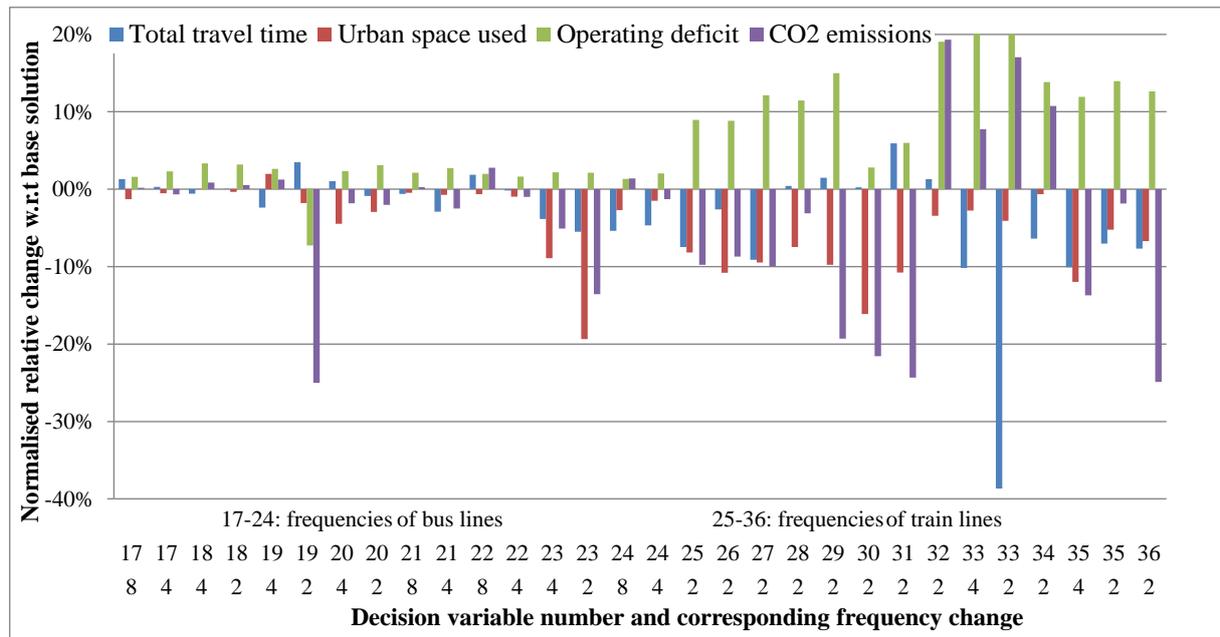
Figure 5.12: NRC_{w,v,∇_v} **for non-binary decision variables**

The results of the sensitivity analysis are plotted with a distinction between binary variables (Figure 5.11) and variables that can take 4 discrete values (Figure 5.12). Although there is large variation in effects of decision variables, some trends can be identified based on these figures:

- A change with respect to the base situation for most decision variables leads to a decrease in some objectives, but to an increase other objectives, typically illustrating trade-offs between objectives. For example, increasing the frequency of the train line between Uitgeest and Amsterdam Centraal (variable 25) leads to lower values for TTT, USU and CE, but to increased OpD. 3 variables (numbers 1, 11 and 37) lead to an increase of all four objective values. Because all objectives are to be minimised, the individual change of these three decision variables each leads to a solution that is dominated by the base solution, for various reasons. Variable 11 represents the P&R facility at Velsen South, which has a negative effect given the low frequency of the related bus line (variable 19) in the base situation. Note that in the Pareto set this P&R is often included and the related bus line has high frequencies (4 or 6 times per hour). Also variable 1 (representing train station Halfweg-Zwanenburg) has a negative effect given the frequencies of the related train services in the base solution. Variable 37 represents a tram line extension to Amsterdam Geuzenveld, which only makes sense in combination with a new train station at this location, which is not included in the base situation.
- Frequencies of bus lines have a smaller influence than frequencies of local train lines. This makes sense, because train lines have larger a passenger capacity and are more expensive to operate.
- Most decision variables result in lower TTT, because the decision variables in principle cause an increase of the PT service network quality (here also note that the measures that deteriorate the PT network are analysed by redirecting the effect, i.e. the

effect is divided by a negative step size ∇_v). In some cases TTT is higher, which may be due to additional stations of PT lines, where the travel time loss of through passengers is larger than the travel time gain of the users of the opened train station. Another mechanism is that improving the quality of PT leads to a shift from car to PT, which may imply a higher travel time for these travellers.

- One specific decision variable (number 6, the opening of Amsterdam Westerpark train station) strongly reduces USU. This may be explained by the location of the station, close to the centre of Amsterdam, in the vicinity of a large number of highly urban zones. Most other variables reduce USU as well, which makes sense because in general the PT network is improved. However, opening some train stations and turning some stations into express train stations results in higher USU. This may be explained by the effect on through passengers, who mainly travel between highly urban areas, and now, because of the higher travel time, switch to the car to make their trip.

- Most decision variables result in a larger OpD, because the decision variables all cause an increase of the PT service network quality, which comes at a cost. Some measures generate more revenues than their marginal costs, resulting in a decrease of operating deficit (for example variable 8, where the train station of Duivendrecht is changed into an express train station, and variable 19, where a low frequency bus service is introduced between IJmuiden and Amsterdam).

- The effect on CE varies strongly among decision variables.
  - It can be observed that turning train stations into express train stations is beneficial for CE, because the marginal emissions of the service network are very small (no additional vehicle kilometres are made), while additional car passengers are attracted to PT. Apparently the burden for through passengers in express trains is smaller than the benefit (in terms of CE, not for all 3 variables in terms of TTT).
  - Opening P&R facilities leads to more CE, because former PT trips (possibly with bicycle as access or egress mode) are replaced by P&R trips. This result was also found by Mingardo (2013), who used the findings of a user survey conducted in nine rail-based P&Rs located around the cities of Rotterdam and Den Haag (The Netherlands) to identify intended and unintended effects of introducing P&R.
  - A small frequency increase on most PT lines is beneficial for CE, because this leads to fewer car kilometres and only limited additional CE by PT vehicles. A larger increase in frequency is not beneficial, because the additional CE by the PT vehicles outweighs the relatively low number of attracted car passengers.

- Variable 13 (P&R facility at the train station of Geuzenveld), has no influence on its own, because in the base network this station is not served by PT and therefore a P&R facility makes no sense.

Correlation between decision variable types and objectives

To identify relations between decision variables and objectives, the values of decision variables are aggregated based on their type (train station, express train station, P&R facility, train frequency or bus frequency). For each solution in the Pareto set the values of the decision variables within a type are summed up, resulting in a representation of a solution containing 5 values (1 value per type) to represent the decision variable values of a solution.

For example, if a solution contains three additional train stations, the new decision variable for train stations equals 3. Note that this representation neglects the fact that some decision variables represent larger measures than other decision variables (for example the train between Zandvoort and Amsterdam Centraal has a shorter route than the train from Uitgeest to Amsterdam Bijlmer).

A correlation matrix is made that has the four objectives in its rows and the aggregated values of types of decision variables in its columns (see Table 5.12), based on all solutions in the Pareto set. The train frequencies have the strongest correlation with the objectives, where TTT, USU and CE are reduced by increasing train frequencies and OpD is increased. Similar but less strong effects are found for bus frequencies, train stations and P&R facilities. Finally, express train stations show different relations: TTT and USU are increased by opening express train stations, and OpD is decreased. This is an indirect effect, which can be explained by the relation with (local) train frequencies: when frequencies of local trains are low (to save OpD), a cost-efficient way to still serve the local stations is to make them into an express train station. In Table 5.13 this is illustrated by the correlation between pairs of values per type of decision variable: the values for express train stations have a negative correlation with all other types. The other types all have a positive correlation with each other. Especially bus frequencies and train frequencies are positively related, indicating that the bus and train routes defined as decision variables are rather complementary than competitive.

Table 5.12: Correlation between values per type of decision variable and objective values

	TTT	USU	OpD	CE
Train stations	-0.50	-0.63	0.57	-0.41
Express train stations	0.32	0.23	-0.30	-0.04
P&R facilities	-0.43	-0.36	0.47	-0.16
Bus frequencies	-0.71	-0.80	0.77	-0.57
Train frequencies	-0.85	-0.93	0.96	-0.78

Table 5.13: Correlation between pairs of values per type of decision variable

	Train stations	Express train stations	P&R facilities	Bus frequencies	Train frequencies
Train stations	1.00	-0.12	0.25	0.52	0.55
Express train stations	-0.12	1.00	-0.19	-0.21	-0.25
P&R facilities	0.25	-0.19	1.00	0.30	0.42
Bus frequencies	0.52	-0.21	0.30	1.00	0.73
Train frequencies	0.55	-0.25	0.42	0.73	1.00

As an example the relation between the total train frequency and USU is shown in Figure 5.13. Each dot in the scatter plot represents one solution, showing the total frequency in that solution (i.e. the sum of the real values of decision variables 25-36) and the corresponding

value for USU. Note that not all PT lines in the network are included as decision variables, and therefore still PT lines are running when the value for total train frequency is equal to 0. The largely negative correlation is confirmed by the plot, where it also becomes clear that increasing the frequency when it is low has a stronger effect on USU than increasing the frequency when it is already high: the marginal benefits of frequency increases become smaller.
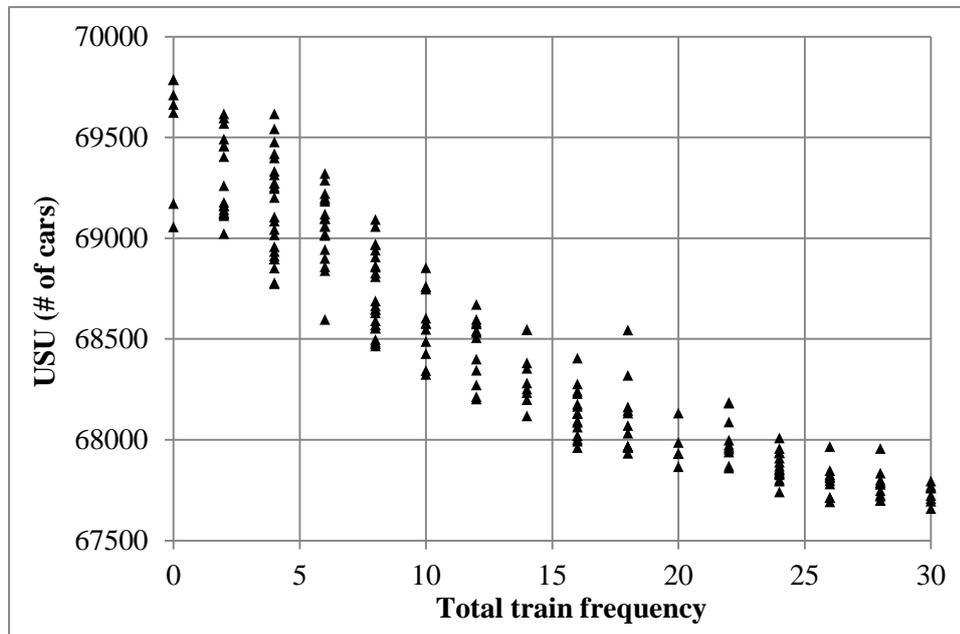
Figure 5.13: Relation between the total train frequency and USU

Percentage of active decision variables in Pareto set

In Figure 5.14 the fraction of nonzero values for decision variables in the Pareto set is shown. Since the Pareto set contains only non-dominated solutions (and therefore all inferior solutions are filtered out), these fractions are an indicator for the effectiveness of each decision variable. In the extreme case that a variable is nonzero in all Pareto solutions, the related measure should always be implemented, regardless which (combination) of the 4 objectives the decision maker finds most important (so called 'no regret' measures). In the other extreme case where a variable is zero in all Pareto solutions, it is sure that this measure should never be implemented (so called 'always regret' measures).

Variable 5 (train station at Nieuw Sloten in Northern direction) does not occur at all in the Pareto set and can therefore be classified as an 'always regret' decision. Some more variables (8 and 9: express train stations at Duivendrecht and Heemstede-Aardenhout and 10, 13, 14 and 15: P&R facilities at Halfweg-Zwanenburg train station, at Amsterdam Geuzenveld train station, at Buikslotermeerplein in Amsterdam North and at Oranjebaan in Amstelveen) occur in only a very small fraction of the Pareto solutions (less than 10%). This means that for most combinations of preferences for objectives these measures should not be included in the final network design. However, specific combinations of preferences for objectives exist that make solutions containing these measures optimal, because the measures occur in the Pareto set.
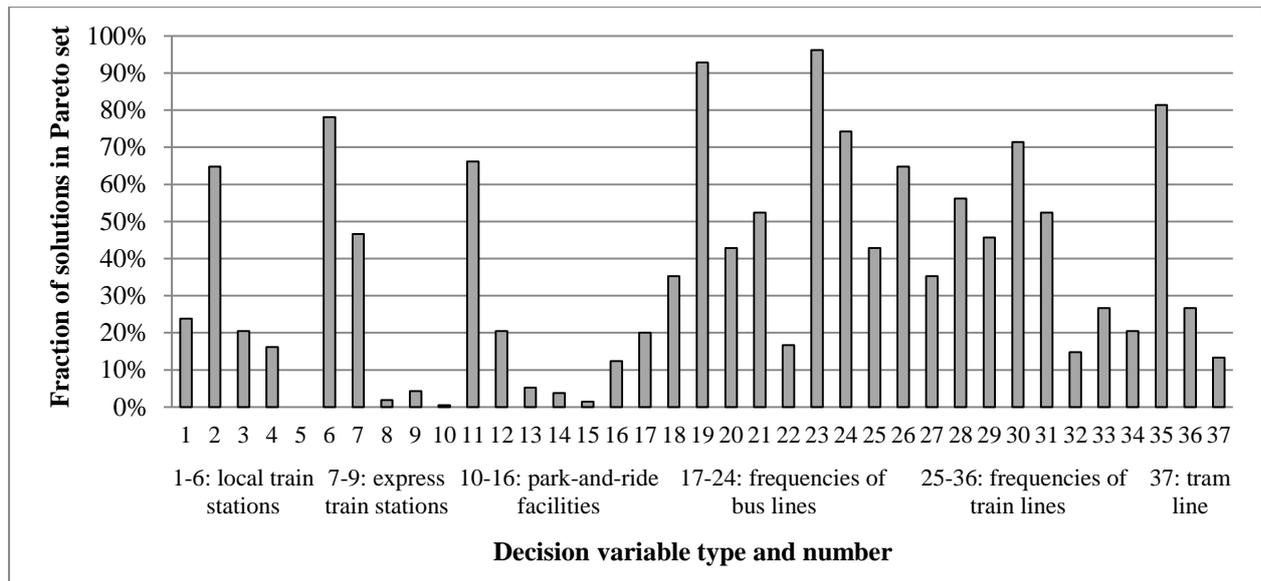
Figure 5.14: Effectiveness of decision variables to reach Pareto-optimality: fraction of solutions in the Pareto set that have a nonzero value for each decision variable

On the other hand, no measures are present in all Pareto solutions, so 'no regret' measures do not exist within the 37 measures investigated. Two measures (variables 19 and 23, representing bus lines between IJmuiden and Amsterdam Sloterdijk and between Schiphol and Amsterdam Sloterdijk) have scores higher than 90%, which could be called almost 'no regret': for most combinations of preferences for objectives these measures should be included in the final network design (but not for all combinations of preferences).

When a distinction is made between types of objectives, P&R facilities in general have low values. The facility associated with variable 11 (Velsen-South) is an exception: it is located directly at a motorway junction and provides a good PT connection to the city centre of Amsterdam. The other P&R facilities with low values are located relatively close to the city centre of Amsterdam: reaching these locations by car involves too much congestion to be an attractive alternative. The new train stations Haarlem South (variable 2) and Amsterdam Westerpark (variable 6) have relatively good scores within the category of train stations. Both stations are located near the city centres of their respective cities. Nonzero frequencies of bus and train lines are in general quite effective to include in Pareto solutions, but there are large differences among individual variables.

A large majority of the variables has a value roughly around 50% and therefore it depends on the objectives that are preferred whether the corresponding measure should be taken or not (and therefore which solution is chosen as a final solution). Therefore, the next step is to make a distinction between solutions based on their objective values.

Percentage of active decision variables in two distinct clusters

In the previous section, decision variable values were analysed for the entire Pareto set. When a certain area of the Pareto front is analysed (i.e. certain ranges for objective values), it is expected that it will become more clear which decision variable values in general correspond with a certain range for objective values. When looking for example at the scatter plot of the

Pareto set for TTT and CE (Figure 5.15), two clusters of solutions can be identified based on their combined TTT and CE values: roughly speaking cluster 1 has a lower value for TTT than average (combined with relatively low values for CE) and cluster 2 has a higher value for TTT than average (combined with relatively high values for CE). The question arises what (type of) measures occur in the two clusters.
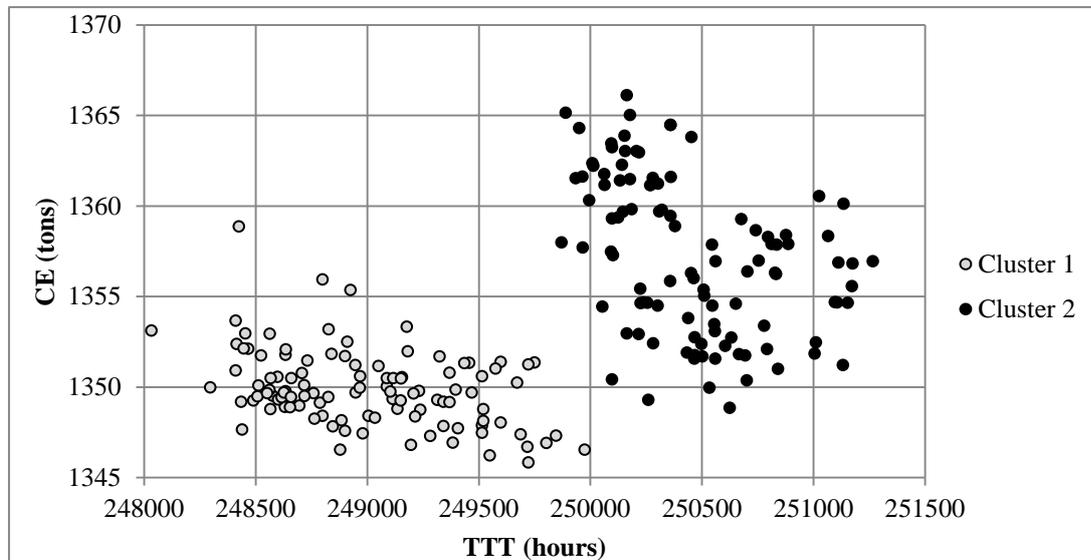
Figure 5.15: Scatter plot for TTT and CE with distinction between two clusters of solutions

In Figure 5.16 the fractions of solutions that have a nonzero value for each decision variable are given per cluster. Generally speaking, cluster 1 contains many more solutions with nonzero decision variable values than cluster 2. This is in line with the expectation, because taking more measures in the multimodal transportation network (i.e. nonzero decision variables) in general has a positive influence on both TTT and CE. For 5 decision variables this is the other way around: in cluster 2 variables 3, 7, 8, 11 and 19 are nonzero more often than in cluster 1. Variable 7 (express train station in Hoofddorp) and variable 11 (P&R facility in Velsen South) have a relatively large difference. For the express train station in Hoofddorp there is a relation with the local trains: in cluster 2 the frequencies of local trains (variables 25-36) are mainly low, making an express station at Hoofddorp more interesting. This does not hold for the other two express train stations (variables 8 and 9), because these locations are not beneficial in the entire Pareto set (very low fractions of nonzero values). The P&R facility at Velsen South (variable 11) is also more present in cluster 2 than in cluster 1. Although it would be expected that this new P&R facility is beneficial for TTT, it is apparently more beneficial for USU (and for OpD in case it attracts many passengers for the related bus line). This P&R facility is located at the bus line IJmuiden – Sloterdijk (variable 19), which has a nonzero value in almost all solutions in both clusters.

When looking at types of decision variables, frequencies of train lines strongly vary among the clusters. This indicates that train frequencies strongly influence TTT and CE. It is striking that for all trains to Almere (variables 32-34) the frequencies are set to 0 for all solutions in cluster 2. This can be explained by the fact that between Amsterdam and Almere alternative travel options exist that are not included as decision variables in the case study. These include

other local trains, buses and express trains. Note that in reality capacity is likely to become a problem when all these three decision variables are set to 0, but this is not included in the lower-level model. For the other variable types, the difference among the two clusters varies strongly over individual measures.
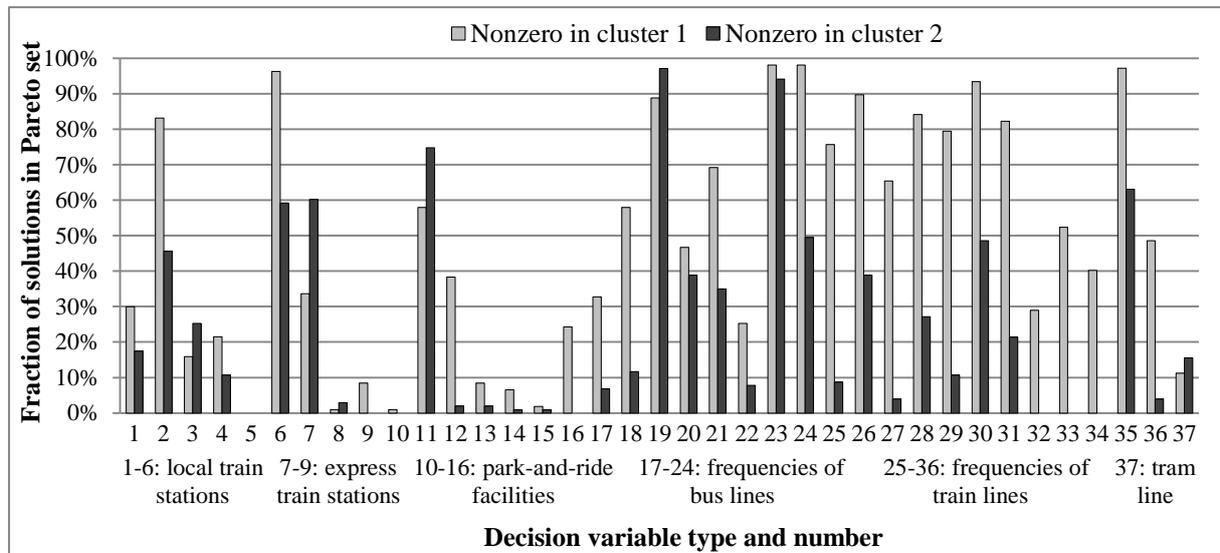
Figure 5.16: Fraction of solutions that have a nonzero value for each decision variable per cluster

5.3.4. Conclusion

The gain per objective with respect to the base solution varies strongly over the objectives. The optimal solutions per objective show that CO_2 emissions (CE) can be improved by only 0.5%, while operating deficit (OpD) can be improved by more than 6%. TTT can be improved by 0.8% and USU by more than 1.5%. These relative impacts should be considered with care, because the gains are relative to the total values for the entire study area. In absolute terms these possible gains are considerable: every AM peak almost 4000 hours of travel time, more than 2000 parked cars and 12 tons of CE (equivalent to the daily CE of more than 500 Dutch households) can be saved. In monetary terms the travel time gain roughly equals 13 million euros per year, assuming 200 working days per year, a value of time of 8 euros / hour and a comparable impact during the PM peak as during the AM peak.

For three out of the four solutions that are optimal for one single objective, another objective value becomes worse than its value in the base solution. Anyhow it appears to be possible to reduce the value of all four objectives concurrently compared to the base solution: 47 out of 210 Pareto-optimal solutions have better values for all objectives (i.e. dominate the base solution). One of those solutions has a concurrent gain for all objectives of at least 0.4%.

From the four objectives, total travel time (TTT), urban space used by parking (USU) and CE are more or less in line with each other and all three are opposed to OpD. When comparing trade-off values between additional monetary expenses (represented by OpD) and the other three objectives, it appears that cost-efficient travel time gains are made possible by the measures studied. A considerable area of the Pareto front has a trade-off value that is lower

than the value of time used in cost-benefit analysis. The price of one hour of travel time gain is lower than the benefit of it, so from the viewpoint of the regional government cost-efficient measures exist within the selection of measures considered. It should be noted that investments in infrastructure (paid by the national government) are not included in the analysis, so whether the measures are cost-efficient from a societal point of view is unsure. The average trade-off value between USU and OpD is a factor 4 higher than the benefit of it, so cost-efficient trade-offs may only occur in a limited part of the Pareto front. Finally the average trade-off value between CE and OpD is more than 10 times higher than the value known from cost-benefit analysis, so using the current valuation of CO_2, additional expenses for measures in the multimodal network are rarely cost-efficient, if any can be found at all. Note that other examples of policy measures to reduce CE from transportation (stimulating electric cars and setting tighter emissions standards for cars) are not cost-efficient either using the current valuation of CO_2.

It is shown that it is possible to combine TTT, USU and CE to calculate trade-off values between OpD and the combination of the other three objectives: for one euro of additional OpD, on average concurrently 0.1 hours of TTT, 0.07 parking spaces of USU and 0.5 kilos of CE are gained. On the other hand, no solution is optimal for more than one objective at the same time, implying that between all pairs of objectives trade-offs exist.

When looking at types of measures, it appears that train frequencies have the largest influence on the objectives, followed by bus frequencies, train stations and P&R facilities. Express train stations play a special role: they are used as a compensation measure in the solutions with low (local) train frequencies, as a cost-efficient way to still serve train stations (by express trains) when the frequency of local trains is getting lower.

For individual measures, in the case study in the Randstad two 'always regret' measures could be identified and no 'no regret' measures. When distinguishing between two clusters of solutions, one with higher TTT and higher CE and one with lower TTT and lower CE, it appears that within the clusters more 'always regret' and more 'almost never regret' measures exist. Finally, sensitivity analysis has provided insight into the relation of individual decision variables with each of the four objectives. USU can be strongly reduced by opening a new train station at Amsterdam Westerpark. P&R facilities in general reduce TTT, but increase CE, because also former PT travellers are attracted to the P&R alternative.

5.4.   Step-by-step pruning to select a final solution

In this section two different step-by-step reduction procedures are presented to come to a final decision for implementation based on the Pareto set. The first procedure puts additional constraints to objective function values. The second procedure fixes certain decision variable values, i.e. by choosing a measure to implement (for example because it is politically desirable for reasons that are not included in the considered objectives). These two approaches are the result of the three interviews with policy officers that prepare decision making at three different local governments in the Netherlands (municipality of Amsterdam, city region of Amsterdam and province of Overijssel). Note that these two approaches may also be combined to make a selection, but for the sake of simplicity this is not done here.

5.4.1. Using values of objective values

Starting from all solutions in the upper left corner of Figure 5.17, one method (that was suggested by policy officers during the interviews) to gradually reduce the number of solutions in the Pareto set is to put additional constraints to objectives after optimisation. In this example, first an additional constraint to (normalised) OpD is set in a way that only solutions with a value lower than 0.4 are included. As a result 137 solutions of the original 210 solutions remain. One more constraint is put to CE: in addition to the constraint to operating deficit, only solutions with a (normalised) value of lower than 0.2 are included. Only 20 solutions now remain in the selection. As can be seen in the lower left corner of the figure, this selection excludes all solutions with very low values for the other two objectives: putting a bound on the values for CE and OpD implies a bound for the other two objectives as well. A closer look at the objective OpD reveals that by the constraint for CE, the best solutions for OpD are now also excluded. Finally, if the decision maker is satisfied by the values for CE and OpD that are set now and still more than one solution remains, a logical final step is to find the best compromise solution from the remaining solutions considering the other 2 objectives. It is learnt from the interviews that in the political context of decision making each political party values objectives differently. As a result, each party represents a certain objective. When in negotiation, no party will accept a solution in which its objective has a very bad score.

A direct method following this line of thought when searching for the best compromise solution is to select the min-max solution as the preferred solution (see Eq. 5.7, where $\bar{Z}_w$ represents the normalised value for objective w and W_C is the compromise subset of objectives, which may also contain all objectives). In this method, for each solution the least scoring objective is leading when selecting the best compromise solution from all Pareto solutions. Note that the normalisation procedure influences the results: choosing a suitable normalisation procedure is relevant, but not considered here. Since the objective values are rescaled from 0 to 1 using the minimum and maximum values in the Pareto set, the absolute difference in objective values is not explicit here anymore (as it was when calculating trade-off values in Section 5.3.2).

$$BCS_{W_C}(P) = \arg \min_{\underline{y} \in P} \left(\max_{w \in W_C} \bar{Z}_w(\underline{y}) \right) \tag{5.7}$$

In Figure 5.18 the min-max solution is plotted for two different subsets of objectives. First, the min-max solution over all four objectives is chosen (the left plot in the figure). This shows that a compromise solution exists with reasonable scores for all four objectives simultaneously: for all objectives this solution has a score in the best 30% of the range covered, i.e. the lower end of the range. Second, the min-max solution for OpD and CE is chosen (the right plot in the figure). This shows that, although OpD and CE are mainly opposed, low values for both objectives are possible simultaneously. However, this comes with a price: especially TTT scores much worse when focussing only on OpD and CE.
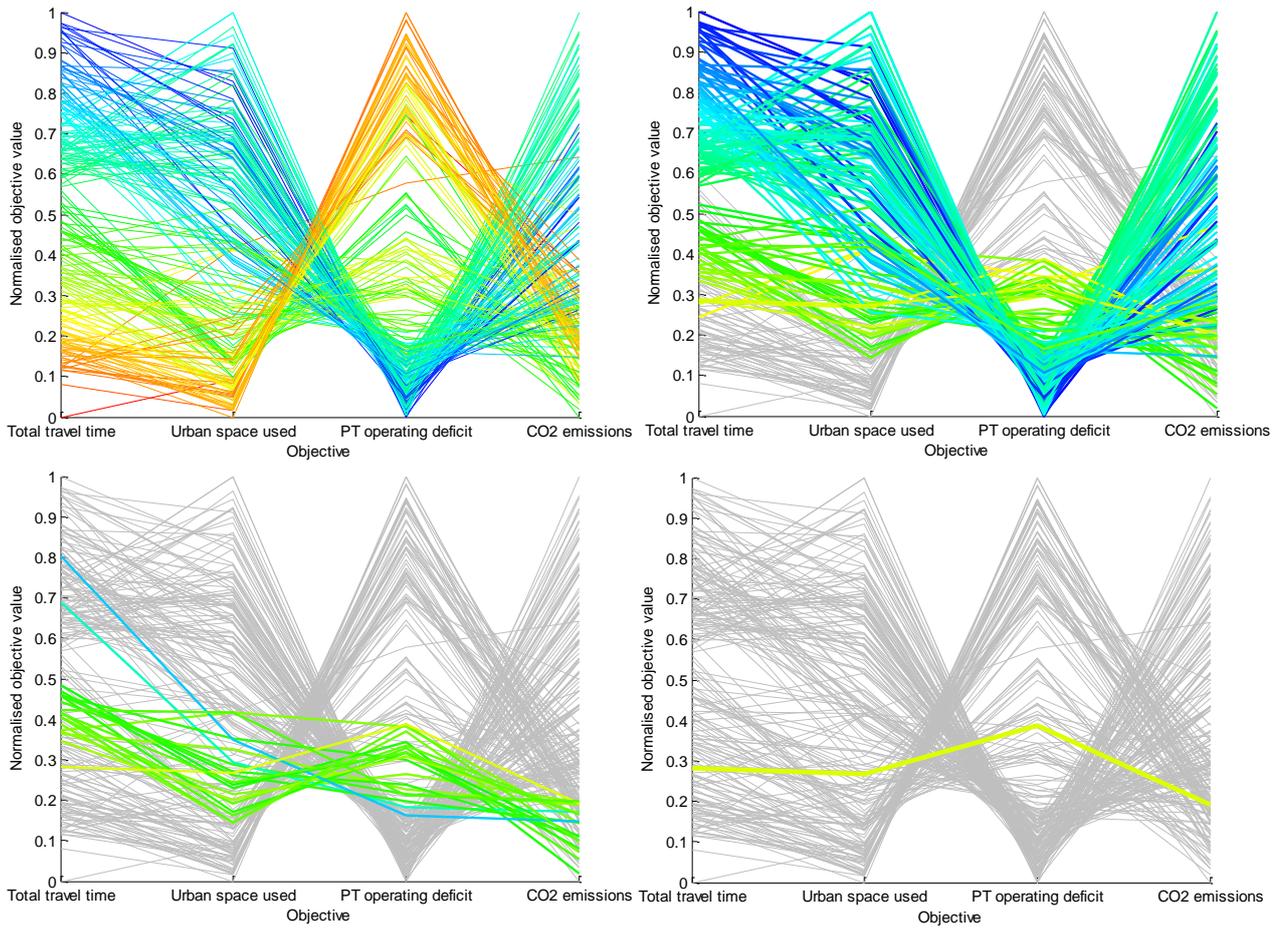
Figure 5.17: Step-by-step reduction from all Pareto solutions to one solution to be implemented by setting bounds to objective values
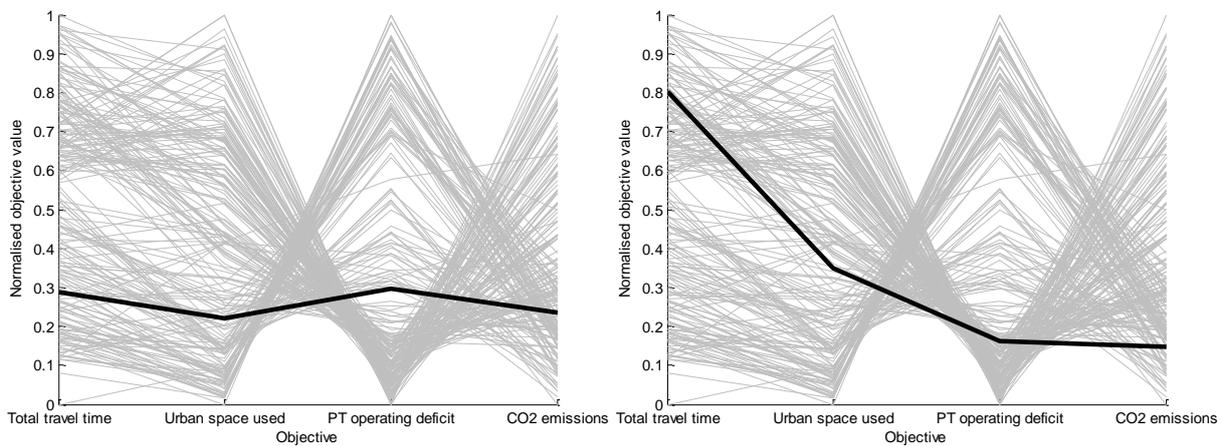
Figure 5.18: Min-max solution for all four objectives (left) and for OpD and CE (right)

5.4.2. Using values of decision variables

Another method that was suggested by policy officers during the interviews to gradually reduce the number of solutions is selecting certain values for decision variables from the Pareto set. This can be relevant in the political context of decision making, where each

political party may have had certain measures in its election-programme and therefore explicitly values certain measures above others. An interactive design process arises, where these political preferences not included in the objective functions come into play in the search for a final network design to implement. This design of network solutions after optimisation has two advantages over a pre-definition of these solutions. Firstly, during the selecting process (i.e. in a workshop), the values for objective values are immediately known for each Pareto solution, since the solutions have already been evaluated using the lower-level model. Consequently, if a certain decision implies very bad scores for an objective that is considered to be important, the decision maker can reconsider it immediately. Secondly, a suboptimal solution (given the four predefined objectives) is never chosen, since all solutions are in the Pareto set (i.e. non-dominated) and therefore all possible as final optimal solution.

An example of this kind of reduction is shown in Figure 5.19. First, only solutions that have a frequency of 6 buses per hour on the bus line between Amsterdam Sloterdijk and Schiphol (called 'Westtangent' in Dutch) are included. This results in 26 solutions that remain from the 210 Pareto solutions, but for all four objectives, both solutions with low values and with high values are still included. A further reduction to 7 solutions is achieved by selecting the solutions that also include a P&R facility along this new bus line, at Schiphol-North. As a result only solutions with high values for OpD and low values for the other three objectives remain (see the upper right plot in Figure 5.19). As a next step selecting the solutions that include the train station in the village of Halfweg (between Amsterdam and Haarlem) results in 3 remaining solutions, which have similar scores for the objectives. Finally, if also the train station of Amsterdam Geuzenveld is included (also between Amsterdam and Haarlem), only one solution remains. The result of choosing these measures for implementation is a low value for TTT, USU and CE, but a very high value for OpD. This example shows that pre-setting only 4 decision variables can already result in selecting only one Pareto solution, with consequences for objective values (in this example a very bad score for OpD). Furthermore, all other 33 decision variables are indirectly fixed to a value in this way. Note that the considered set is a result of a heuristic, so it is an approximation of the Pareto set that probably contains much fewer solutions than the number of true Pareto solutions. If the true Pareto set would have been known, probably many more decision variables would need to be fixed until only one Pareto solution remains.
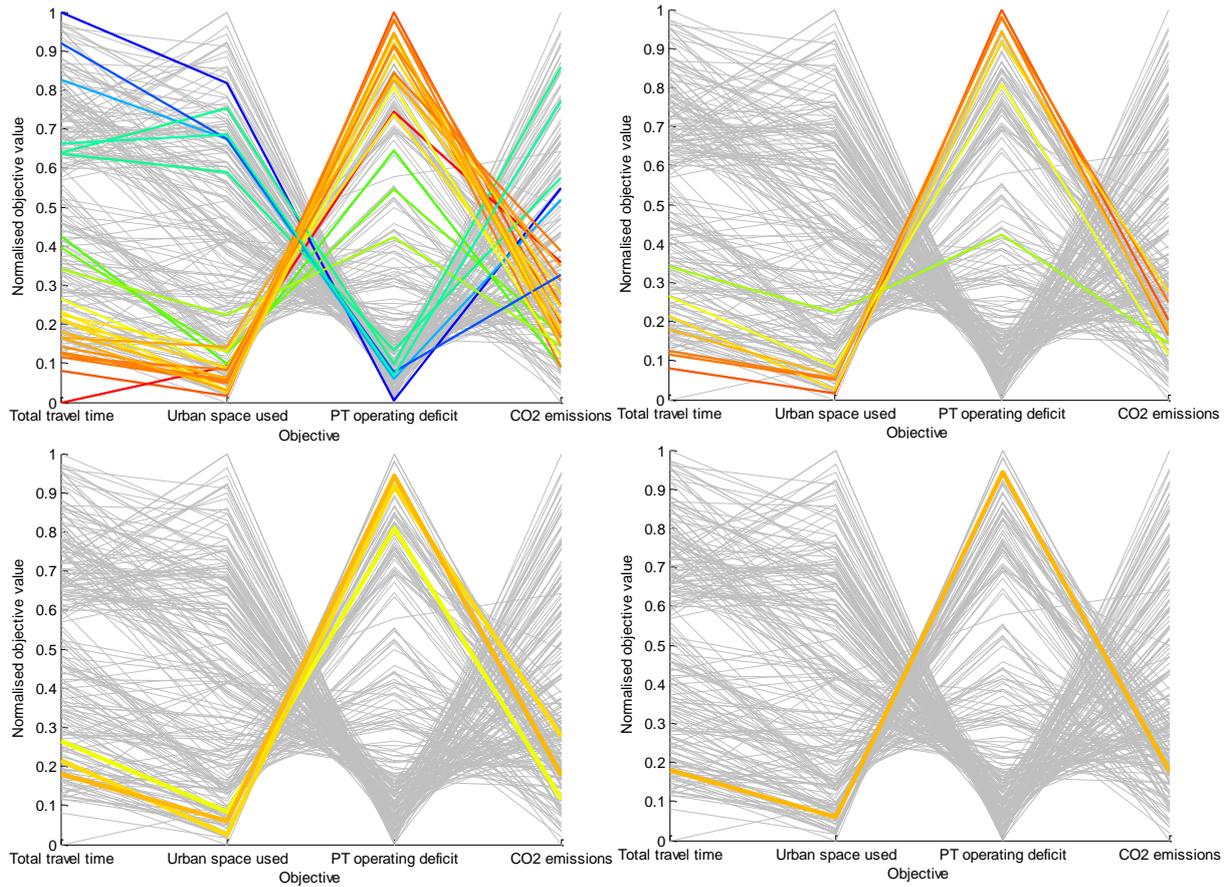
Figure 5.19: Step-by-step reduction to one solution to be implemented by selecting values for decision variables

5.4.3. Defining priorities for objectives

When looking for a best compromise solution, not all objectives are necessarily equally important, especially if these objectives are normalised. Instead of using equal preferences when determining the min-max solution, differentiated preferences can be used. In a parallel coordinate plot this can be visualised as in Figure 5.20: all 4 objectives are 'pushed down' as much as possible, but instead of a horizontal line a differentiation is made per objective. In the figure, the relative values of CE and TTT are equally important and the relative values of USU and OpD are also equally important. The first two are seen as more important than the latter two (since the constraint for the first two is tighter). Both relative difference in importance (Eq. 5.8, using coefficients RW_w) and absolute difference (Eq. 5.9, using coefficients AW_w) can be used to choose the compromise solution. Note that this approach is different from using the compensation principle in the form of a weighted sum of the objectives, which is related to the trade-off values found in Section 5.3.2.

$$RBCS_{W_C}(P) = \underset{\underline{y} \in P}{\arg\min} \left(\max_{w \in W_C} RW_w \bar{Z}_w(\underline{y}) \right) \qquad (5.8)$$

$$ABCS_{W_C}(P) = \underset{\underline{y} \in P}{\arg\min} \left(\max_{w \in W_C} AW_w + \bar{Z}_w(\underline{y}) \right) \qquad (5.9)$$
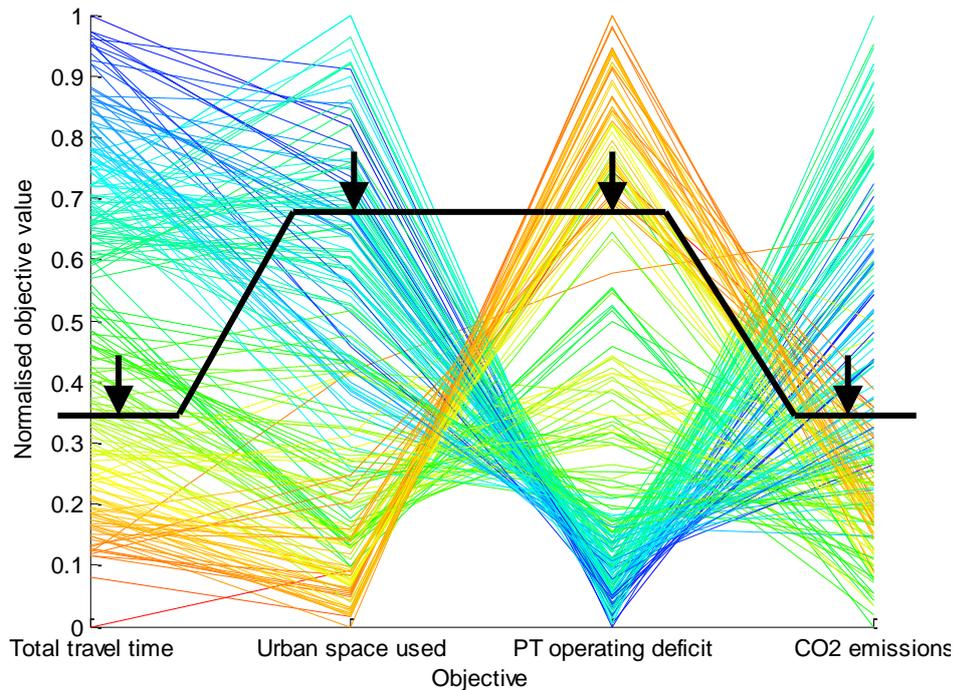
Figure 5.20: Constraints after optimisation, differentiated per objective functions

5.4.4. Conclusion

Application of a step-by-step pruning method showed to be effective to choose a final solution. Especially the combination of additional constraints and making a final decision by identifying the min-max solution showed to be useful to underpin a decision for a final solution to be implemented, which can act as a compromise between political parties (where each party stands for a certain objective).

In the case study in the Randstad it is shown that a compromise solution exists that has a reasonable score of around 0.3 for the normalised value for all four objectives, so it is possible to simultaneously satisfy all four objectives to a large extent.

5.5. Systematic pruning to reduce the number of choice options

As stated before, the set of Pareto-optimal solutions *P* can be large and therefore difficult to comprehend, manage and interpret. The step-by-step method presented in Section 5.4 to reduce the number of solutions results in focusing on the areas of the Pareto set that the decision maker finds the most interesting. This comes with the risk of overlooking interesting solutions if decision makers are driven by their initial thoughts. Systematic pruning can be used to circumvent this potential difficulty, because these methods result in a comprehensive subset of Pareto solutions, which are representative for the whole Pareto set. In this section the convex hull method, K-means clustering (Taboada et al., 2007) and practically insignificance trade-off (PIT) filter (Mattson et al., 2004) are used to prune the set. All of them are based on certain assumptions concerning the decision making process. The first two methods are provided for comparison reasons. To overcome some drawbacks of the PIT filter,

three new variants are presented and tested. For simplicity only three objectives are considered: TTT, CE and OpD.

5.5.1. Systematic pruning methods

Convex hull method

This method assumes that within the decision making process the objectives are eventually linearly weighted, e.g. using the weighted sum method. In that case, some of the Pareto solutions are irrelevant and never chosen as the final design (see Figure 5.21). The only relevant solutions are vertices of the convex hull, and therefore selected by this method. The convex hull method is a straightforward concept and the weighted sum method is widely used. However, there is neither control on the number of remaining solutions, nor on the level of equal distribution of solutions in the objective space. Furthermore, it assumes certainty concerning the performance of the solutions (as assessed by the lower-level model in terms of objective value).
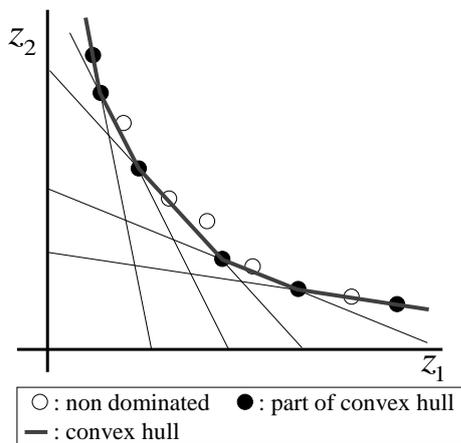
Figure 5.21: Visualisation of the pruning method convex hull for the two objective case

K-means clustering method

A clustering method assumes that within the decision making process distinct solutions in the objective space are relevant. In this pruning method a data mining clustering technique is used to cluster the similar solutions. In one specific clustering method, the K-means clustering method, data is grouped into k sets. The grouping is done by calculating the centroid for each group, while minimising the within cluster variance (the squared distance between each centre and its assigned data points). Each observation is assigned to the group with the closest centroid, and this process is repeated until the cluster assignments do not change. For each cluster the centroid is chosen as a representative solution. The algorithm is implemented as in Taboada et al. (2007). It uses a random selection of centroids as a start, from which it locally optimises the within cluster distance. For this reason, the procedure is repeated multiple times to increase the chance that the global optimum is found. This method is effective in reducing the number of solutions, because it will surely produce a selection of exactly k solutions. However, it has a few drawbacks as well: it does not guarantee an even spread of selected

solutions, it is sensitive to outliers, the selection of the representative solution is arbitrary and as the number of clusters changes cluster memberships can change in an arbitrary way.

PIT-filter

The method practically insignificant trade-off (PIT) filter (Mattson et al., 2004) assumes that within the decision making process the regions of the Pareto set that entail significant trade-off between objectives are the most interesting. In this pruning method the user has to define insignificant trade-off per objective. This is made operational by two parameters. Δ^t specifies the insignificance: if two solutions differ less than Δ^t from each other in one objective, the two solutions are considered as equal on that objective. Δ^r specifies the minimum level of spread along the Pareto set: if two solutions differ more than Δ^r they are considered to be different, also if all other objectives differ less than Δ^t. This parameter is needed if it is desirable that no big gaps exist in the Pareto set. Δ^r is always greater than or equal to Δ^t. In Figure 5.22 Δ^t and Δ^r are visualised. Subsequently the regions of practically insignificant trade-off can be determined for each solution and other solutions which fall within these regions are removed. This method guarantees a representation of the complete Pareto set and provides the solutions that are conceptually the most interesting for analysis. A drawback of the method is that the predefinition of insignificance is arbitrary.

The PIT filter is implemented as in Mattson et al. (2004) as PIT-0. In this filter the order of the solutions influences the outcome, so it is important to mention that in this case the Pareto set is sorted based on the values of the first objective function (TTT) before the filter is applied. To prevent any influence of the chosen starting solutions, an initiation phase is added, which assures that the extreme points of the Pareto set are kept in the reduced selection (i.e. the best solutions in the set with respect to each individual objective). This is important, because these solutions point out the boundaries of the Pareto set, which should be available to the decision maker in case he or she finds that particular objective very important. In this initiation phase, the PIT filter is applied to the solutions that score best in each of the objective functions individually. This results in a maximum of W selected solutions and already some rejected solutions in the PIT regions of these selected solutions. PIT-1 is a combination of this initiation phase and PIT-0.

In PIT-2 the condition for rejection is made tighter: only one objective is allowed to have a greater difference than Δ^t when comparing two solutions. This is done, because it assures that no insignificant trade-off is made between any combination of two objective. For example, if the objective values of a solution are within a distance of $(\Delta^t, \Delta^r, \Delta^r)$ from a base solution $\underline{Z}(\underline{y}_i)$, it is still possible that a significant trade-off exists between objectives 2 and 3, because both differences exceed the insignificance parameter Δ^t. To be sure that only insignificant trade-offs exist, the difference between the two solutions should not be greater than $(\Delta^t, \Delta^t, \Delta^r)$. In that case, both objective 1 and 2 have an insignificant trade-off with respect to the significant objective 3 and only one of the two solutions needs to be selected. Just like in PIT-1, in PIT-2 the solutions are sorted arbitrarily based on the value of the first objective and deal with the solutions in that order.

Next, it could be argued that the value of objective 3 is the only significant value in this case, so the value of objective 3 should determine which of the two solutions is the best. In PIT-3 the solutions with the best value for this significant objective are kept, instead of letting this

depend on the arbitrary order of solutions. This increases the number of selected solutions, because every objective value has to be checked for the best value, which sometimes results in selection of more than 1 solution within one PIT region (if another solution than the central solution is selected, or if a solution is selected based on its score in another overlapping PIT region).
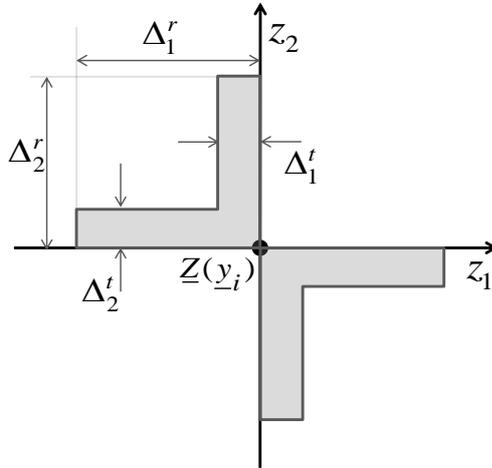
Figure 5.22: 2-dimensional visualisation of the Region of Practically Insignificant Trade-off for solution $\underline{Z}(\underline{y}_i)$ (shaded, defined by an insignificance parameter Δ^t and a spread parameter Δ^r for each objective)

5.5.2. Performance metrics

To compare the various pruning methods, three performance metrics are defined. The metrics are partly the same as the metrics used in chapter 4 and only focus on objective space. A new metric is defined to assess the evenness of the distribution of solutions along the Pareto front. This is relevant when assessing pruning results, because one aspect of representativeness is that all areas of the Pareto set are included in the pruned set as much as possible. Computation time is not included in the description of the metrics, because none of the methods encounters computational problems when applied to a Pareto set of a few hundred solutions.

Number of solutions in reduced set

The number of remaining solutions gives an indication to what extent the method is able to fulfil its primary goal: reducing the number of Pareto-optimal solutions, to give a more concise overview of the Pareto set.

Hypervolume

Hypervolume (While et al., 2006) is used to assess the pruning methods (in chapter 4 it had previously been used to assess the performance of solution algorithms, see Table 4.1 for the definition). The higher the value the better the pruned Pareto-optimal set is capable of representing the original set.

Equal distribution (Gini index)

The level of equal distribution of the selected solutions indicates to what extent the selection is an evenly spread set in the objective space. This is made operational by the Gini index (Gini, 1912), a well-known index to indicate the level of equal distribution within a dataset. In economics it is widely used to indicate the level of income inequality. This index is calculated for the distances from all solutions to their nearest neighbour solution in different directions. The use of the nearest neighbour in each direction, has an advantage over the more traditional measures spacing metric and diversity metric, which are for example used in Wismans et al. (2011), because these measures lack the ability to detect large gaps in the Pareto set.

The neighbours per direction are made operational by defining spaces Ω_i^g consisting of a number of convex cones g for every solution $\underline{y}_i$ with $\underline{Z}(\underline{y}_i)$ as origin, adding up to a whole space. In the 3 objective case $2^3 = 8$ convex cones Ω_i^g are distinguished for every solution, in such a way that each Ω_i^g represents an increase or a decrease / equal value per objective with respect to solution i. This is visualised for the 2-dimensional case in Figure 5.23. Note that two of the cones (both for the 2-dimensional case in the figure and for the general more dimensional case), with positive or negative changes for all objectives, are always empty, because the Pareto-optimal set only contains non-dominated solutions. Now the nearest neighbour from solution $\underline{Z}(\underline{y}_i)$ is searched in each direction g and each distance $d_{i,g}^{NB}$ to this neighbour is calculated by Eq. 5.10. Here, $e_{ii'}$ represents the Euclidian distance from solution i to i' in the objective space (normalised using the minimum and maximum values in the Pareto set per objective).

$$d_{i,g}^{NB} = \min_{i':i'\in\Omega_i^g, i'\neq i} e_{ii'} \qquad (5.10)$$
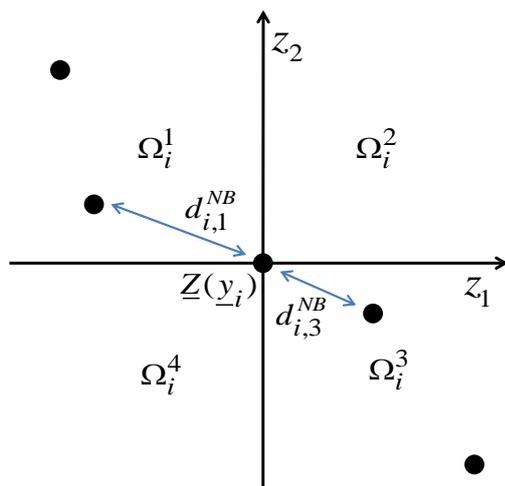
Figure 5.23: 2-dimensional visualisation of the division of a whole space into 4 distinct convex cones to identify neighbours of solution $\underline{Z}(\underline{y}_i)$

This assures that a gap in the Pareto set is detected if it exists, because a neighbour is found in each direction. If Ω_i^g is empty (i.e. no solution exists in direction g), $d_{i,g}^{NB}$ does not exist and no value is included for that direction. After calculating the values of $d_{i,g}^{NB}$ for each solution in P, the values $d_{i,g}^{NB}$ are sorted by increasing value. This results in a dataset $d_{\breve{i}}^{NB}$, where $\breve{i}$ is the

rank index and D^{NB} is the number of observations in the dataset. Based on this sorted set, the Gini index χ is calculated using Equation (5.2), taken from Thon (1982), where the second formulation is equivalent to the first formulation. A low value represents an even distribution and is to be preferred.

$$\chi = \frac{1}{D^{NB}} \left(D^{NB} + 1 - 2 \left(\frac{\sum_{\breve{i}} (D^{NB} + 1 - \breve{i}) d_{\breve{i}}^{NB}}{\sum_{\breve{i}} d_{\breve{i}}^{NB}} \right) \right) = \frac{2 \sum_{\breve{i}} \breve{i} d_{\breve{i}}^{NB}}{D^{NB} \sum_{\breve{i}} d_{\breve{i}}^{NB}} - \frac{D^{NB} + 1}{D^{NB}} \qquad (5.2)$$

5.5.3. Pruning results

The various PIT filters are applied to the Pareto set P, containing 210 solutions. Various settings for the insignificance parameter Δ^t and the minimum spread parameter Δ^r are chosen: 4 different values for Δ^t (0.01; 0.02; 0.05; 0.10) and 4 different values for Δ^r (0.1; 0.2; 0.5; 1.0). This results in 16 different parameter settings being tested for every variant of PIT. The values of Δ^t and Δ^r are normalised in the objective space: for each objective it is a fraction of the range covered by P for that objective. The PIT filters select fewer solutions when Δ^t or Δ^r is increased. Note that $\Delta^t = \Delta^r$ implies the situation without distinction between insignificant trade-off and spread. Furthermore, note that $\Delta^r = 1$ implies no minimum spread, covering the complete normalised objective space. In that case, the pruning takes place only based on insignificance.

Figure 5.24 shows the results. Note that no results are presented for PIT-0, because in this numerical example its results are the same as the results for PIT-1. This means that in this case PIT-0 accidentally retained the extreme values. As expected, PIT-1 selects fewer solutions than PIT-2 with the same parameter values, and PIT-2 selects fewer solutions than PIT-3. PIT-3 is not able to reduce the number of solutions substantially. This makes PIT-3 far less effective than the other 2 PIT variants, and makes it impossible to compare the performance of PIT-3 with the other filters for lower numbers of selected solutions.

Furthermore, Figure 5.24 contains the pruning results after applying the K-means clustering method in the normalised objective space and after applying the convex hull method. K-means clustering needs the number of desired selected solutions as an input. 16 values of k are explored (6, 7, 8, 9, 10; 11; 12; 13; 15; 17; 20; 25; 30; 35; 40; 50; 60; 70). Only one result is shown for the convex hull method, because the number of resulting solutions is a result of this method and cannot be influenced by any parameter.
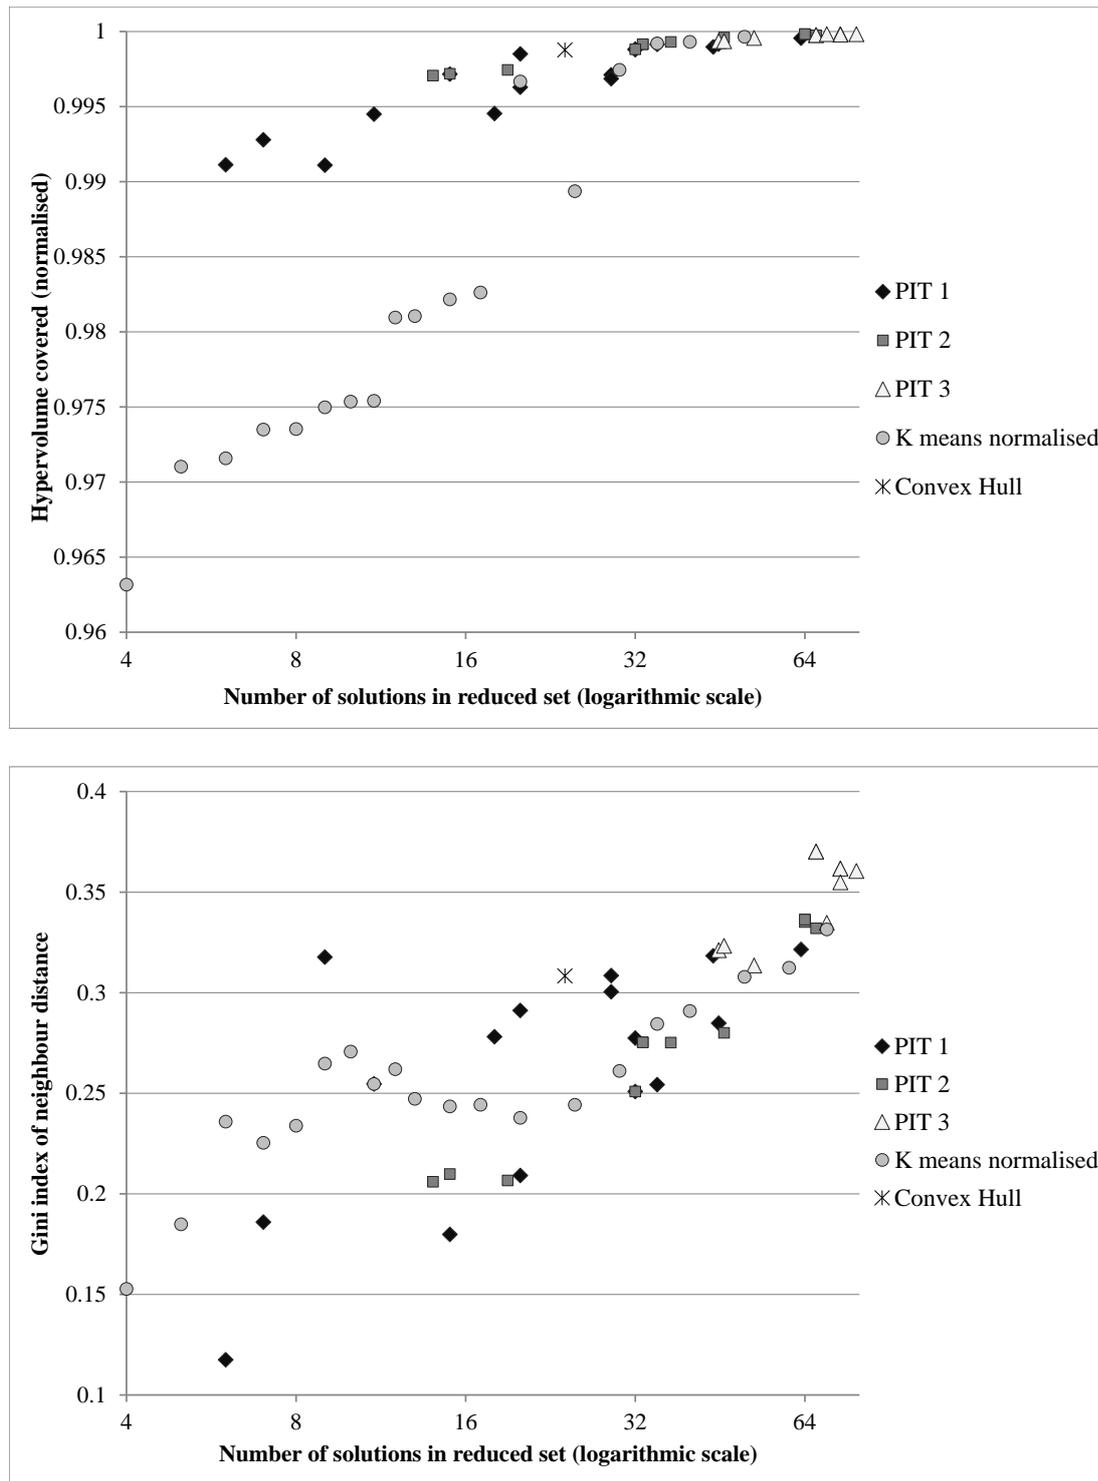
Figure 5.24: The pruning results for the various pruning methods. Above: the relation between the relative hypervolume covered by the pruned set and the number of solutions in the pruned set. Below: The relation between the Gini index of the pruned set and the number of solutions in the pruned set

The following observations can be made when looking at the results in Figure 5.24:

- There is a clear positive relation between the hypervolume covered and the number of solutions: it makes sense that more solutions are able to cover a larger space.
- The PIT filters clearly cover larger hypervolumes than K-means clustering for comparable size of the pruned set. This can be explained by the fact that K-means selects cluster centroids, while PIT selects the edge solutions where significant trade-off takes place, including the extreme solutions.
- PIT-2 in most cases slightly covers larger hypervolumes than PIT-1and PIT-3 for comparable size of the pruned set, or covers comparable hypervolume with a smaller size of the pruned set.
- Having fewer solutions in the pruned set roughly leads to a lower Gini index and consequently to a more equal distribution of solutions through the objective space, because the pruning methods mainly select those solutions that are distinct from the other solutions.
- When looking at comparable sizes of the pruned set, K means and PIT-1 on average have comparable values for the Gini index (although K means shows less variation in these values than PIT-1), while PIT-2 has lower (and thus better) values.
- The convex hull method is well able to reduce the number of solutions. It selects solutions that cover a large volume considering the small number of selected solutions. This makes sense, since the vertices of the convex hull point to the desired direction, considering all objectives. The convex hull method has a relatively high Gini index, given the number of solutions in the reduced set. This result is likely to be very case-specific: compared with the other methods, this method does not aim at an equal spread.

In Figure 5.25 a selection of the results of the pruning processes are plotted using scatter plots, showing all 119 Pareto solutions based on the three objectives and those Pareto solutions selected by the pruning method.

When comparing PIT-1 and PIT-2 in Figure 5.25C and 5.25D, it can be observed that with these specific parameters, both methods select 15 solutions. Both achieve an even spread along the Pareto set. With 15 remaining solutions, the figure shows that both PIT filters are very capable of retaining the main characteristics of the Pareto set in the pruned set. The corresponding result of the k-means algorithm with $k = 15$ in Figure 5.25A fails to include solutions on the edge of the set and has a less even spread. When comparing convex hull (Figure 5.25B) with PIT, it can be observed that convex hull clearly selects the solution on the border of the Pareto set, while PIT and K-means select a more evenly spread set of solutions.

In Figure 5.25E it can be observed that PIT-3 selects 45 solutions, despite the high values for Δ^t and Δ^r. If it is compared with the choice of the K-means algorithm in Figure 5.25F, it can be seen that also in this case the PIT filter selects solutions that are closer to the edges than K-means.

A: K-means clustering, $k = 15$

B: Convex hull

C: PIT-1, $\Delta^t = 0.10$, $\Delta^r = 0.2$

D: PIT-2, $\Delta^t = 0.10$, $\Delta^r = 0.50$

E: PIT-3, $\Delta^t = 0.10$, $\Delta^r = 0.5$

F: K-means clustering, $k = 45$
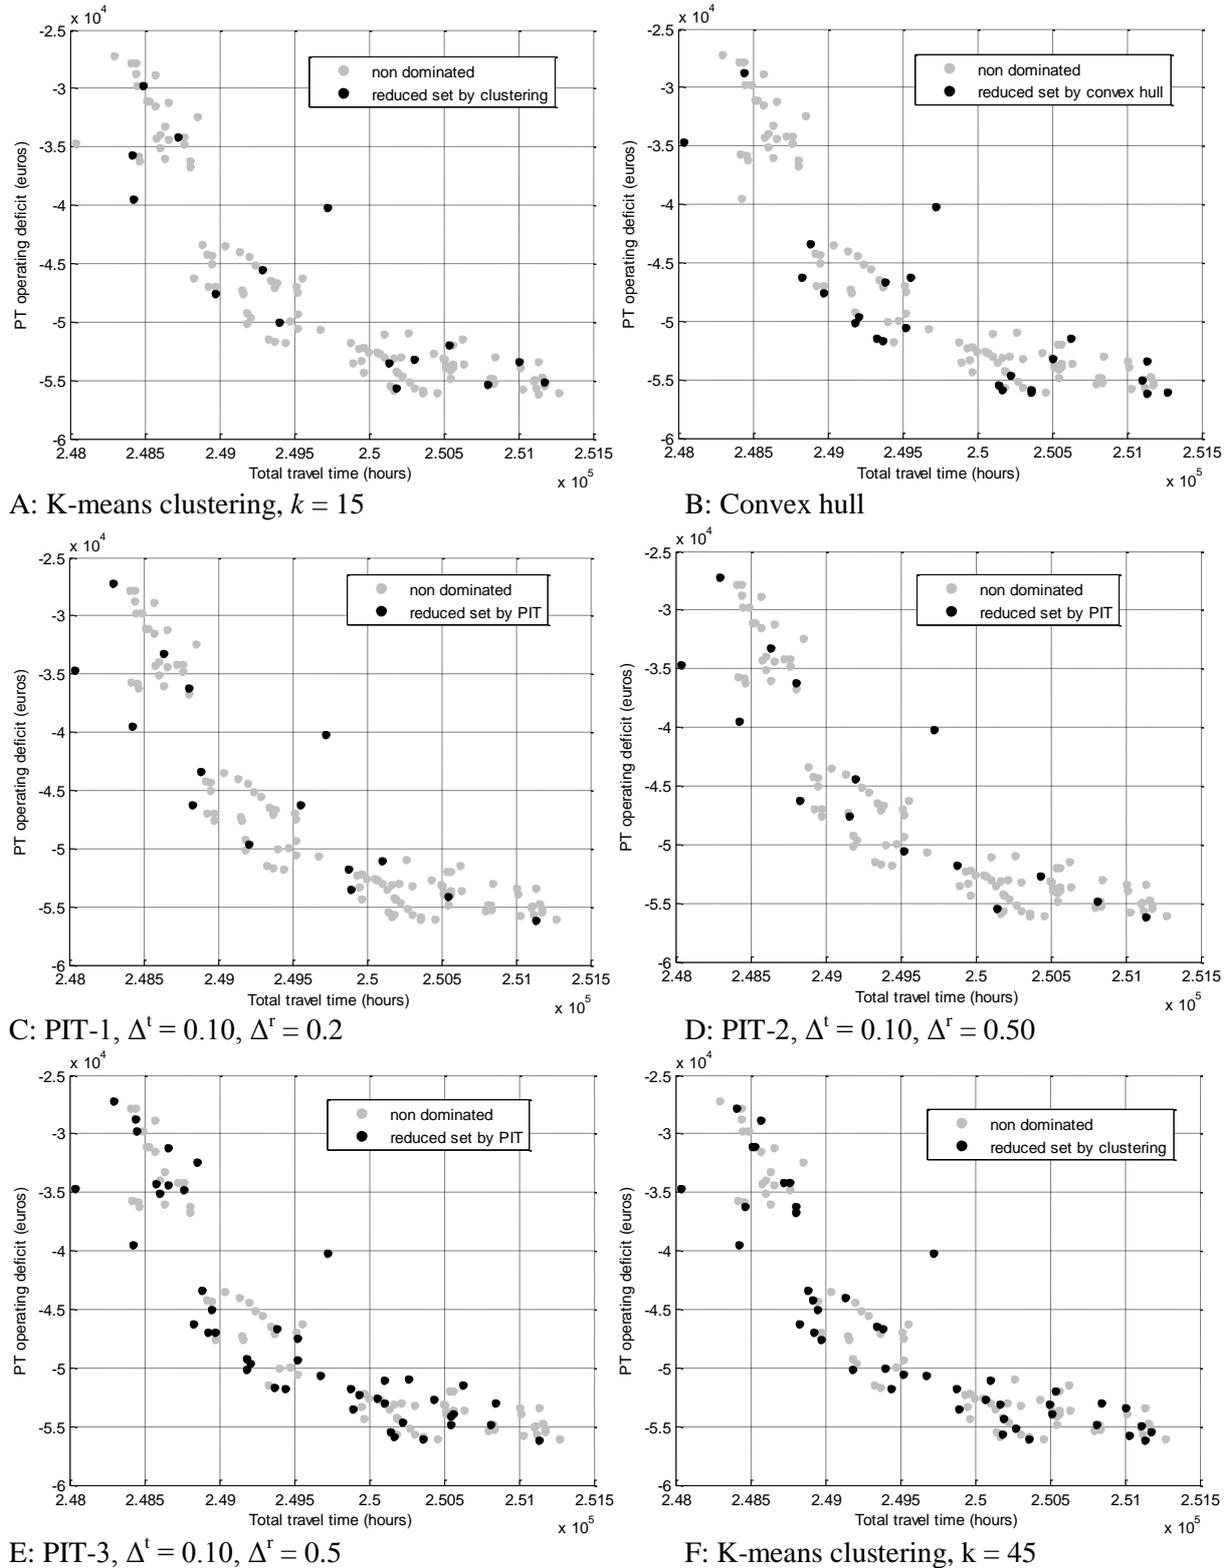
Figure 5.25: Selection of solutions by the different pruning methods

5.5.4. Conclusion

The application of pruning methods to the Pareto set results in a significant reduction of the number of solutions, while still representing the complete Pareto set. This prevents potential difficulties when a large number of Pareto network designs are to be compared to choose a final one, because the decision maker can determine the desired level of reduction by setting appropriate parameter values in the pruning methods.

The studied pruning methods convex hull, K-means clustering and PIT each have their own strengths: convex hull selects solutions that cover a large hypervolume, K-means clustering operates for every desired value of k and PIT achieves an equally distributed set (a low Gini index). The PIT filter outperforms K-means clustering in terms of hypervolume and the PIT-2 filter outperformed K-means clustering in terms of the Gini index. So if the focus of the decision maker is on the trade-off regions of the Pareto set, the PIT filter turned out to be a strong pruning method in this case to identify a limited number of distinct solutions that cover a large hypervolume.

The first conceptual improvements made to the PIT filter (PIT-2) made the filter slightly less effective in pruning the set, but based on the case study results the new filter is recommended over the old one, because it outperformed PIT-1 in terms of hypervolume and Gini index. Moreover, it has a superior theoretical foundation: it assures that no significant trade-off exists between a selected and a rejected solution in the 3 or more dimensional case. The second conceptual change (PIT-3) decreases the effectiveness of the filter strongly, so based in this case study it is not recommended to apply that filter further. However, it would still be interesting to further test these filters on other real-world Pareto sets from MO optimisation problems.

5.6. Conclusions

In this chapter several methods were applied to make the Pareto set resulting from the case study in the Randstad more useful as decision support information. First, conclusions are drawn concerning methodology development. Second, conclusions are drawn on the underlying design problem, the design of a multimodal passenger transportation network. The latter conclusions are mainly case-specific for the situation in the Northern part of the Randstad area in the Netherlands, but may be generalised to some extent to multimodal passenger transportation networks in general.

5.6.1. Methodology development

Interviews with policy officers showed that a Pareto set is valued positively as decision support information, because it enables an interactive process, where the consequences of certain choices can be demonstrated directly. Several visualisation methods and analytical methods turned out to make the Pareto set easier to understand and more useful to guide the decision maker to come to a final solution for implementation. The first method sets bounds to one or more objective values or selects a decision variable that is politically desirable, resulting in a step-by-step pruning process. This showed to be very helpful to finally select one solution for implementation. In this process, the parallel coordinate plot helps to get an impression of the data and to visualise the position of one or more solutions in the entire set.

The second method determines trade-off values between objectives, where scatter plots help to visualise trade-offs between 2 objectives (or a decision map to visualise 3 objectives). These trade-offs are seen as marginal costs, and therefore a decision maker can easily judge whether additional investments are worthwhile. The third method identifies the min-max solution and is useful in the search for a compromise solution. This showed to be useful in the political context of decision making, where each political party values each objective differently. When in negotiation during the political debate, no party will accept a solution in which its objective will not be reached at all.

Policy makers also indicated that their task becomes easier if the number of solutions is reduced systematically by the modeller / researcher. This is done by systematic pruning, which results in a comprehensive subset of Pareto solutions to represent the whole Pareto set. An existing pruning method (PIT-filter) is enhanced. The enhanced method (called PIT-2) outperforms other known methods using indicators that make 'systematic' operational, in terms of Pareto set attainment and spread. Application of this method to the case study provides an overview of 15 solutions as a representative subset of the entire Pareto set of 210 solutions, which is judged to be a reasonable number by the interviewed policy officers.

All developed methods in principle are ready to be applied to any practical choice situation where many options are considered and each option has scores for multiple objectives. However, to be more useful in practice, it is recommended to develop an interactive decision support tool in a software environment, which contains the developed methods to visualise and analyse the Pareto set. This tool would enable changing setting of the methods interactively, providing the opportunity for a direct feedback loop between the modeller and the decision maker, for example in a workshop setting. This interactive decision support tool could include the possibility to:

- select a pair of objectives to be plotted in a scatter plot and interactively click on an area of the Pareto front to reveal the trade-off values between a pair of objectives for the corresponding pair of solutions;
- compare these trade-off values with benchmark values (plotted as the slope of a line);
- select a third objective to be represented by colours (a scatter plot is changed into a decision map) and interactively change the colour scale;
- change the order of objectives when visualising the Pareto set in a parallel coordinate plot;
- prune solutions based on post-optimisation constraints on objective values and on decision variable values (based on preferences of policy officers / decision makers);
- highlight decision variable values on a geographical map after selecting a certain Pareto solution (for example in a parallel coordinate plot or in a scatter plot);
- have the possibility to interactively combine two or more objectives to one objective using weight factors and do all other analyses with the new combined objective;
- interactively select areas of the Pareto set and plot the shares of active decision variables based on the selection (also possible for the entire set);
- determine the compromise solution (i.e. the min-max solution) based on all or a subset of objectives, including the possibility interactively give different priorities per objective;

- change the parameters of systematic pruning methods, to interactively see the consequences for the resulting pruned set;
- have the possibility to aggregate the values of sets of decision variables of a similar type (for example the total number of active P&R facilities);
- present correlation matrices between pairs of objectives, between pairs of decision variables or between objectives and decision variables.

5.6.2. Multimodal passenger network design: case study in the Randstad

The objectives total travel time (TTT), urban space used (USU) and CO_2 emissions (CE) are all mainly in line with each other and opposed to operating deficit (OpD). The trade-off value between TTT and OpD is smaller than the value of time used in cost-benefit analyses on more than half of the Pareto front, so from the viewpoint of the regional government cost-efficient measures exist within the selection of measures considered. When TTT, USU and CE are combined to one objective value and compete against OpD, an average trade-off value is found as follows: for one euro of additional OpD, on average simultaneously 0.1 hours of TTT, 0.07 parking spaces of USU and 0.5 kilos of CE are gained.

The min-max solution has a reasonable score of around 0.3 for the normalised value for all four objectives, so it is possible to satisfy all four objectives to a large extent simultaneously. However, it was found that in this trade-off solution, in one part of the study area no measures are taken and in the rest of the study area a lot of measures are taken, so the distribution of the benefits is not even, which may lead to a discussion regarding equity.

When only solutions with low CE are selected, it is still possible to cover largely diverse scores for OpD and TTT. Furthermore, when searching for trade-off solutions between these three objectives, small losses in OpD and TTT result in a relatively large gain in CE.

When looking at individual measures, in the case study in the Randstad two 'always regret' solutions could be identified and no 'no regret' solutions. The decision variables in the case study that represent train routes and bus routes are complementary, rather than competitive. Furthermore, a sensitivity analysis provided insight into the relation of individual decision variables with each of the four objectives. This analysis showed that opening P&R facilities leads to more CE, because former PT trips are replaced by P&R trips. It further showed that increasing train frequencies and, to a smaller extent, bus frequencies in general reduces TTT, USU and CE, but increases OpD.

The measures (all related to multimodal trip making) that were selected in the case study can only contribute to small relative improvements with respect to the base network for the sustainability objectives that were defined (CE can be improved by 0.45%, TTT can be improved by 0.79% and USU by 1.5%). However, in absolute terms these possible gains are considerable: every AM peak almost 4000 hours of travel time, more than 2000 parked cars and 12 tons of CE (equivalent to the daily direct CE of more than 500 Dutch households) are saved. Furthermore, it is possible to reduce all four objectives simultaneously with at least 0.4% for each objective. In other words, a reduction of at least 0.4% is possible for TTT, CE and USU, with 0.4% less costs.

Finally, comparing the results of solving the defined multimodal MO-NDP with results that may be reached by TOD or road pricing, leads to the conclusion that redesigning the multimodal network is the only type of measure that enables simultaneous improvement for all considered aspects of sustainability, due to the adopted MO optimisation approach. However, when larger improvements in environmental or social sustainability are aspired, other types of measures are more effective, especially when the car alternative is made less attractive.

Chapter 6: Long-term robustness

In Chapter 5 the optimisation problem was solved for the second case study, using a demand prediction for the situation in 2030. This resulted in a Pareto set, which was analysed in detail to reveal information from it: the extent to which sustainability objectives can be improved by facilitating multimodal trip making and how the objectives and network designs relate. However, because the future is inherently associated with uncertainty, it is important to investigate the robustness of the Pareto set considering this uncertainty. In this chapter the long-term robustness is tested by applying the modelling framework for different scenarios regarding future transportation demand. The similarities and differences between the resulting Pareto sets are analysed. Finally, two other issues related to robustness of the optimisation result are discussed: the reversibility of measures and the effect one of the four objectives becoming obsolete due to unknown societal developments.

This chapter is partly based on

Brands T., L. J. J. Wismans, E. C. v. Berkum (2014). Multi-Objective Optimization of Multimodal Passenger Transportation Networks: Coping with Demand Uncertainty. OPT-i, 1st International Conference on Engineering and Applied Sciences Optimization. M. Papadrakakis, M. G. Karlaftis and N. D. Lagaros. Kos Island, Greece, pp. 547-561

6.1. Introduction

In the previous chapters the optimisation problem was defined and solved, resulting in a Pareto set. This Pareto set provides insight into which (re)designs of the multimodal passenger transportation network in the Randstad result in the best scores for various aspects of sustainability, into how much the new designs are able to improve the existing situation and into how these scores and designs relate. This insight is provided for the chosen definition of the optimisation problem, which is the combination of objectives, decision variables and lower-level transportation model. The latter includes input data like transportation demand, choice parameters, model structure and general parameters like oil prices. Each of these input data is subject to a certain level of uncertainty, for example because a certain external development or behavioural parameter is assumed. Tahmasseby (2009) defines robustness as follows: 'A transport system is said to be more robust, if it is capable of coping well with variations (sometimes unpredictable) in its operating environment with minimal damage, alteration, or loss of functionality.' In this chapter the focus is on long-term robustness, the extent to which a network is able to cope with uncertainty in developments over decades (i.e. variation of input data). The effects of day-to-day variability or the effects of incidents on the network performance are not studied.

In the modelling framework in this research objective values are used to determine the scores of network designs. Uncertainty in input data will propagate through the model: these influence the outcomes of the transportation model, which is used to determine the objective values. Effects of uncertain input on the output of transportation models received considerable attention in literature, which is reviewed in Rasouli and Timmermans (2012). They found that previous literature comprises a wide variety of sources of uncertainty, ranging from methodological reflections on the nature of choice models and model uncertainty in a fundamental statistical sense to concrete studies of a specific source of uncertainty. In most studies only a single source of uncertainty is examined and not much effort is paid to systemically apply variation to the factors of interest in a more sophisticated way. From this, the authors conclude that although these studies articulate and illustrate the issue at hand, evidence remains rather sketchy. Specifically for the traditional 4-step model they found that uncertainty often increases in the first 3 steps and reduces in the assignment step, due to capacity constraints and aggregation of OD flows. This suggests that the lower-level model in this research (where steps 3 and 4 of the 4-step model are carried out) may be limitedly sensitive for transportation demand input, and therefore the effect on the objective values is expected to be small. However, this lower-level model is used as a part of the modelling framework, so in the end the influence on the optimisation result is of interest. Besides the effect on the objective values, the uncertainty in model input also possibly affects the optimal physical transportation network designs for a given set of decision variables. This is relevant, because it would mean erroneous decisions could be made.

Robustness is incorporated in NDPs in various ways in literature. Three main approaches are identified: robustness can be included in one or more objective functions during optimisation, it can be included by creating an adaptable design to be flexible for future circumstances or it can be assessed after optimisation. Examples of multi-objective (MO) network design studies can be found for each of those three approaches to include robustness.

For the first approach several examples can be found in literature. Santos et al. (2009) define robustness as one of the objective values in terms of reserve capacity in the network. In that case, the interpretation is the robustness of transportation networks themselves: the extent to which these networks perform well under disruptive circumstances. Examples are a blocked link due to an accident or extreme demand due to an event with a lot of visitors, like a concert. Although such a definition would make sense in the multimodal context of the problem considered in this research, methods to cope with this interpretation of robustness of transportation networks in a multimodal context do not yet exist according to Van Nes et al. (2007). Sharma et al. (2009) use an additional objective, which is variance of the total travel time next to travel time itself, to integrate robustness for demand uncertainty into the network design problem (NDP). Compare et al. (2015) present an MO optimisation problem for gas turbine maintenance, where the cumulative distribution function is calculated for each objective resulting from uncertain parameter values. Knowing this function, any statistical measure may be used during optimisation (for example standard deviation or percentile values). However, the authors go one step further and define a new concept of Pareto ranking within the solution approach. This ranking concept is based on the entire cumulative distribution functions of the objectives and uses the probability that a solution is dominated by another solution. Zhu et al. (2014) argue that values of a single objective function for multiple scenarios for the future (e.g. travel time in scenario 1 and travel time in scenario 2) can also be seen as separate objectives. Using these objectives in MO optimisation results in a Pareto set that shows trade-offs between being optimal for one future scenario or for another scenario. If multiple objectives are to be included this approach is less suitable, because the number of objectives will increase by a factor equal to the number of scenarios. Alicino and Vasile (2014) start by translating an optimisation problem that includes uncertainty into a single or MO min-max problem equivalent to a worst-case scenario optimisation problem. All the studies in this first category result in networks that are designed to be robust and consequently have higher scores on indicators for robustness. The main drawback of this approach is that the computation time is larger, because indicators for robustness have to be calculated during the optimisation procedure, which implies that multiple runs of the lower-level model are needed per solution.

The second approach is described in Ukkusuri and Patil (2009), where a flexible investment scheme is introduced as a decision variable that is an answer to demand uncertainty. The concept of designing strategies instead of single solutions is also called Dynamic Adaptive Policy Pathway (Haasnoot et al., 2013). Rules for the phasing of implementation of measures are part of the design problem, where realisation of future measures may depend on developments that took place earlier. This implies that the network has to be assessed in terms of objective values for multiple moments in time and for multiple scenarios for the future. When combined with an MO optimisation algorithm this approach can lead to a Pareto-optimal set of pathways that is flexible with respect to future developments. This is an elegant method to cope with uncertainty, because it explicitly includes robustness in the network design. Main drawbacks are the definition and interpretation of the NDP becoming more complicated and computation time for solving it becoming larger, because for each solution the lower-level model needs to be run for multiple moments in time and for multiple scenarios for the future.

The third approach is demonstrated by Kasprzyk et al. (2013), who integrate uncertainty in an iterative decision-making process. First the Pareto set is determined by MO optimisation. After that, the ranges that objectives may cover under different (uncertain) circumstances are calculated for every Pareto solution by evaluating the objective value for a range of so-called 'states of the world', generated by Monte Carlo simulation. This results in standard deviation values next to average values for each objective function. These standard deviation values can then be included in the process of choosing a final solution. The advantage of this approach is that only Pareto solutions need to be calculated for multiple scenarios for the future, instead of all solutions during the optimisation process. The drawback of this approach is that these values are not used during optimisation to guide the optimisation process towards robust solutions. As a result, potentially robust solutions not part of the initial Pareto set are not considered in this approach.

The optimisation problem in the second case study involves long computation times, making the third approach the only feasible approach in terms of computation times. The approach of testing robustness of optimal solutions after optimisation is extended by solving the optimisation problem for three different possible future transportation demand scenarios, which may occur due to various possible future socioeconomic conditions. This enables testing to what extent Pareto-optimal solutions for the reference situation investigated in Chapter 5 still perform well under a different demand. It is tested whether the Pareto set is robust with respect to an uncertain demand. A method is presented that enables analysis of differences between Pareto sets resulting from optimisation processes with different demand assumptions. When these differences are analysed, it is important to take into account the measures represented by the decision variables. Some measures involve investments in infrastructure, implying that these measures cannot be easily reversed. Other measures involve PT operations, and can be reversed more easily. Finally for MO problems, the sensitivity for the set of considered objectives can be seen as long-term robustness as well, because priorities may change in the future, for example due to technological developments. Robustness for these changing circumstances is tested by omitting one objective and analysing the effects on the optimisation result. CO_2 emissions are omitted from the analysis, as an example of a possible future situation where superior technology leads to unlimited supply of renewable energy. Altogether, in this chapter a first step is made to incorporate robustness into the design of transportation networks in a multimodal and MO context.

The remainder of this chapter is structured as follows. In Section 6.2 the effect of different demand circumstances on the optimisation results is tested. First attention is given to the notation needed and to the demand forecasts that are used as input. Then the methods and indicators that are used in the robustness analysis are defined. The section concludes with results on the comparison between optimisation results using different demand input. In Section 6.3, the reversibility of measures is discussed: in case a different situation occurs in the future, some measures are easier to reverse than others. In Section 6.4 the effect of one objective function becoming obsolete in the future is analysed. In Section 6.5 conclusions are given.

6.2.   Performance of Pareto sets for different transportation demand scenarios

The robustness of the modelling framework is tested by investigating the sensitivity of the resulting Pareto set with respect to transportation demand input. The extent to which the optimisation result (i.e. the Pareto set) still performs Pareto-optimally under different demand input is tested: the objective values in Pareto sets optimised for a certain demand scenario are compared with Pareto sets optimised for another demand scenario, but assessed using the first demand scenario. Furthermore, differences and similarities between these Pareto sets in the solution space are identified.

6.2.1.   Notation

As defined earlier in Section 4.3.1, the set of N^j solutions $P^j = \left\{ \underline{y}_1^j, \underline{y}_2^j, \cdots, \underline{y}_{N^j}^j \right\}$ is defined as the Pareto set resulting from optimisation process j, which includes all non-dominated solutions with respect to all calculated solutions ϕ^j. Each optimisation process j has a corresponding transportation demand q^* for which that Pareto set was optimised. In addition, for each Pareto-optimal solution $\underline{y}_i^j \in P^j$ the objective values are calculated for the other demands q, resulting in Q objective vectors $\underline{Z}^q(\underline{y}_i^j)$ for each solution i. These objective vectors consist of W objective functions, denoted by index w, for each demand q: $\underline{Z}^q(\underline{y}_i^j) = \left\{ Z_1^q(\underline{y}_i^j), \cdots, Z_w^q(\underline{y}_i^j), \cdots, Z_W^q(\underline{y}_i^j) \right\}$. The Pareto set optimised for demand q^*, but containing objective values calculated for demand q is denoted as $P_{q^*,q}^j$.

6.2.2.   Demand forecasts

A demand forecast for the year 2030 is used as the reference situation in the second case study. This forecast is also used in practice and is based on a scenario study regarding the expected spatial economic developments in the Netherlands (CPB et al., 2006). In this chapter the resulting Pareto sets from solving the NDP using this demand scenario and two other realistic demand scenarios are compared. Besides the demand forecast for 2030, these scenarios involve a demand forecast for 2020 and a transit-oriented development (TOD) demand forecast.

The demand forecast for 2020 is based on the same scenario study and represents a situation with slower growth and therefore a lower total transportation demand. In Table 6.1 the relative change with respect to the 2030 demand scenario is given for aggregated OD pairs. The geographical locations of these aggregated zones can be found in Figure 6.1. The table shows that especially the areas of Almere, Hoofddorp / Schiphol and Amsterdam South and East have a smaller demand in the 2020 scenario. Some zones have a little larger demand in the 2020 scenario due to larger average household size in 2020 compared to 2030. In general the structure of the OD matrix remains comparable: most changes are relatively small, and where larger changes occur the distribution over the destinations remains similar.

This 2020 scenario can be seen as a lower boundary scenario for the use of PT, because the planned car infrastructure extensions in combination with low growth imply relatively low travel times for car. This leads to lower PT use, which makes the measures considered in the optimisation process less effective.

Figure 6.1: Aggregated zones used in the aggregated OD matrices in Tables 6.1 and 6.2. The zones are indicated by town or city name(s) and include their surroundings

Table 6.1: Percentage change of 2020 demand scenario w.r.t. 2030 demand scenario

	Haarlem	Beverwijk	Zaanstad	Amsterdam West	Amsterdam Centre	Amsterdam S. and E.	Schiphol Hoofddorp	Almere	Leiden R. D.H.	Utrecht H. A.	Rest of N-H	Rest of Netherlands	Total
Haarlem	-2	-2	-5	-2	3	-1	-4	-4	-2	-1	-4	1	-2
Beverwijk	0	1	-1	1	6	2	-1	-1	4	1	0	1	1
Zaanstad	2	1	0	3	7	4	-1	-2	6	4	-2	5	2
Amsterdam West	-2	-4	-5	0	-1	-1	-6	-3	-1	-1	-4	0	-1
Amsterdam Centre	0	-1	-1	2	1	1	-5	-4	1	2	-3	2	0
Amsterdam South and East	-4	-6	-7	-4	-6	-4	-8	-11	-4	-6	-10	-5	-6
Schiphol Hoofddorp	-8	-7	-11	-5	-7	-7	-2	-16	-5	-7	-10	-16	-5
Almere	-25	-34	-26	-5	-14	-14	-22	-19	-23	-19	-33	-28	-19
Leiden Rotterdam Den Haag	-4	-3	-4	-1	0	-1	-5	-4	0	-1	-2	-1	-2
Utrecht Hilversum Amersfoort	-2	-4	-1	2	1	-1	-6	-9	-2	-2	-2	4	-2
Rest of North-Holland	-4	-1	-2	1	1	1	-3	-6	1	0	-4	-1	-1
Rest of The Netherlands	0	-2	1	2	1	1	-9	-8	2	3	-2	0	-3
Total	-3	-1	-2	0	-1	-2	-5	-14	-2	-5	-4	-11	-4

The TOD demand forecast (used earlier in Chapter 5) contains the same level of economic growth as the demand forecast for 2030 and therefore the total demand for transportation is approximately equal. However, all growth (like new jobs and new houses) takes place in the vicinity of train stations, instead of throughout the whole region, resulting in a demand pattern that fits the PT network better. Note that TOD requires a strong and consistent policy and is difficult to realise (Tan et al., 2013), but it is useful as an upper boundary scenario in the robustness analysis. To increase the realism of the used TOD scenario, only new developments are shifted to the vicinity of train stations. Existing houses and jobs are maintained at their original locations, which will be the majority of the built-up area by 2030. Furthermore, some train station areas have already been entirely developed and therefore are not able to absorb growth, which is taken into account in the used TOD scenario. In Table 6.2 the relative change with respect to the 2030 demand scenario is given for the aggregated OD pairs. Also in this scenario the structure of the OD matrix remains comparable: most changes are relatively small, and where larger changes occur the distribution over the destinations remains similar. Especially in the area of Schiphol / Hoofddorp more jobs are located, because these locations are easily accessible by train and still have space for additional developments. This results in more arrivals during the modelled AM peak period. Fewer jobs are located in Almere, Beverwijk and Zaanstad. Fewer residents are located in Almere and Hoofddorp and more in Amsterdam South and East. Also within the aggregated zones residents are shifted from outskirts of Hoofddorp and Almere to neighbourhoods closer to train stations in the same town; this is not visible in this aggregated table.

Table 6.2: Percentage change of TOD demand scenario w.r.t. 2030 demand scenario

	Haarlem	Beverwijk	Zaanstad	Amsterdam West	Amsterdam Centre	Amsterdam S. and E.	Schiphol Hoofddorp	Almere	Leiden R. D.H.	Utrecht H. A.	Rest of N-H	Rest of Netherlands	**Total**
Haarlem	2	1	-2	0	3	2	0	1	0	5	3	8	2
Beverwijk	4	-3	-3	-5	1	0	-1	-6	4	2	3	-3	0
Zaanstad	-2	-5	1	6	3	-1	7	-3	3	1	3	3	2
Amsterdam West	7	-4	-2	-1	0	-1	5	0	1	0	0	3	0
Amsterdam Centre	-1	-7	-6	-1	-1	-2	17	0	-1	0	-3	3	0
Amsterdam South and East	8	0	-1	4	3	3	13	4	5	5	1	9	5
Schiphol Hoofddorp	-4	-8	-11	-7	-2	-6	-2	-5	0	-5	-3	4	-3
Almere	1	-7	-5	4	5	4	24	-4	-2	-6	-3	-8	-2
Leiden Rotterdam Den Haag	-2	-6	-9	-5	-2	-2	7	-6	-1	0	-3	4	0
Utrecht Hilversum Amersfoort	0	-5	-6	0	2	0	6	-11	2	2	0	7	0
Rest of North Holland	0	-4	-7	-2	-1	-2	2	-7	0	2	-1	6	-1
Rest of The Netherlands	0	-5	-6	1	1	0	22	-10	4	3	1	0	0
Total	2	-3	-3	0	0	0	7	-5	0	0	0	0	0

This TOD scenario can be seen as an upper boundary scenario for the use of PT, because the demand pattern associated with TOD is focused on establishing high demand relations which are especially well accessible using PT, making PT more attractive to use.

Altogether three transportation demand scenarios are considered (Q=3): the reference scenario (2030 demand), a scenario that is disadvantageous for PT (2020 demand) and a scenario that is advantageous for PT (TOD demand). The optimisation process is executed twice for each transportation demand, because the solution method is stochastic in nature. This results in 6 Pareto sets being calculated (J=6).

6.2.3. Indicators

The same indicators are used as in Section 4.3.2. A special situation occurs when using set coverage, because this indicator is defined to compare two Pareto sets and in this robustness analysis the set that is compared with (set $P^{j'}$ or set 2) is not necessarily a Pareto set. Not all solutions in set 2 necessarily perform Pareto-optimally, because the solutions in such a set were optimised for a different demand. In Figure 6.2 this is visualised for the 2-dimensional case, where set 1 is a Pareto set and set 2 contains solutions that are dominated by a solution in its own set.

In set 1 four different types of solutions may occur based on their dominance relation with solutions in the other set:

- solutions that dominate a solution in set 2 (circle);
- solutions that are dominated by a solution in set 2 (square);
- solutions that are neither dominated by nor dominate a solution in set 2 (star);
- solutions that are dominated by a solution in set 2 and at the same time dominate another solution in set 2 (plus).

Especially the third type of solution is of interest for the robustness analysis. The more solutions in set 2 are within the 'star' type, the more robust the modelling framework can be qualified, because the 'star' type solutions perform Pareto-optimally (so perform optimally under a different demand scenario) without undermining solutions in set 1 (that were optimised for that specific demand and are therefore expected to perform at least as well as solutions not optimised for that specific demand). For the first type of solutions (circle) it is expected that these mainly occur in the Pareto set optimised for the corresponding demand. The second type of solutions (square) is mainly expected to occur in Pareto sets optimised for a different demand. Finally 'plus' type solutions can only occur in the Pareto set optimised for the corresponding demand, because this can only happen if the other set is not a Pareto set (i.e. optimised for a different demand). Set coverage $CTS(set1, set2)$ equals the fraction of solutions in set 2 that is dominated by at least 1 solution in set 1 (corresponding to the number of open squares divided by the total number of open symbols in Figure 6.2).
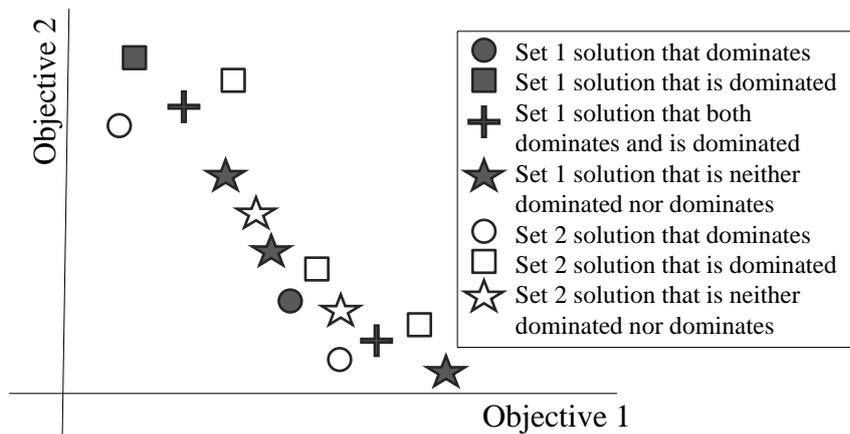
Figure 6.2: Different types of solutions that can be distinguished concerning set coverage

6.2.4. Results

In this section the optimisation results using different demand forecasts are presented. Firstly, the relation between 2020 demand objective values and 2030 demand objective values is given to illustrate the impact of a different transportation demand. Secondly, the objective values are compared by presenting the minimum values per objective, the hypervolume values and the set coverage for pairs of sets. For each of the three demand forecasts, two outcomes of the optimisation process with a different random seed are presented (6 Pareto sets in total). Thirdly, the decision variables of Pareto sets that were optimised for a different demand are compared by using the defined indicators.

In Figure 6.3 the relations between objective function values of the same network designs are shown, calculated for two different demand scenarios (2020 demand and 2030 demand). The plot for total travel time (TTT) indicates that the effectiveness of the measures for TTT depends on the occurring demand: some network designs perform well for the 2020 demand and others for the 2030 demand. This is caused by network dynamics related to travellers' behaviour that make the objective TTT highly nonlinear. Solutions with an elaborate PT network perform well in the situation with high demand, because these networks have PT as a congestion-free alternative to cope with this demand. Solutions with a limited PT network perform relatively well in the situation with low demand, because the car network has enough capacity to provide low travel times. On the other hand, it can be observed that operating deficit (OpD) for 2020 demand almost shows a linear relation with OpD for 2030 demand. First, operating costs do not depend on demand. Second, for a given network marginal transportation demand will be distributed with almost fixed fractions to the travel alternatives (PT or car), leading to an almost linear increase of operating revenues. For the same reason urban space used by parking (USU) almost shows a linear relation when 2020 and 2030 demand are compared. An additional reason is that the difference in demand between 2020 and 2030 is limited in the subset of zones that are highly urban (used to quantify this objective function), because most developments take place in other zones that are not yet entirely built-up. The relation between 2020 demand values and 2030 demand values for CO_2 emissions (CE) is somewhere in between, showing an approximately linear relation. Especially the relation between 2020 demand values and 2030 demand values for TTT indicates that

solutions optimised for one demand probably do not perform optimal any more when assessed using a different demand.
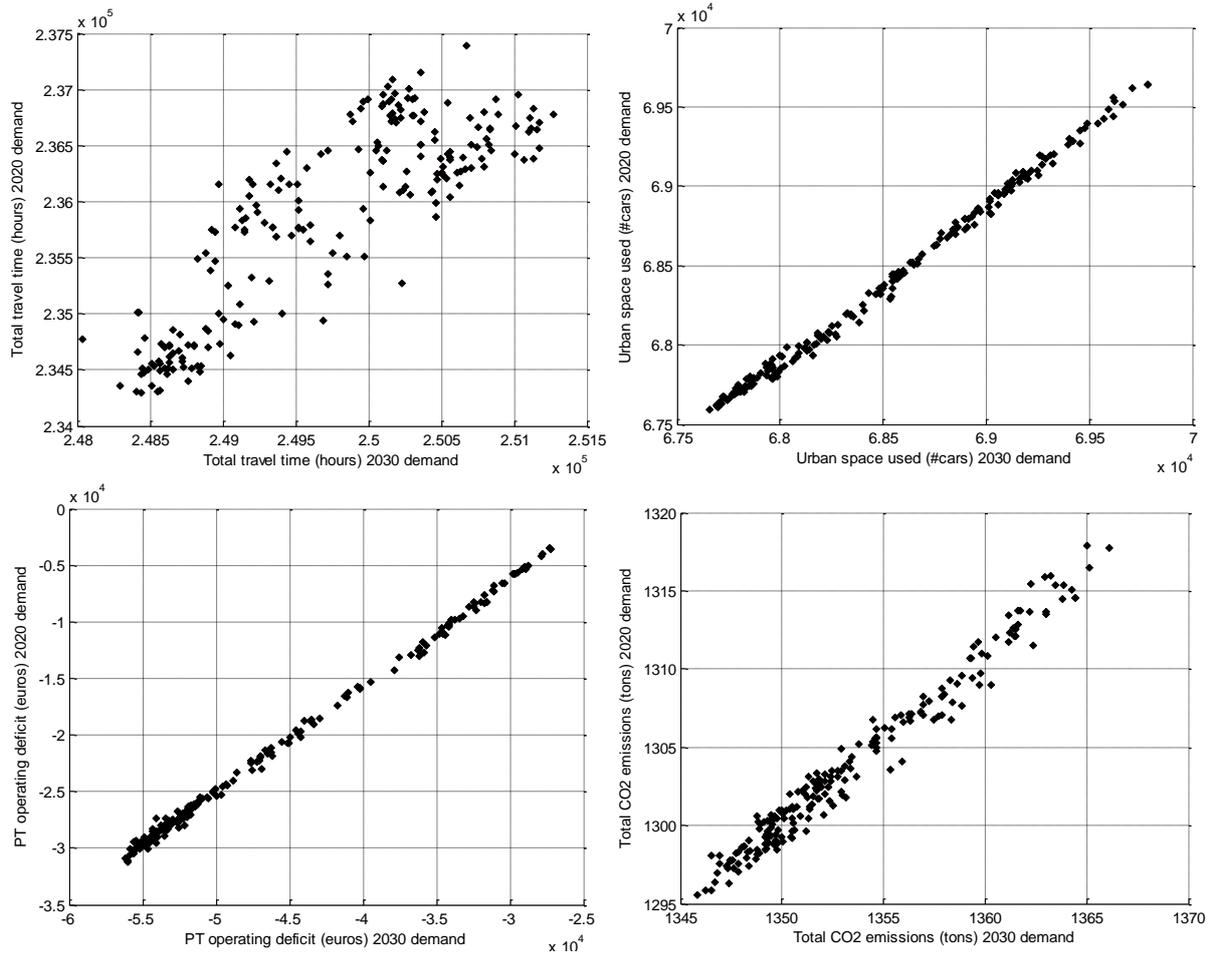
Figure 6.3: Relation between the values of a) TTT b) USU c) OpD d) CE for two demand scenarios. Each dot in the figure represents one solution (one multimodal transportation network design) assessed using 2020 demand (vertical axes) and assessed using 2030 demand (horizontal axes).

Number of solutions calculated

The used optimisation algorithm ε-NSGAII terminates after a fixed number of restarts. Due to this stop criterion, in each optimisation process a different number of solutions is calculated, because the restarts do not occur after a fixed number of iterations, but depend on whether ε-progress is made. In Table 6.3 the number of calculated solutions per optimisation process is shown, as well as the number of resulting Pareto-optimal solutions N_j.

Minimum values per objective

In Table 6.4 the relative minimum values per objective are shown for each Pareto set, compared to the minimum known value over all 6 sets that are calculated for a certain demand q (over a column in Table 6.4). It can be observed that for TTT, the best or near-best values occur in the sets that were optimised for the corresponding demand. This is in line with the

observation in Figure 6.3 that TTT highly depends on the demand input, so that optimising for a different demand results in different optimal solutions. For the other objectives, the best values also occur in sets that were not optimised for the corresponding demand. This is also in line with the observations in Figure 6.3, because a solution that performs optimally using one demand is likely to perform similarly when a different demand is used. This can also be observed directly in Table 6.4: the scores of the same Pareto set (in a row) are comparable for the three different demand results.

Table 6.3: number of solutions calculated during each optimisation process and number of resulting Pareto-optimal solutions N_j

q^*	j	Number of calculated solutions during optimisation process j	Number of Pareto solutions N_j
2020	1	3151	263
	2	2533	264
2030	3	2384	210
	4	1979	211
TOD	5	1914	260
	6	3328	308

Table 6.4: Lowest values per objective for each Pareto set $P_{q^*,q}^j$, percentage change w.r.t. the minimum value over all Pareto sets. Pareto set $P_{q^*,q}^j$ was optimised for demand q^* with the objective function values calculated using demand q

		TTT			USU			OpD			CE		
		q			q			q			q		
q^*	j	2020	2030	TOD	2020	2030	TOD	2020	2030	TOD	2020	2030	TOD
2020	1	0.00	0.15	0.05	0.11	0.13	0.14	0.00	0.56	0.86	0.05	0.02	0.00
	2	0.01	0.15	0.00	0.07	0.06	0.04	0.66	0.00	0.65	0.03	0.00	0.00
2030	3	0.00	0.00	0.04	0.04	0.00	0.00	0.22	0.75	0.00	0.06	0.04	0.09
	4	0.00	0.10	0.00	0.04	0.06	0.06	2.53	1.18	0.69	0.00	0.02	0.05
TOD	5	0.04	0.09	0.01	0.00	0.06	0.01	1.80	1.09	0.81	0.12	0.10	0.05
	6	0.01	0.19	0.02	0.07	0.03	0.02	3.32	1.87	1.33	0.07	0.11	0.08

Hypervolume

In Table 6.5 the relative hypervolume values are shown for each Pareto set $P_{q^*,q}^j$ as a percentage change w.r.t. the maximum value over all 6 sets calculated for a certain demand q (over a column in Table 6.5, recall that the hypervolume indicator is to be maximised). Consequently, the best value is 0, and all other values are negative values. For comparison reasons, the relative hypervolume value of the random set of solutions used to start up the optimisation algorithm is also shown. In all cases the best performing Pareto set (per column) is the set that was optimised for the corresponding demand. On the other hand, the hypervolume values of Pareto sets optimised for a different demand are only slightly lower.

Table 6.5: Hypervolume values for each Pareto set $P^j_{q*,q}$, which was optimised for demand $q*$ and for which the objective function values are calculated for demand q, percentage change w.r.t. the maximum value over all Pareto sets

		q		
$q*$	j	2020	2030	TOD
2020	1	0.0	-1.6	-1.0
	2	-0.3	-1.0	0.0
2030	3	-0.1	0.0	-0.3
	4	-0.9	-1.2	-0.5
TOD	5	-1.0	-1.8	0.0
	6	-1.7	-2.7	-0.6
Start set		-72.7	-71.0	-71.1
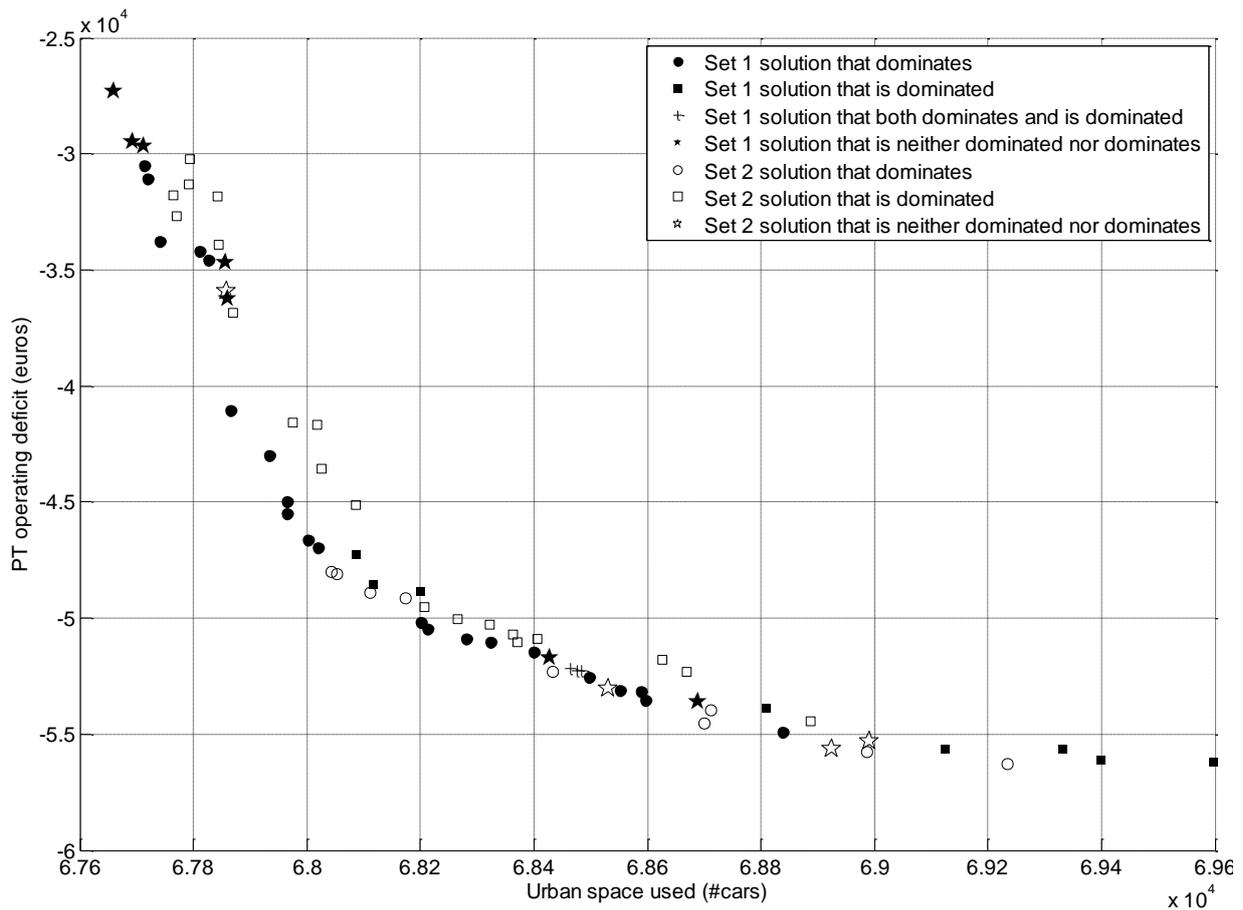
Figure 6.4: 2-dimensional example plot showing the different types of solutions

Dominance relations

In Figure 6.4 an example is shown of a classification into the types of solutions that were defined in Section 6.2.3, using the results of the case study in a 2-dimensional plot for USU and OpD. In this example, set 1 is a Pareto set optimised for these 2 objectives for the 2030 demand. Set 2 is the set that was optimised for the 2020 demand, but assessed using the 2030 demand. Therefore set 2 is no Pareto set any more: it contains solutions that are dominated by

solutions in the set itself. Furthermore, it mainly contains solutions that are dominated by set 1, but it also contains solutions that dominate at least one set 1 solution.

In Table 6.6 the average shares of the different types of solutions are shown when sets optimised for the corresponding demand are compared with sets optimised for a different demand. All combinations over j and j', with P_{q*}^{j} as a first set and $P_{q*',q'}^{j'}$ as a second set are included in the analysis, in a way that $q* = q'$ and $q* \neq q*'$ (implying also $j \neq j'$). For example, the 2 Pareto sets for the 2030 demand are compared with 4 other sets: 2 sets optimised for the 2020 demand and 2 sets optimised for TOD demand, all 4 assessed using the 2030 demand. This results in 8 combinations for the 2030 demand. Similarly, also 8 combinations for the 2020 demand and TOD demand exist, leading to a total of 24 combinations.

Table 6.6: Shares of different types of dominance relations

	circle	star	plus	square
Sets P_{q*}^{j} optimised for corresponding demand $q*$	48%	30%	9%	13%
Sets $P_{q*',q'}^{j'}$ optimised for a different demand $q*'$ assessed using demand $q' = q*$	15%	28%	-	57%

The shapes (circle, star, plus, square) correspond with the possible relations between solutions presented in Figures 6.2 and 6.4. The results show that only 43% of the solutions still perform Pareto-optimally under different demand circumstances (circle or star in the second type of sets), so a majority of these solutions does not perform Pareto-optimally any more when a different demand occurs. On the other hand, it is good to see that only 22% of the solutions optimised for that specific demand is dominated by a solution that was optimised for a different demand (sum of the shares of plus and square type of solutions). Concluding, a large difference can be observed in coverage between the two types of sets. However, it is also relevant to know to which extent a solution is not Pareto-optimal. As is indicated in Figure 6.4 and by the hypervolume scores, the differences in attained objective values are small, leading to the conclusion that the sets are all near-optimal.

Decision variable comparison

Finally, the values of decision variables for solutions in pairs of Pareto sets are compared. A distinction is made between sets where $q* = q*'$ and sets where $q* \neq q*'$. In words, the first pairs of sets were optimised for the same transportation demand and differences between the sets are caused by the random nature of the solution algorithm. The second pairs of sets were optimised for different transportation demand, so these differences are both caused by the random nature of the solution algorithm, as well as by this different demand input. The results for three indicators (as defined in Section 4.3.2) are shown in Table 6.7. These numbers indicate that the resulting decision vectors only have slightly more differences when optimised for a different demand than differences that are caused by the random nature of the heuristic.

Table 6.7: Differences between sets concerning decision variables: comparison between pairs of sets optimised for the same transportation demand and pairs of sets optimised for a different transportation demand

	$q* = q*'$	$q* \neq q*'$
$AD(P_{q*}^j, P_{q*'}^{j'})$	8.9	9.0
$AND(P_{q*}^j, P_{q*'}^{j'})$	4.1	4.3
$AFD(P_{q*}^j, P_{q*'}^{j'})$	0.123	0.132

6.3. Reversibility of measures

In this section the reversibility of measures is discussed. Some measures involve new infrastructure, like new train stations and P&R facilities, which in principle is built for decades. On the other hand, frequencies of bus lines can be adjusted annually (with the introduction of a new timetable) or even more often (by introducing an intermediate timetable change). The same should hold for train frequencies, but it may be necessary to invest in rail tracks and rolling stock to realise a higher frequency. As described before, assessing the capacity of the train tracks is outside of the scope of this research, but if investments are needed, it would be a waste of money if in the end the higher frequencies turn out not to be necessary. The same holds for express train status of train stations: it is not assessed whether additional investments are needed to dispatch all the trains, but in principle express train status can easily be reversed. Concluding, the decision variables are grouped into three classes: easy to reverse (bus frequencies), easy to reverse, but with risk on disinvestment (train frequencies and express train status of stations) and hard to reverse (new train stations, P&R facilities and tram line extensions).

Using these three types of measures regarding reversibility, the differences over the three demand scenarios concerning the effectiveness of each measure are studied. Similar as in Section 5.3.3, the fraction of decision variables in the Pareto sets that have a nonzero value is used as an indicator for the effectiveness of the measure over all four objectives. The fractions are determined considering all solutions included in the two known Pareto sets for each demand scenario and plotted in Figure 6.5.

Regarding bus frequencies (easy to reverse, the upper graph in the figure), some decision variables show very small variation and others show much larger variation. This is not a problem, because if another demand scenario appears to come true in the future, these measures are easy to reverse.

For the second category of measures that are easy to reverse with the risk of disinvestments (the middle graph in the figure) also large differences can be observed over the measures. Some measures have very similar fractions over all three demand scenarios, but others have larger differences. Variables 7 (express train status of Hoofddorp) and 33 (local train from Hoofddorp to Almere) have large differences for the TOD demand scenario. A closer look shows that these variables are related: in the TOD scenario more local trains are running, making Hoofddorp less suitable to be an express train station. A similar situation exists for

variable 30 (train from Zandvoort to Amsterdam) and variable 20 (bus from Zandvoort to Amsterdam). Although these measures perform differently over the demand scenarios, they seem to be at least partly interchangeable, and therefore to some extent robust for demand uncertainty. In Section 5.3.3 it was observed that the effects of changed train frequencies on objective values in general are larger than the effects of changed bus frequencies, so the measures that are easiest to reverse have a smaller effect.
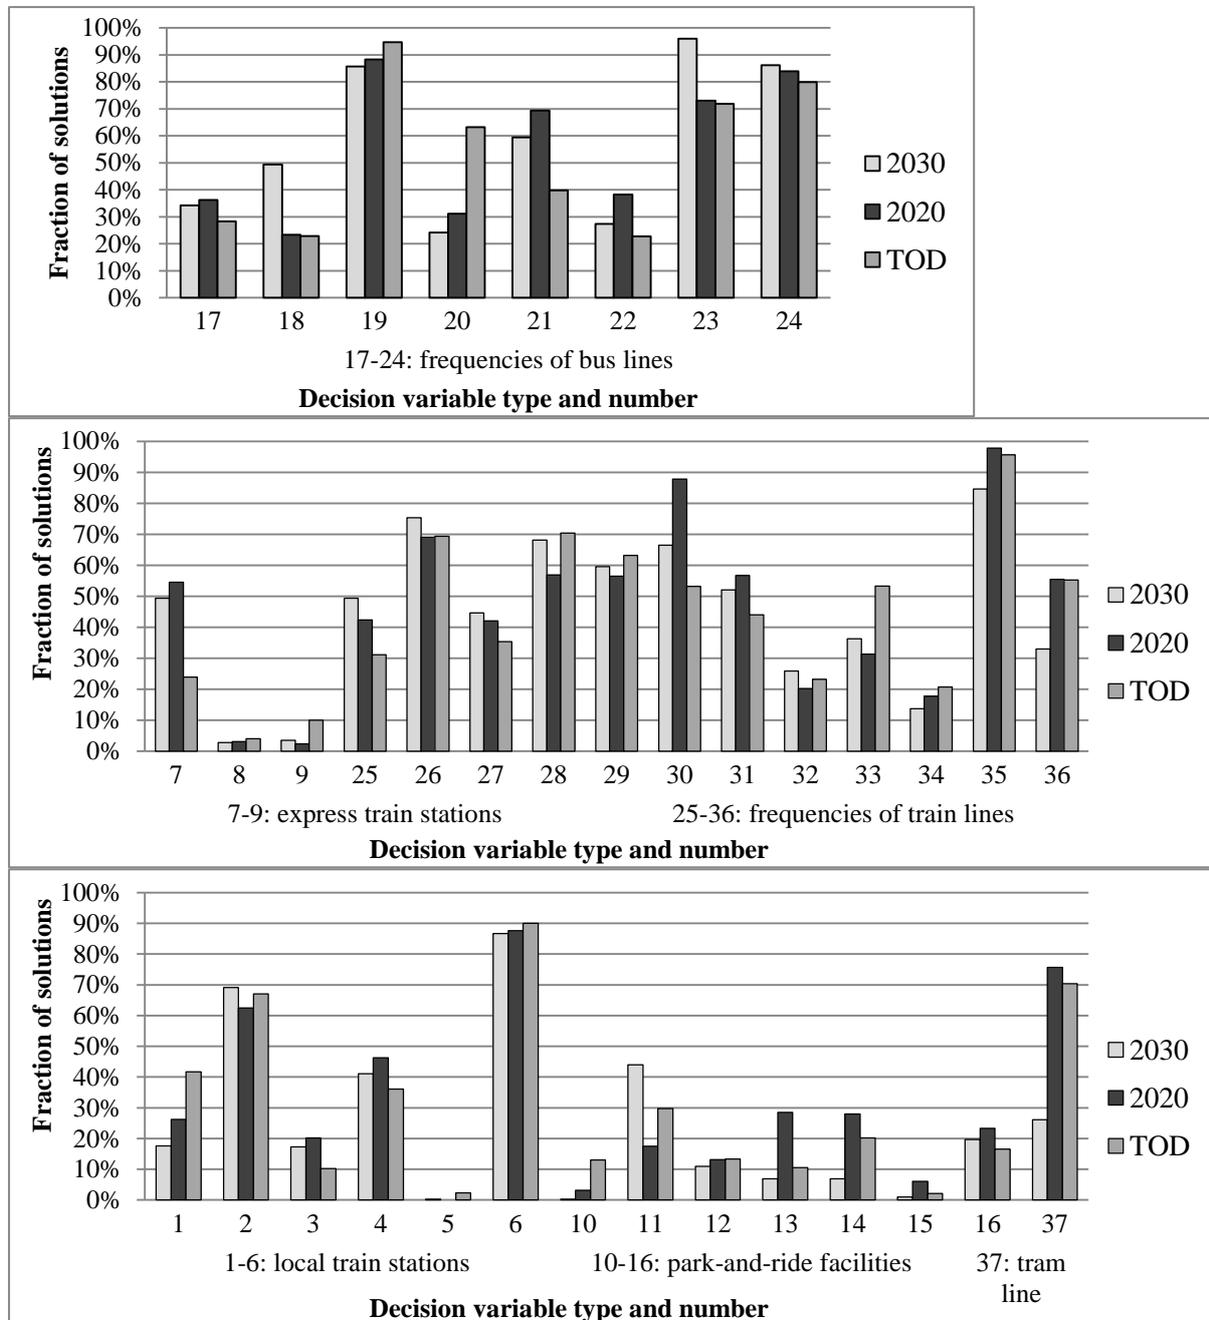

Figure 6.5: Fraction of solutions that have a nonzero value for each decision variable per demand scenario for a) measures that are easy to reverse b) measures that are easy to reverse, but with risk on disinvestment and c) measures that are hard to reverse

Finally for measures that are hard to reverse again large differences are observed over the measures. Variable 6 (train station Amsterdam Westerpark) is effective in all three demand scenarios and variable 5 (train station Nieuw Sloten direction North) is not effective in all three demand scenarios. However, for variable 1 (train station Halfweg-Zwanenburg) and for most P&R locations the effect differs over the scenarios. The largest difference is observed for variable 37 (tram line extension in Amsterdam Geuzenveld), which may be a result of random processes during optimisation, because this variable has very little influence on each of the objective functions (see Section 5.3.3).

Concluding, the robustness for demand uncertainty of measures in each category of reversibility is as follows. Considering that bus frequencies are easy to reverse, there is no reason not to implement these measures. Furthermore, investments in increased train frequencies and express status of stations have a low risk of being regretted when a different transportation demand occurs, because these measures do not vary strongly over the demand scenarios or are interchangeable. Moreover, these measures are important to achieve improvements, because they are the most effective in terms of objective values (see Section 5.3.3). Finally, more care should be taken when investing in new train stations and P&R facilities, because these measures are hard to reverse and in some cases are less effective under different demand circumstances.

6.4.   One obsolete objective function

In this section the effect of one objective function becoming obsolete in the future is analysed. Next to the influence of demand uncertainty, as was discussed in Sections 6.2 and 6.3, long-term robustness of the optimisation result relates to the objectives included in the problem definition. Future, yet unknown societal developments may cause objectives to become irrelevant, or may introduce new relevant objectives in the problem. In this section the objective CE is omitted from the analysis, as an example of a possible future situation in which superior technology leads to unlimited supply of renewable energy.

Table 6.8: Ranges covered by the 4 objective and the 3 objective Pareto set

		TTT	USU	OpD	CE
Minimum, relative w.r.t. base	4 objectives	-0.79%	-1.54%	-6.19%	-0.45%
	3 objectives	-0.79%	-1.54%	-6.19%	-0.40%
Maximum, relative w.r.t. base	4 objectives	0.50%	1.55%	2.96%	1.05%
	3 objectives	0.50%	1.55%	2.96%	1.05%

The Pareto set resulting from the original 4-objective optimisation problem contains 210 solutions. When CE is omitted as objective, 131 solutions remain Pareto-optimal considering the 3 remaining objectives TTT, USU and OpD. In Table 6.8 the ranges covered by both sets are shown. As expected, the improvement of CE with respect to the base solution becomes smaller: 0.45% improvement in the original result and 0.40% in the Pareto set that is only based on 3 objectives. This shows that still a large share of the improvement is achieved, which can be explained by the relations between the objectives: CE is especially in line with USU, but also with TTT. This is also the reason for the ranges covered by the other three

objectives to remain the same, although it was expected that at the maximum values of the other objectives would decrease to some extent. The best scores for the other three objectives remain the same by definition. The consequence of this observation is that optimising with respect to CE is not likely to be regretted, because it does not lead to worse maximum values for the other three objectives.

It is relevant to know whether certain measures become more or less effective, due to CE becoming irrelevant. In Figure 6.6 the nonzero fractions for each decision variable are plotted for both the original 4-objetive Pareto set as for the new 3-objective Pareto set. The effect of omitting CE from the analysis turns out to be small: the fractions change only slightly. This may be explained by the small range covered for CE in the optimisation problem: it is difficult to influence CE with the selected measures, so it does not have a strong impact on the structure of the Pareto set. This was earlier observed in Section 5.2.2, where solutions with low CE were found along almost the entire range covered for OpD and TTT.
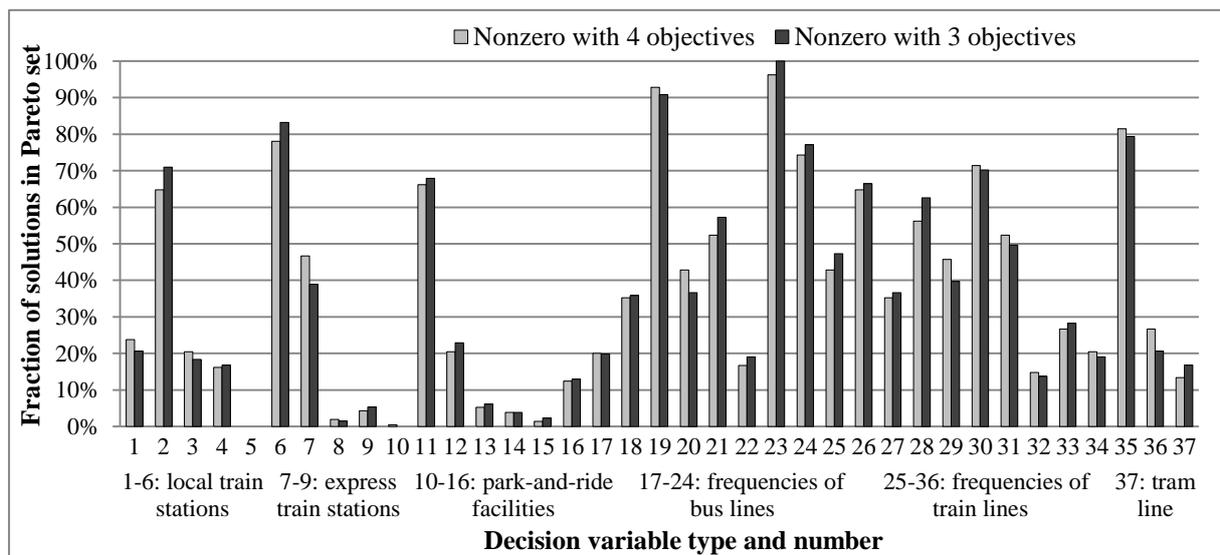
Figure 6.6: Fraction of solutions that have a nonzero value for each decision variable, with distinction between the original 4-objective problem and the 3-objective problem that remains after omitting the objective CO_2 emissions

6.5.   Conclusions

In the multimodal passenger transportation NDP, the demand forecast used is an important assumption. Demand forecasts always contain a certain amount of uncertainty. In this chapter the impact of different demand assumptions on the results of MO optimisation of a multimodal transportation network is shown. Furthermore, the impact of one objective becoming obsolete is analysed.

For the case study it can be concluded that a different, realistic transportation demand forecast has a strong influence on the Pareto-optimal performance of solutions in the set: a majority of the solutions does not perform Pareto-optimal any more if assessed using a different transportation demand. On the other hand, the loss in attainment of objective values is small: although performance is no longer optimal in most cases, loss in performance is limited.

Furthermore, the decision variables only show slightly more differences when optimised for a different demand than differences that are caused by the random nature of the used optimisation heuristic. This indicates that the resulting decision variables are not strongly influenced by the exact transportation demand scenario. An important reason for this is that all three tested demand scenarios are realistic forecasts: as long as the structure of the used demand scenarios is similar, the same measures will stay most effective. This is a promising observation when using the modelling framework in planning practice, because making policies for the future inherently involves uncertainty in autonomous developments like the demand for transportation.

Care should be taken when investing in new train stations and P&R facilities, because these measures are hard to reverse and in some cases are less effective under different demand circumstances. The opposite is true for easy-to-reverse measures, like bus frequencies and to a lesser extent train frequencies and express train stations: a varying performance is not a problem. When such a measure appears to be performing badly due to changing transportation demand, it can just be reversed.

Concluding, the influence of demand uncertainty on the optimisation outcomes is not big, as long as the structure of the used demand scenarios is similar: the decisions to be taken are comparable and the objective values achieved are still near-optimal. In other words, the measures that are optimal in one future situation, are near-optimal in a different future situation. Also, the impact of omitting one objective function from the analysis is small in this case, because the chosen objective (CE) is largely in line with other objectives.

Chapter 7: Conclusions and recommendations

In Chapter 1 the context, research objective, research approach and scope were presented. In Chapter 2 a review was given of earlier research on the network design problem (NDP) and on modelling the behaviour of travellers in a multimodal network. Together with the scope defined in Chapter 1, this led to the formulation of a mathematical optimisation problem in Chapter 3, comprising objective functions, decision variables and a lower-level behavioural model. In Chapter 4 the solution approach was described and tested, comprising a heuristic that enables approximating the Pareto front. In Chapter 5 the solution approach was applied to the case study of the Randstad area, which led to a Pareto set. Several methods to derive information from the Pareto set and to reduce the number of solutions in the set were presented in the same chapter. Finally, the long-term robustness was assessed in Chapter 6, by investigating to what extent the optimisation results are influenced by different scenarios for future transportation demand or by a different number of objectives.

This final chapter presents the conclusions and recommendations. It has three sections. Firstly, conclusions are drawn, which relate to the research objective and research questions. This includes the conclusions regarding the case study in the Randstad, which are also relevant for practice. Secondly, the implications for the research field and for practice are given. Thirdly, directions for further research follow from the results and conclusions.

7.1. Conclusions

This research contributes to the goal to attain a more sustainable transportation system in the Randstad, which is also the goal of the Sustainable Accessibility of the Randstad research programme. The objective of this research was *to determine which (re)designs of the multimodal passenger transportation network in the Randstad contribute best to improving various aspects of sustainability and to provide insight into the extent to which these aspects are improved and into how scores on these aspects and designs relate.*

From a series of interviews with stakeholders a multi-objective optimisation problem is defined. A genetic algorithm is used to solve this problem to identify the resolution of the selected measures related to multimodal trip making. The corresponding network designs are also determined. Methods to analyse the resulting Pareto set and to help the decision maker to choose a final solution for implementation are developed and applied. This provides the desired insight into how objectives relate and what kind of measures can be used to optimise certain objectives in the Randstad case study. The main insights that have been provided are:

- It is possible to improve all aspects of sustainability simultaneously in comparison to the current transportation network: the current design is not part of the Pareto set.
- Improving travel time further can be done cost-efficiently, but reducing CO_2 emissions is expensive when using measures related to multimodal trip making.
- Increasing frequencies appears to be more effective to improve sustainability than introducing P&R facilities and train stations.
- The influence of demand uncertainty on the optimisation outcomes is not large, as long as the structure of the used demand scenarios is similar.

To reach these main insights, several steps were taken by answering the research questions that were formulated in Chapter 1. In the remainder of this section conclusions are drawn considering each of these research questions. First, the definition of the optimisation problem is summarised and reflected upon. Second, conclusions related to the solution algorithm are drawn. Third, conclusions from the robustness analysis are given. Fourth, conclusions are drawn concerning methodology development on how the Pareto set can become more useful as decision support information. Finally, conclusions are drawn specifically for the case study in the Randstad, which can be used as decision support information. These findings are particularly interesting for practitioners.

7.1.1. Problem formulation

Improving the public transport (PT) network is chosen as a solution direction, because it is assumed that a shift from private to public modes of transportation can improve social and environmental sustainability and possibly at the same time improve economic sustainability. The polycentric nature of the Randstad area (the main urban area in The Netherlands) and the expectation that large-scale investments in infrastructure are not feasible in the near future are main reasons to focus on facilitating multimodal trips. In such a multimodal trip a traveller combines private modes (i.e. car or bicycle) and public modes (i.e. train, metro, tram, bus). One important issue in facilitating multimodal trip making is to improve the transfer between modes. This can be achieved by improving existing transfer locations or creating new ones and by creating a PT network that is attractive enough. To that end, the current transportation

network is redesigned. To investigate which design serves the objectives best, a mathematical optimisation approach is adopted, formulated as a multi-objective (MO) NDP. The MO-NDP enables identification of the dependencies and trade-offs between objectives. It is defined as a bi-level mathematical optimisation problem, taking into account the behavioural response of travellers on measures taken. It also offers the possibility to investigate the resolution of the different decision variables for the various objectives. A multimodal passenger network is studied, so the problem is an MO, multimodal passenger NDP, which is mathematically defined by choosing objectives, decision variables and behavioural models that make the lower-level optimisation problem operational.

The single most important objective for each of the three main aspects of sustainability is included in the definition of the NDP. This definition is based on interviews and workshops with policy officers. Total travel time represents accessibility (economic sustainability), which is suitable because fixed demand is assumed. Urban space used by parking represents social sustainability. CO_2 emissions represent climate impact and are directly related to energy use (environmental sustainability). Finally, operating deficit represents cost-efficiency and is included as the fourth objective. This is the difference between PT operating costs (including depreciation of investments in park-and-ride facilities) and operating revenues and represents the viewpoint of the regional government, who are responsible for these costs. Investments in rail infrastructure are excluded, because these are paid by the national government. The stakeholders preferred costs to be included as an objective rather than as a budget constraint, to reveal trade-off information, for example to identify the marginal costs needed to reach additional travel time savings.

Explicitly taking into account the most important objective per main aspect of sustainability implies that other objectives related to sustainability are not taken into account. However, it is expected that also other sustainability objectives are improved, like noise hindrance, traffic safety and emissions of local pollutants, like NO_X, SO_2, CO, particulate matter and hydrocarbons. The reason why this is expected is because just like CO_2 emissions, these objectives relate to the number of vehicle kilometres, yet with different emission factors. An important difference between CO_2 emissions and these other objectives is that the location of the emission / exposure is relevant: emission / exposure in populated areas is more harmful than in rural areas. This is related to the objective of urban space used by parking, because the number of car trips to and from highly urban zones has a relation with the number of car kilometres driven in highly populated areas. Besides, local pollutants are expected to be less relevant in the future, because their emissions are expected to be strongly reduced by technological developments.

The decision variables are also defined in conjunction with the stakeholders. All potentially successful locations for new park-and-ride (P&R) facilities and new local train stations in the study area are included as decision variables. Also possible changes in service lines are included as decision variables in the form of express train status of stations, frequencies of major bus lines and frequencies of local train lines. All these measures involve operating costs, which affect the objective operating deficit. These measures are small-scale measures compared to the entire multimodal network that is modelled, which contains for example fixed express train services and a fixed car network. This results in relatively small changes in objective values, because the objective values are determined for the entire network.

Redesigning the multimodal network (i.e. implementing measures that facilitate multimodal trip making) causes changes in the costs of mode and route choice alternatives. To determine the behavioural response of travellers to these changes, a transportation model is used that includes mode and route choice. It turned out that existing modelling components could be combined into a lower-level model that is easy to operate and has a low enough computation time. This model includes traditional flow-dependent static user equilibrium for car, taking congestion in the car network into account, combined with static PT modelling with multiple access and egress modes and multiple routing, to take various mode chains into account. These mode chains (for example using car or bicycle to access a train station) are choice options for modal split, next to the traditional unimodal choice options (car-only trips and PT trips that only involve walking for access and egress).

7.1.2. Solution algorithms

The MO-NDP is proven to be NP-hard, so heuristics are needed to approximate the solution of the problem. Various studies showed that (specific variants of) genetic algorithms are suited to solve MO or multimodal NDPs, because these algorithms have a low risk of ending up in a local minimum, do not require the calculation of a gradient and are able to produce a diverse Pareto set. A solution method that only uses a limited number of function evaluations to produce near-optimal solutions is selected, because in both case studies the lower-level model involves large computation times. The NSGAII algorithm is a popular MOEA that was previously successfully used to solve MO-NDPs, while ε-NSGAII, developed some years later, proved to outperform NSGAII for other types of problems. Therefore, the performance of the MOEAs NSGAII and ε-NSGAII is compared when applied to the multimodal passenger transportation NDP. The experiments confirm earlier findings that ε-NSGAII outperforms NSGAII for this optimisation problem too, especially in the early stage of running the algorithm. This faster convergence of ε-NSGAII is particularly useful for cases with objective functions that are highly computationally expensive, like in the case study in this research. For this case study it is demonstrated that convergence is almost reached within a given computation time of a maximum of 3000 function evaluations. This shows that the modelling framework is appropriate for case studies with realistic size: the case study covers an entire urban region in the Netherlands, containing all relevant decision variables in the study area. Concluding, the genetic algorithm ε-NSGAII turns out to be a suitable optimisation algorithm to find an approximation of the Pareto set for this case study with realistic size. It is likely that this algorithm will perform well on comparable MO-NDPs with expensive objective function evaluation.

The optimisation outcomes are used to derive information to support decision making, so it is relevant to know if the used solution approach has any implications on how to interpret these outcomes. It is inherent in any genetic algorithm to contain stochastic processes: each run of the algorithm has a random seed as a starting point for the genetic operators. Furthermore, genetic algorithms contain parameters for which values have to be determined: appropriate parameter settings for a specific case study are not known before running the algorithm multiple times. The experiments show that the values of the decision variables are substantially different for multiple runs of the algorithm. As a result the advised decisions are influenced by the stochastic principles of the algorithm. This is related to the flat shapes of the objective functions in the chosen definition of the optimisation problem: similar objective

values can be achieved by different measures in the multimodal network. The experiments confirm this: the differences in the attained objective values among multiple random outcomes of the optimisation process are small, both within runs with the same parameters settings and between runs with different parameter settings. This leads to the conclusion that the method is suitable for the purpose it serves, i.e. providing a set that is sufficiently close to the Pareto optimal set and that can thus be input for decision support information.

7.1.3. Long-term robustness

In the multimodal passenger NDP, transportation demand is an important input variable. Demand forecasting is necessary, because the long lifespan of infrastructure implies that the transportation network is designed for a long period of time. Forecasting is inherently associated with uncertainty, which means it is relevant to know to what extent resulting network designs are robust for demand uncertainty. Pareto sets are constructed for the base demand forecast for 2030 and for two different realistic future developments. As a lower growth scenario the demand forecast for 2020 is used. In addition, a demand forecast based on the alternative spatial planning concept transit-oriented development is used. In this scenario all new residents and jobs in the region until 2030 are developed in the vicinity of train stations, instead of spread out throughout the whole region in the base scenario. This can be seen as an upper scenario in terms of potential PT use.

It turns out that the loss in objective function values is small if a Pareto set optimised for a certain future demand is assessed using one of the two different transportation demand scenarios. Although performance of these designs is no longer Pareto-optimal in most cases, loss in performance is limited. Furthermore, the values of decision variables only show slightly more differences when optimised for a different demand than differences that are caused by the random nature of the used genetic algorithm. This indicates that the resulting values for decision variables are not strongly influenced by transportation demand. These two observations underpin that the influence of demand uncertainty on the optimisation outcomes is not big: the decisions to be taken are comparable and the objective values achieved are still near-optimal. In other words, measures that are optimal in one future situation, are nearly optimal in a different future situation. An important reason for this is that all three tested demand scenarios are realistic forecasts: as long as the structure of the used demand scenarios is similar, the same measures will stay most effective. This is a promising observation when using the modelling framework in planning practice, because making policies for the future inherently involves uncertainty in autonomous developments like the demand for transportation. When a distinction is made between measures that are easy to reverse, easy to reverse with risk of disinvestments or hard to reverse, no clear differences can be observed between these three classes. In all classes measures are observed that perform similarly in the three demand scenarios, but measures that have a varying performance throughout the three demand scenarios are also observed. Care should be taken when implementing measures that are hard to reverse and have a varying performance over the demand scenarios. On the other hand, for easy-to-reverse measures a varying performance is not a problem: when such a measure appears to be performing badly due to changing transportation demand, it can just be reversed.

Although the experiments were only conducted for one case study, similar results are expected for other NDPs as long as the structure of the demand does not vary strongly: in NDPs those measures that influence the largest flows generally are the most effective.

7.1.4. Methods for decision support

Interviews with policy officers showed that a Pareto set is valued positively as input for decision support information, because it enables an interactive process, where the consequences of certain choices can be demonstrated directly. Several visualisation methods and analytical methods turn out to make the Pareto set easier to understand and more useful to guide the decision maker to come to a final solution for implementation. The first method sets bounds to one or more objective values or selects a decision variable that is politically desirable, resulting in a step-by-step pruning process. This has showed to be very helpful to finally select one solution for implementation. In this process, the parallel coordinate plot helps to get an impression of the data and to visualise the position of one or more solutions in the entire set. The second method determines trade-off values between objectives, where scatter plots help to visualise trade-offs between 2 objectives (or a decision map to visualise 3 objectives). These trade-offs are seen as marginal costs, and therefore a decision maker can easily assess whether additional investments are worthwhile. The third method identifies the min-max solution: for each solution the least scoring objective is leading to select the best solution from all Pareto solutions. This is useful in the search for a compromise solution in the political context of decision making, where each political party values each objective differently. When in negotiation during the political debate, no party will accept a solution whose objective is not reached at all.

Policy makers also indicated that their task becomes easier if the number of solutions is reduced systematically by the modeller / researcher. This is done by systematic pruning, which results in a comprehensive subset of Pareto solutions that represents the whole Pareto set. An existing pruning method (PIT-filter) is enhanced. The enhanced method (called PIT-2) outperforms other known methods using indicators that make 'systematic' operational, in terms of Pareto set attainment and spread. Application of this method to the case study provides an overview of 15 solutions as a representative subset of the entire Pareto set of 210 solutions, which is regarded as a reasonable number by the interviewed policy officers.

7.1.5. The case study in the Randstad area

The goal to (re)design the multimodal passenger transportation network in the Randstad to improve various aspects of sustainability is reached by solving an MO-NDP. The outcome of this optimisation process is a Pareto set, from which information on the inherent structure of the problem and its solutions is derived, using several methods for decision support. This results in case-specific insight into which (re)designs of the multimodal passenger transportation network in the Amsterdam metropolitan area best improve sustainability. The resolving power of facilitating multimodal trip making is presented, which is the main research question of this research.

Objective values

The measures (all related to multimodal trip making) manage to achieve a reduction for all four objectives simultaneously in the case study compared to the base situation. The relative improvements with respect to the total values of the sustainability objectives in the base network can be at least 0.4% for each objective. In other words, one specific design (i.e. Pareto solution) shows a reduction of at least 0.4% for total travel time (TTT), urban space used (USU) and CO_2 emissions (CE), with a 0.4% reduction in costs. In total, 47 out of 210 Pareto-optimal solutions have better values for all objectives than the base solution (i.e. dominate the base solution). Consequently, there is a range of solutions to choose from, also if only a solution with improvements in all objectives is accepted by the decision maker. For the individual objectives, larger improvements are possible: CE can be improved by 0.45%, TTT can be improved by 0.79% and USU by 1.5%, but not without deteriorating at least one other objective; in some cases even a deterioration with respect to the base network. Note that the objectives are network-wide measures: if TTT decreases, on average travel times decrease, but for individual travellers the travel time may increase. The contribution of the measures is small in relative terms, because the objective functions are defined as the sum over the entire multimodal transportation network in the study area, implying large aggregate values. The measures at transfer locations and in the PT network typically only influence a limited part of the network (the car network is hardly influenced). In absolute terms these potential gains per objective are considerable: every AM peak almost 4000 hours of travel time, more than 2000 parked cars or more than 12 tons of CO_2 emissions (equivalent to the daily direct CO_2 emissions of more than 500 Dutch households) can be saved.

Relations between objectives

The objectives TTT, USU and CE are mainly in line with each other and opposed to operating deficit (OpD). However, there is no single solution that is optimal for all three objectives simultaneously: pairs of trade-off solutions still exist. An average trade-off value of €6.90 euro/hour is found for TTT and OpD, which means that on average one hour of travel time is saved by spending €6.90 on operating PT. This value is smaller than the value of time used in Dutch cost-benefit analyses, so from the viewpoint of the regional government cost-efficient measures exist within the selection of measures considered. It should be noted that investments in infrastructure (paid by the national government) are not included in the analysis, so whether the measures are cost-efficient from a societal point of view is unsure. To save one kilo of CE, on average €1.10 needs to be spent. This is much more than alternative measures to reduce the transportation system's CO_2 emissions, like tighter CO_2 emission standards for cars (€0.11 per kilo) or stimulating electric cars (€0.35 per kilo). It should be noted that this figure neglects effects on other objectives, which on average are also positive in the case of the multimodal NDP. When effects of additional expenditures on operating PT on TTT, USU and CE are considered simultaneously, each additional euro results on average in a marginal gain of 0.098 hours of TTT, 0.067 parking spaces of USU and 0.45 kilos of CE. This insight illustrates the value of explicitly taking into account multiple objectives: if, for example, only travel time had been included in the analysis, the positive effects on other objectives would have been neglected when making a choice for a final network design.

When only Pareto solutions with low CE are selected (i.e. with less CE than in the base situation), still almost the entire Pareto front for OpD and TTT is available to choose from. This provides the insight that when only these three objectives are considered, there is no reason to choose a solution with higher CE. Finally, the best compromise solution (i.e. min-max solution) over all four objectives is in the best 30% of the range covered by each objective, so with equal weights for the normalised objective functions it is possible to satisfy all four objectives to a large extent simultaneously.

Relations between measures and objectives

Concerning the relation between individual decision variables and each of the four objectives, earlier findings are confirmed that in the Dutch context opening P&R facilities leads to more CE, because not only former car trips are attracted to P&R, but also former PT-only trips, possibly with bicycle as access or egress modes. Furthermore, increasing train frequencies and to a smaller extent bus frequencies in general reduces TTT, USU and CE, but increases OpD. Increasing frequencies in general appears to be more effective than P&R facilities and new train stations, indicating that stimulating multimodal trips is done best by improving the quality of the PT network itself, rather than by introducing additional transfer locations. Another reason to prefer frequency increases over new train stations and P&R facilities is that the first are easier to reverse in case in the future circumstances occur that differ from the initial predictions.

Relations between measures

In Pareto-optimal network designs with low frequencies on local train lines, which is beneficial for OpD, usually more stations acquire express train status to improve other objectives, by compensating for the low frequency at these stations in a cost-efficient way. Train routes and bus routes turn out to be complementary rather than competitive, because high frequencies occur simultaneously on train and on bus lines in Pareto-optimal network designs. This means that for the network in the case study, bus lines and train lines can strengthen each other. When looking at individual measures, in the case study in the Randstad from all 37 possible measures in the case study only one measure (a new train station at Nieuw Sloten in northern direction) is identified as 'always regret', i.e. it does not occur in any Pareto solution). No 'no regret' measures (occurring in all Pareto solutions) are found. This indicates that whether a measure performs optimally depends on the objectives that are preferred.

Performance of other types of measures

Redesigning the multimodal network enables simultaneous improvement for all considered aspects of sustainability, due to the adopted MO optimisation approach. On the other hand, the contribution of this type of measures to for example CO_2 emission reduction targets is limited, because the high-emission alternative (i.e. the car-only trip) remains attractive for the traveller. Therefore the impact of network-wide measures that make the car as an alternative less attractive is demonstrated, although these measures are probably difficult to implement in practice. The same transportation model is used, despite this model excluding behavioural responses that become relevant for this kind of measures like elastic transportation demand. This shows that transit-oriented development achieves both larger improvements in

environmental sustainability and a small decrease of delay times. Road pricing achieves much larger improvements in environmental sustainability and also improves social sustainability, because on the entire network the car alternative becomes less attractive. This is also the reason why average travel times increase due to road pricing in a multimodal context, because former car users switch to slower PT alternatives. On the other hand, travel times on the road network decrease due to less congestion.

7.2.   Implications

The implications (or lessons learnt) relate to the challenges and contributions that were identified in Chapter 1, covering several aspects of the MO-NDP. First, the implications for the academic world are given. Second, the implications for policy-making practice are highlighted, which are useful for policy officers in the field of transportation and decision makers at related authorities.

7.2.1.   Implications for the research field

The Pareto set has traditionally been seen as the final product of MO optimisation. However, to be useful as decision-support information, further analysis is needed. Although within the MO optimisation research community this issue is widely recognised to be urgent, attention to it in literature has been limited until now. In this research first steps are made to carry out these analyses and to define them formally. For future MO optimisation studies these definitions can directly be used to present the Pareto results in a more comprehensive way.

The solution direction to stimulate multimodal trip making only leads to relatively small improvements for the sustainability objectives. However, the related measures involve limited costs as well, and therefore the solution direction is promising when looking for cost-efficient measures to improve sustainability of the transportation system. For future NDP studies that involve multimodal transportation networks, this research provides an operational lower-level model that explicitly takes into account mode chains during mode and route choice, resulting in accurate network flows. The model has a low computation time and is available in an easy-to-use software environment.

NSGAII is currently the standard MO algorithm used to solve engineering problems. This research provides additional evidence that other, newer algorithms like ε-NSGAII perform better and are ready to be applied in practical optimisation problems. In current literature most attention goes to comparing the objective value attainment of several variants of the algorithm when comparing the performance of algorithms. However, this research demonstrated that random processes in NSGAII have effects on the optimisation outcome as well, especially on the values of decision variables. When sensitivity of algorithms to input data or to these random processes is investigated, it is also important to compare results with respect to the decisions to be made. The indicators developed in this research (average distance between Pareto solutions, average distance to the nearest Pareto solution and difference between fractions of nonzero decision variables in the Pareto sets) are an important step to pay more attention to differences between decision variable values of solutions in the Pareto sets.

Transportation demand is important input for the MO-NDP. The effects of uncertainty in demand input on the outcomes of an optimisation problem are analysed. This provides first insights into what happens with a Pareto set when different circumstances occur than it was optimised for. Although the performance of most network designs is no longer Pareto-optimal in most cases when a different demand occurs, loss in performance is limited, as long as the structure of the used demand scenarios is similar. The indicators developed in this research provide the necessary methods to cope with the additional difficulty to compare pairs of sets of solutions instead of just pairs of solutions.

7.2.2. Implications for practice

The results obtained in this thesis can be translated into practical implications for policy officers in the field of transportation and decision makers at related authorities. The modelling framework could be used in an orientating planning phase, for example when deciding on a long-term policy document. On the other hand, the framework comes with concrete measures and with information on to what extent these measures contribute to achieving policy objectives. This helps policy makers to set realistic goals and to start putting the policy document into operation. The modelling framework can help to bridge the gap between visionary policy documents and concrete measures in the transportation network. It contributes to attaining better network designs, when used to support a planning process of a government or transport authority. This would change the process to come to a more sustainable transportation system.

Traditionally a set of network design solutions is formulated by expert judgement and only these solutions are assessed using sustainability objectives. From this limited set of solutions the best performing solutions is selected, but there is no guarantee that it performs Pareto-optimally. It is likely that better solutions are overlooked, because the behavioural response of travellers (represented by the lower-level model) is complicated and therefore cannot entirely be included in the expert judgement. The optimisation approach presented in this research overcomes this issue, because its outcome is a Pareto set of network designs. This Pareto set contains the best sustainability scores instead of only sustainability scores for the limited set of evaluated solutions. Expert judgement is still used in an earlier phase to set the objectives, decision variables and constraints, but the influence of expert judgement on the outcome is ruled out, preventing any influence of implicit preferences or mistakes of the experts.

Based on the Pareto set, information on the problem is revealed, like interdependencies and trade-off values between objectives and relations between measures and objectives. To reveal this information, proper methods are essential to guarantee an easy and correct interpretation of the results by the decision maker. Such methods to analyse the resulting Pareto set are developed and applied in this research. These methods can be included in an interactive decision support tool, which can be used in any practical situation where many options are considered and each option has scores for multiple objectives. In that case, a decision maker receives direct feedback concerning the influence of specific measures on the objective values. Based on this information, a decision maker can choose a final network design for implementation, depending on the priorities given to the objectives.

When the modelling framework is used in practice, a (transportation) model to represent the behavioural response of travellers should be available / adapted / built, in such a way that it suits the area of study and the measures included in the definition of the decision variables. The computation time of this model determines the computation time needed for function evaluation, so it directly influences the computation time of the entire optimisation process. Parallel computing can be used to reduce computation time if it turns out to be too large.

Furthermore, communication between the researcher / modeller and the policy officer and / or decision maker is essential when the modelling framework is used in practice. Before the optimisation process takes place, the problem has to be defined. This includes definition of objectives, definition of decision variables, including their possible values, and choosing assumptions regarding input, for example transportation demand. If needed, constraints can be defined, too. After the optimisation process the outcomes can be analysed in detail by using the methods in this research. It would be effective if these methods could be applied directly during a meeting to interactively cope with feedback from the decision makers. Planning the policy making process has consequences for the computation time that is available for the optimisation process and therefore for the quality of the optimisation outcome. It may be desirable to repeat the optimisation process with a different definition, for example if the resolving power of the combined measures is not enough to meet a policy goal that was set, asking for different (more rigorous) measures (i.e. decision variables).

Based on the case study results the following implications can be given concerning the resolving power measures. Opening P&R facilities leads to lower TTT, but also to more CE, because not only trips that were previously made by car are attracted to P&R, but also trips that were previously made solely by PT (possibly with bicycle as access or egress mode). Furthermore, increasing train frequencies and to a smaller extent bus frequencies in general reduces TTT, USU and CE, but increases OpD. Increasing frequencies appears to have a larger effect on the sustainability objectives than opening new P&R facilities and new train stations, indicating that multimodal trips can be stimulated better by improving the quality of the PT network itself than by introducing additional transfer locations. Another reason to prefer frequency increases over new transfer locations is that frequency increases are easier to reverse, making these measures easier to adjust to changing circumstances in the future. Although the optimisation outcome depends on the demand forecast that is used as input, the sensitivity of the result for the exact demand forecast used is small, as long as the forecast does not contain a substantial change in the structure of the demand matrix.

In most countries, a lot of stakeholders are involved in operating PT. In the Netherlands, a main distinction can be made between the national level (the national rail network) and the regional level (bus, tram, metro and local train lines). On both levels, a governmental body is responsible for operating PT and contracts a PT operating company to actually operate PT. For the regional level, The Netherlands is divided in approximately 30 regions. In principle, it would be best to optimise the PT network as one integrated network, not bothering about levels or regions. In practice however, local and national interests often are not in line with each other, leading to the situation that each level and region optimises its own network, using its own objectives. This is often suboptimal: the case study for example showed that bus and train lines can strengthen each other. Another mechanism is the way the PT network is financed: regional governments are responsible for the expenses on operating regional PT, but

investments in new rail infrastructure are the responsibility of the national government. Therefore the regional government may see this infrastructure as 'free' in their decision making process, which may lead to suboptimal decisions for the system as a whole. The modelling framework in this research can be applied at any level and in any region. Furthermore, it is possible to apply the framework before tendering the contract (from the perspective of the government) or during a contract (from the perspective of the PT operator). However, the overall, network-wide optimisation result will be better for society as a whole if the problem is defined as an integrated NDP. Such an NDP includes both levels of the PT network and possibly includes more than one region.

In this case study the NDP is limited to decision variables related to multimodal trip making. The best scores that are found show how much improvement is possible with respect to the current multimodal transportation network. This resolving power of the measures in a study area can be seen as an indicator of effectiveness, where trade-off values can be seen as indicators for efficiency. The framework could therefore be used to compare urban regions with each other using these indicators of effectiveness and efficiency, for example by a higher level government to allocate budgets to these urban regions.

This research demonstrates that measures related to multimodal trip making can contribute to reaching sustainability goals by making travel alternatives that include PT more attractive. On the other hand, the contribution to for example CO_2 emission reduction targets of this type of measures is limited (a reduction of less than 1% could be achieved in the case study), because the high-emission alternative (i.e. the car-only trip) remains attractive for the traveller. Especially in CO_2 emissions, on the long term large reductions are needed: a 40% reduction (compared to 1990) for 2030 and an 80-95% reduction for 2050. Although it is likely that technological improvements will strongly improve the environmental performance of cars in the future, a car trip is still likely to have a higher energy use per traveller than a PT alternative, because PT vehicles will also be improved by technological developments. In addition to this, if renewable energy is used a PT alternative will still be preferred from this perspective, because also generation of renewable energy has negative impacts on society (for example land use by wind turbines). If this leads to the conclusion that it is desirable to reduce car use more strongly, it is necessary to make the car alternative less attractive (next to making alternatives that include PT more attractive). This can be made operational in several forms (for example road pricing, parking pricing, fuel tax or closure of car infrastructure).

7.3.   Recommendations for further research

Most recommendations for further research in this section relate to a specific component of the modelling framework. A majority of the recommendations involve making the components more elaborate, but in some cases it might be interesting to investigate whether making the components more concise is beneficial. Additional recommendations are made regarding the definition of the optimisation problem and the process to come to the optimisation results.

Lower-level transportation model

The modelling techniques used to model combined mode and route choice may be improved further to better handle the behavioural complexity of multimodal trip making, resulting in more accurate network flows. A possible research direction is the super-network concept (Sheffi, 1985), where all constraints on mode composition are dropped by allowing for the full diversity of multimodal trip chains. Recently, application of this concept to real size networks has become possible (Van Eck et al., 2014). Apart from the complexity related to trip chains, attributes in the generalised costs function may be extended. In addition to the attributes considered in this research (travel time, waiting time, number of transfers and distance based fare), the comfort level of various types of PT services and transfers plays a role. This for example depends on leg space, availability of Wi-Fi, availability of real-time travel information or waiting facilities at stops. This has a relation with vehicle occupancy: taking into account the capacity of PT vehicles during passenger assignment would improve the lower-level model results, especially in network designs with low frequency services on busy corridors. Both comfort and crowding may have an influence on the value of time. Consequently, the generalised costs function changes and insight is needed into how these additional attributes influence the mode choice model: how the new generalised costs relate to the generalised costs of the non-PT travel alternatives. A completely different approach is to include (daily) activity patterns in the models, taking aspects like vehicle availability and combinations of multiple trips in a tour into account. In that case, a choice for a certain mode for a morning trip directly influences the mode choice options for the other trips to be made in the tour. In fact, the mode choice is not done on a trip level, but on a tour level. Such a tour-based approach would be a useful extension as it would be possible to track vehicle availability along the whole tour.

Apart from these possible improvements in modelling route and mode choice, the behavioural response may be extended, leaving the assumption of a fixed total transportation demand. Extending the lower-level model with departure time choice modelling, destination choice modelling and / or elastic demand modelling would improve the realism of the model. This is especially relevant when the NDP is extended including more rigorous measures like new rail infrastructure or road pricing.

Although it is expected that a more elaborate lower-level model leads to better optimisation results, this has not been established. It might even be that simpler models in the end lead to better optimisation results, because a lower computation time for the lower-level model enables more function evaluations during optimisation. Therefore, more research is needed to indicate whether the conceptual limitations of the used model are negligible at the full network scale (leading to the same optimal network designs) or systematically bias the assignment results (leading to different network designs being optimal). This also has a relation with the translation from assignment results to objectives: decisions are based on objective values, which are derived network characteristics, such as travel time, operating deficit, car use in urban areas and CO_2 emissions. Incorrect prediction of network use does not necessarily mean that the aggregated objective value is wrong, for example because the location of CO_2 emissions is not relevant. To investigate whether more elaborate or simpler behavioural response models are desired, it is recommended for future research to apply various lower-level models in the modelling framework and compare the output in terms of

network designs, objective values and computation times. A hybrid approach could be to apply a simple lower-level model first to achieve first improvements using limited computation time. The outcome of this first stage can be used as a starting point for further, more accurate optimisation using a more elaborate lower-level model.

Solution algorithm in the upper level

When using the solution algorithm ε-NSGAII, ε-settings can be used to put emphasis on specific objectives, so it would be interesting to investigate the influence of different settings of ε on the results. These settings also strongly influence computation times, since convergence of the algorithm depends on them. Furthermore, it would be good to do more similar tests on different real-life design problems, because results of evolutionary algorithms are always case-specific and applications of ε-NSGAII are scarce in literature, especially in the field of transportation.

Although ε-NSGAII performed in a satisfactory way in the large case study, it is of interest whether other heuristics can be found with better performance than ε-NSGAII, for example the reference point based algorithm NSGAIII (Deb and Jain, 2014), which performs well in cases with four or more objectives. Another example is the self-adaptive algorithm named Borg (Hadka and Reed, 2014), which explicitly provides the opportunity of parallelisation to decrease computation time. ε-NSGAII also has the possibility of parallelisation, but the Borg algorithm explicitly benefits from parallelisation by introducing the concepts of master and slave and multi-master optimisation.

Influence of stochastic processes in the solution approach

The fact that genetic algorithms include stochastic processes is often easily dealt with by investigating average values, standard deviation values or other statistical metrics. However, when used in real-world decision making in the field of transportation network design, these algorithms can only be run a few times or even only once due to high computation times in relation to limited time available for the decision process. The advised decisions turn out to be influenced by the stochastic processes in the used genetic algorithm, because different solutions were found in the Pareto sets with comparable objective values. When taking decisions in other case studies based on Pareto sets that are constructed by such an algorithm, knowledge is needed on the sensitivity of decision variables with respect to the objective values, to identify important and less important decision variables. Research on this important aspect of genetic algorithms is needed for more case studies to be able to draw more general conclusions. The indicators developed in this research can be used to analyse these differences between decision variable values of solutions in the Pareto sets. Development of more indicators to compare decision variable values of Pareto sets is an interesting topic for further research.

Robust optimisation

The robustness of the optimisation result with respect to travel demand is investigated by comparing optimisation outcomes for different demand scenarios as input. However, if more computational power is available in the future, an interesting research direction will be to include several demand inputs in the lower-level model. These demand inputs could involve

several scenarios for the future demand forecast or could involve day-to-day dynamics to model an entire month or year, instead of only an average working day. This enables determining average values and other statistical measures, like standard deviations or percentile values for each objective over all demand scenarios. These values can be used directly to guide the optimisation process towards network designs that are robust for uncertainty in demand input. Another possibility is to define a separate objective for each demand scenario in the MO optimisation problem, enabling an explicit trade-off between being optimal for several future situations. This however results in many more objectives to be included in the optimisation problem: the number of objectives will increase by a factor equal to the number of scenarios.

Methods for decision support

The PIT-2 filter turns out to be the preferred method to systematically prune the Pareto set to come to a comprehensive subset of Pareto solutions that is representative for the whole Pareto set. The conceptual improvements made the PIT-filter slightly less effective in pruning the set, but the new filter is recommended over the old one based on the case study results. Moreover, it has a superior theoretical foundation: it assures that no significant trade-off exists between a selected and a rejected solution in the 3 or more dimensional case. The second conceptual change (PIT-3) decreases the effectiveness of the filter strongly, so based on this case study it is not recommended to apply that filter any further. It would be useful to continue testing the performance of these filters on other real-world Pareto sets from MO optimisation problems to verify whether these results are more generally applicable. Also the other defined methods to analyse the Pareto set need more tests and applications to other MO optimisation problems. Making the Pareto set more useful to support decision making is widely recognised to be urgent in the MO optimisation research community. However, attention to it in literature has been limited until now, so future research into more and further improved methods is favourable.

Methods for decision support enable a decision maker to make a choice. The methods do not actually make or predict the choice. The field of choice modelling aims to predict choices that are made. Behavioural scientists study how and why people make choices. The focus of choice modellers and behavioural scientists is usually on large groups that together make large numbers of choices, for example travellers or consumers. Although this is different from the choices made during a policy-making process, the theories and methods available in the field of choice modelling may be applicable. In that case unlikely options can be removed from the choice set, so that the decision maker is not distracted by choices he or she is not likely to make anyway. When applied in a real-world decision making process, the actual choices made may be compared with the choices predicted by the models. Concluding, combining the field of choice modelling with research in MO optimisation is an interesting research direction to provide information for decision support.

Interaction with decision makers

MO optimisation is well-known in academic literature, but is mainly described theoretically and is rarely applied in practice. On the one hand there is a lack of understanding on the side of the decision makers and they have to get used to a new, different way of working. On the other hand, the aspirations of decision makers are often not known to researchers either.

Therefore researchers are recommended to interact with decision makers from the field, using the visualisation techniques and methods for decision support. This results in more realistic definition of future optimisation problem, as well as in a better understanding of the results of MO optimisation by practitioners.

Different study area

In general a vast number of stakeholders are involved in operating PT. The modelling framework in this research is applied using the interests of a single stakeholder, which may be different when another stakeholder is consulted. Moreover, for society as a whole it would be best to optimise the PT network as one integrated network. This relates to the costs function, which only included the costs relevant for the regional government, neglecting costs to be paid by the national government. Defining the optimisation problem in such a way that costs and revenues for all relevant governments and other parties in the region under study are taken into account will lead to a solution that is optimal for society as a whole. Determining the borders of the study area will stay an issue here, especially in highly populated areas like the Netherlands and surrounding countries.

Different objectives

In the current definition of the optimisation problem, a selection of four objectives is included. It is expected that other objectives related to sustainability, like noise hindrance, traffic safety and emissions of local pollutants, are roughly in line with the chosen objectives CO_2 emissions and urban space used. However, this is not tested and would be interesting to analyse in future research.

Furthermore, in the current definition the four objectives are defined as network-wide values, i.e. the distribution of effects over the study area is not accounted for. A relevant question is whether the benefits of the network designs are spread evenly over de region or are concentrated at certain locations. This issue is related to equity: the question is whether the benefits of public expenses are fairly distributed over society, for example geographically, but also over users of various modes of transportation.

Technological developments

The optimisation problem is defined in such a way that it complies with the current insights in future developments. The effects of measures, i.e. the way the objective values are determined, may change due to technological developments, like hydrogen cars or electric cars. The availability of new modes can also be included, like the introduction of the electric bicycle, the possibility for the traveller to take a bicycle into the PT vehicle or demand responsive PT systems, e.g. sharing automated vehicles. This may lead to different network designs, and it is relevant to known to what extent these developments influence the optimisation result.

Different measures

The best scores found in the case study after redesigning the multimodal network show limited relative improvement, while road pricing or transit-oriented development can achieve larger improvements. If larger improvements are aspired, a favourable extension of the NDP

is to include other types of measures in the multimodal NDP, including their costs, both for the government and for society. This would result in the combination of measures related to multimodal trip making with other measures in one integral multimodal NDP. Also for this optimisation problem with changed boundaries the effectiveness and efficiency of the corresponding measures could be determined and compared to the values of the measures in the original NDP. These measures could be pricing of roads, parking or PT or large-scale investments in fast PT systems like high-speed rail or spatial developments. This would require extension of the behavioural response model: taking other types of measures into account in the formulation of the mathematical optimisation problem requires incorporating (more rigorous) behavioural responses than mode / route choice only, like the effects on trip distribution and trip generation.

Whether the definition of the NDP is extended or not is a result of the political process. Cars play a major role in the current transportation system and therefore in the possibility to develop economic activities. Cars also generate an important share of the government's tax revenues. Economic sustainability is valued strongly in current society, so although desirable from an environmental sustainability point of view, putting restrictions on car use is therefore not politically feasible these days. Whether this is going to change in the future depends on the extent to which politicians keep to their opinions. If car trips are made less attractive and at the same time enough capacity and quality is provided in the PT network (including the connections between the PT network and the bicycle and car networks), improving environmental sustainability may be possible without harming economic sustainability. A shift of economic activities from outskirts of urban areas to city centres and other locations in the vicinity of PT stations may help economic sustainability in a situation where private car use is made less attractive. Such policy may become politically feasible in the future, because the attitude with regard to private car use seems to be changing: starting from 2003, people under 40 gradually travel fewer kilometres as a car driver (KiM, 2014). Furthermore, it is likely that environmental problems will become more urgent in the future.

The solution direction presented in this research is only one step on the long road to a sustainable transportation system in the Randstad. To achieve a sustainable transportation system, cooperation between all stakeholders involved is necessary. Furthermore, this solution direction needs to be combined with other measures or developments, for example technological improvements on the environmental performance of vehicles, a different spatial planning and a different attitude of travellers (and therefore of politicians) towards mobility.

References

Albarello N., A. Guillaume, A. Ohayon (2014) Optimization of a Satellite System using Multi-Objective Evolutionary Algorithms. OPT-i, 1st International Conference on Engineering and Applied Sciences Optimization. M. Papadrakakis, M. G. Karlaftis and N. D. Lagaros. Kos Island, Greece.

Alicino S., M. Vasile (2014) An Evolutionary Approach to the Solution of Multi-Objective Min-Max Problems in Evidence-Based Robust Optimization. IEEE Congress on Evolutionary Computation. Beijing, China, IEEE Press: pp. 1179-1186.

Alt B., U. Weidmann (2011) A stochastic multiple area approach for public transport network design, *Public Transport* 3 (1), pp. 65-87.

Arentze T. A., E. J. Molin (2013) Travelers' preferences in multimodal networks: Design and results of a comprehensive series of choice experiments, *Transportation Research Part A: Policy and Practice* 58, pp. 15-28.

Atashkari K., N. Nariman-Zadeh, A. Pilechi, A. Jamali, X. Yao (2005) Thermodynamic Pareto optimization of turbojet engines using multi-objective genetic algorithms, *International Journal of Thermal Sciences* 44 (11), pp. 1061-1071.

Beirão G., J. A. Sarsfield Cabral (2007) Understanding attitudes towards public transport and private car: A qualitative study, *Transport Policy* 14 (6), pp. 478-489.

Beltran B., S. Carrese, E. Cipriani, M. Petrelli (2009) Transit network design with allocation of green vehicles: A genetic algorithm approach, *Transportation Research Part C: Emerging Technologies* 17 (5), pp. 475-483.

Ben-Akiva M., M. Bierlaire (1999) Discrete choice methods and their applications to short term travel decisions, R. W. Hall, *Handbook of transportation science*, Springer, pp. 5-33.

Benjamins M., C. Lindveld, R. v. Nes (2002) Multimodal travel choice modeling: a super-network approach. 81st Annual Meeting of the Transportation Research Board Washington DC, USA.

Bos I. D., R. Van der Heijden, E. J. Molin, H. J. Timmermans (2004) The choice of park and ride facilities: an analysis using a context-dependent hierarchical choice experiment, *Environment and Planning A* 36 (9), pp. 1673.

Bovy P. H. L., S. Hoogendoorn-Lanser (2005) Modelling route choice behaviour in multi-modal transport networks, *Transportation* 32 (4), pp. 341-368.

Brands T., E. v. Berkum, D. H. van Amelsfort (2009) Optimal toll design in dynamic traffic networks using a pattern search approximation algorithm. *88th Annual Meeting of the Transportation Research Board* Washington DC, USA.

Bunschoten T., E. Molin, R. v. Nes (2013) Tram or bus; does the tram bonus exist? European Transport Conference. Frankfurt, Germany, AET.

Cantarella G. E., A. Vitetta (2006) The multi-criteria road network design problem in an urban area, *Transportation* 33 (6), pp. 567-588.

CBS (2013a) StatLine: Luchtverontreiniging, feitelijke emissies door alle bronnen (in Dutch). Den Haag/Heerlen, CBS.

CBS (2013b) StatLine: Mobiliteit in Nederland; mobiliteitskenmerken en vervoerwijzen, regio's (in Dutch). Den Haag/Heerlen, CBS.

Cepeda M., R. Cominetti, M. Florian (2006) A frequency-based assignment model for congested transit networks with strict capacity constraints: characterization and computation of equilibria, *Transportation Research Part B: Methodological* 40 (6), pp. 437-459.

Chakroborty P. (2003) Genetic Algorithms for Optimal Urban Transit Network Design, *Computer-Aided Civil and Infrastructure Engineering* 18 (3), pp. 184-200.

Chaudhari P. M., R. V. Dharaskar, V. M. Thakare (2010) Computing the Most Significant Solution from Pareto Front obtained in Multi-objective Evolutionary, *International Journal of Advanced Computer Science and Applications* 1(4).

Chen A., J. Kim, S. Lee, Y. Kim (2010) Stochastic multi-objective models for network design problem, *Expert Systems with Applications* 37 (2), pp. 1608-1619.

Cipriani E., S. Gori, M. Petrelli (2012) Transit network design: A procedure and an application to a large urban area, *Transportation Research Part C: Emerging Technologies* 20 (1), pp. 3-14.

Coello Coello C. A. (2006) Evolutionary multi-objective optimization: a historical view of the field, *Computational Intelligence Magazine, IEEE* 1 (1), pp. 28-36.

Compare M., F. Martini, E. Zio (2015) Genetic algorithms for condition-based maintenance optimization under uncertainty, *European Journal of Operational Research* (In press).

CPB, MNP, RPB (2006) Welvaart en Leefomgeving: Een scenariostudie voor Nederland in 2040 (in Dutch). Den Haag/Bilthoven, Centraal Planbureau, Milieu- en Natuurplanbureau en Ruimtelijk Planbureau.

Daniëls B., R. Koelemeijer, F. D. Longa, G. Geilenkirchen, J. v. d. Meulen, K. Smekens, J. v. Stralen (2014) EU-doelen klimaat en energie 2030: Impact op Nederland (in Dutch), ECN and PBL.

Deb K. (2001) *Multiobjective Optimization Using Evolutionary Algorithms*, John Wiley & Sons, Ltd, Chichester.

Deb K. (2003) Unveiling innovative design principles by means of multiple conflicting objectives, *Engineering Optimization* 35 (5), pp. 445-470.

Deb K., S. Bandaru, D. Greiner, A. Gaspar-Cunha, C. C. Tutum (2014) An integrated approach to automated innovization for discovering useful design principles: Case studies from engineering, *Applied Soft Computing* 15 (0), pp. 42-56.

Deb K., H. Jain (2014) An Evolutionary Many-Objective Optimization Algorithm Using Reference-Point-Based Nondominated Sorting Approach, Part I: Solving Problems With Box Constraints, *Evolutionary Computation, IEEE Transactions on* 18 (4), pp. 577-601.

Deb K., M. Mohan, S. Mishra (2005) Evaluating the ε-Domination Based Multi-Objective Evolutionary Algorithm for a Quick Computation of Pareto-Optimal Solutions, *Evolutionary Computation* 13 (4), pp. 501-525.

Deb K., A. Pratap, S. Agarwal, T. Meyarivan (2002) A fast and elitist multiobjective genetic algorithm: NSGA-II, *Evolutionary Computation, IEEE Transactions on* 6 (2), pp. 182-197.

dell'Olio L., J. L. Moura, A. Ibeas (2006) Bi-level mathematical programming model for locating bus stops and optimizing frequencies, *Transportation Research Record: Journal of the Transportation Research Board* 1971 (1), pp. 23-31.

European Commission (2014) "2030 framework for climate and energy policies." from http://ec.europa.eu/clima/policies/2030/index_en.htm.

Fan W., R. Machemehl (2006) Optimal Transit Route Network Design Problem with Variable Transit Demand: Genetic Algorithm Approach, *Journal of Transportation Engineering* 132 (1), pp. 40-51.

Farahani R. Z., E. Miandoabchi, W. Y. Szeto, H. Rashidi (2013) A review of urban transportation network design problems, *European Journal of Operational Research* 229 (2), pp. 281-302.

Fiorenzo-Catalano S. (2007) *Choice set generation in multimodal transport networks*, TU Delft, TRAIL Thesis Series, T2007/6.

Fiorenzo-Catalano S., S. Hoogendoorn-Lanser, R. v. Nes (2003) Choice set composition modelling in multimodal travelling. *10th International Conference on Travel Behaviour Research*, Lucerne, Switserland.

Fisk P. (2010) *People Planet Profit: How to Embrace Sustainability for Innovation and Business Growth*, Kogan Page Publishers.

Florian M., I. Constantin (2012) A note on logit choices in strategy transit assignment, *EURO Journal on Transportation and Logistics* 1 (1-2), pp. 29-46.

Gallo M., B. Montella, L. D'Acierno (2011) The transit network design problem with elastic demand and internalisation of external costs: An application to rail frequency optimisation, *Transportation Research Part C: Emerging Technologies* 19 (6), pp. 1276-1305.

Gao Z., J. Wu, H. Sun (2005) Solution algorithm for the bi-level discrete network design problem, *Transportation Research Part B: Methodological* 39 (6), pp. 479-495.

García R., A. Marín (2002) Parking Capacity and Pricing in Park'n Ride Trips: A Continuous Equilibrium Network Design Problem, *Annals of Operations Research* 116 (1-4), pp. 153-178.

Gini C. (1912) Variabilità e mutabilità (Variability and Mutability), *C. Cuppini, Bologna*, pp. 156.

Guihaire V., J.-K. Hao (2008) Transit network design and scheduling: A global review, *Transportation Research Part A: Policy and Practice* 42 (10), pp. 1251-1273.

Haasnoot M., J. H. Kwakkel, W. E. Walker, J. ter Maat (2013) Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world, *Global environmental change*.

Hadka D., P. Reed (2014) Large-scale parallelization of the Borg multiobjective evolutionary algorithm to enhance the management of complex environmental systems, *Environmental Modelling & Software* (Available online 6 November 2014).

Hamdouch Y., M. Florian, D. W. Hearn, S. Lawphongpanich (2007) Congestion pricing for multi-modal transportation systems, *Transportation Research Part B: Methodological* 41 (3), pp. 275-291.

Hamer P. (2010) Analysing the effectiveness of park and ride as a generator of public transport mode shift, *Road & Transport Research: A Journal of Australian and New Zealand Research and Practice* 19 (1), pp. 51-61.

Hekkenberg M., M. Verdonk (2014) Nationale Energieverkenning 2014. Petten, Energieonderzoek Centrum Nederland.

Hensher D. A., J. King (2001) Parking demand and responsiveness to supply, pricing and location in the Sydney central business district, *Transportation Research Part A: Policy and Practice* 35 (3), pp. 177-196.

Hettenhausen J., A. Lewis, S. Mostaghim (2009) Interactive multi-objective particle swarm optimization with heatmap-visualization-based user interface, *Engineering Optimization* 42 (2), pp. 119-139.

Hoogendoorn-Lanser S., R. van Nes, S. Hoogendoorn (2006) Modeling Transfers in Multimodal Trips: Explaining Correlations, *Transportation Research Record: Journal of the Transportation Research Board* 1985, pp. 144-153.

Huijts N. M. A. (2013) *Sustainable Energy Technology Acceptance: A psychological perspective*, TU Delft, TRAIL Thesis Series, T2013/10.

Huisman C., A. De Jong, C. Van Duin, L. Stoeldraijer (2013) Regionale prognose 2013–2040 (in Dutch), PBL/CBS.

IEA (2010) Energy Technology Perspectives 2010, Scenarios and Strategies to 2050. Paris, International Energy Agency.

INFRAS (2007) Artemis: Assessment and reliability of transport emission models and inventory systems. ROAD EMISSION MODEL – Model Description. Workpackage 1100. Deliverable No. 13, INFRAS.

IPCC (2007) Climate Change 2007: Synthesis Report. Contribution of Working Groups I, II and III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. P. Core Writing Team, R.K and Reisinger, A. Geneva, Switzerland, IPCC.

Johnson D. S., J. K. Lenstra, A. H. G. R. Kan (1978) The complexity of the network design problem, *Networks* 8 (4), pp. 279-285.

Kasprzyk J. R., S. Nataraj, P. M. Reed, R. J. Lempert (2013) Many objective robust decision making for complex environmental systems undergoing change, *Environmental Modelling & Software* 42 (0), pp. 55-71.

Kieft S. (2013) Handboek Venom (in Dutch). The Netherlands, Stadsregio Amsterdam.

KiM (2014) Mobiliteitsbeeld 2014 (in Dutch). Den Haag, Kennisinstituut voor Mobiliteitsbeleid (KiM).

Kollat J. B., P. M. Reed (2006) Comparing state-of-the-art evolutionary multi-objective algorithms for long-term groundwater monitoring design, *Advances in Water Resources* 29 (6), pp. 792-807.

Kuik O., L. Brander, R. S. J. Tol (2009) Marginal abatement costs of greenhouse gas emissions: A meta-analysis, *Energy Policy* 37 (4), pp. 1395-1403.

Kulturel-Konak S., D. Coit, F. Baheranwala (2008) Pruned Pareto-optimal sets for the system redundancy allocation problem based on multiple prioritized objectives, *Journal of Heuristics* 14 (4), pp. 335-357.

Lozano A., G. Storchi (2002) Shortest viable hyperpath in multimodal networks, *Transportation Research Part B: Methodological* 36 (10), pp. 853-874.

Martens K. (2007) Promoting bike-and-ride: The Dutch experience, *Transportation Research Part A: Policy and Practice* 41 (4), pp. 326-338.

Mattson C. A., A. A. Mullur, A. Messac (2004) Smart Pareto filter: obtaining a minimal representation of multiobjective design space, *Engineering Optimization* 36 (6), pp. 721-740.

Mauttone A., M. Urquhart (2009) A multi-objective metaheuristic approach for the Transit Network Design Problem, *Public Transport* 1 (4), pp. 253-273.

Miandoabchi E., R. Farahani, W. Dullaert, W. Y. Szeto (2012) Hybrid Evolutionary Metaheuristics for Concurrent Multi-Objective Design of Urban Road and Public Transit Networks, *Networks and Spatial Economics* 12 (3), pp. 441-480.

Mingardo G. (2013) Transport and environmental effects of rail-based Park and Ride: evidence from the Netherlands, *Journal of Transport Geography* 30 (0), pp. 7-16.

Ministry of Infrastructure and the Environment (2012) Structuurvisie Infrastructuur en Ruimte, Nederland concurrerend, bereikbaar, leefbaar en veilig (in Dutch). Den Haag, Ministry of Infrastructure and the Environment.

Mouter N., J. A. Annema, B. van Wee (2013) Ranking the substantive problems in the Dutch Cost–Benefit Analysis practice, *Transportation Research Part A: Policy and Practice* 49 (0), pp. 241-255.

Nguyen S., S. Pallottino (1989) Hyperpaths and shortest hyperpaths, B. Simeone, *Combinatorial Optimization*, Springer, pp. 258-271.

Nguyen S., S. Pallottino, M. Gendreau (1998) Implicit Enumeration of Hyperpaths in a Logit Model for Transit Networks, *Transportation Science* 32 (1), pp. 54-64.

Nielsen O. A. (2000) A stochastic transit assignment model considering differences in passengers utility functions, *Transportation Research Part B: Methodological* 34 (5), pp. 377-402.

Paulley N., R. Balcombe, R. Mackett, H. Titheridge, J. Preston, M. Wardman, J. Shires, P. White (2006) The demand for public transport: The effects of fares, quality of service, income and car ownership, *Transport Policy* 13 (4), pp. 295-306.

Pryke A., S. Mostaghim, A. Nazemi (2007) Heatmap visualization of population based multi objective algorithms. *Evolutionary multi-criterion optimization*, Springer.

Pucher J., R. Buehler (2009) Integrating bicycling and public transport in North America, *Journal of Public Transportation* 12 (3), pp. 79-104.

Rasouli S., H. Timmermans (2012) Uncertainty in travel demand forecasting models: literature review and research agenda, *Transportation Letters: the international journal of transportation research* 4 (1), pp. 55-73.

Sager T. (2013) The Comprehensiveness Dilemma of Cost-Benefit Analysis, *EJTIR* 13 (3), pp. 169-183.

Santos B., A. Antunes, E. Miller (2009) Multiobjective Approach to Long-Term Interurban Multilevel Road Network Planning, *Journal of Transportation Engineering* 135 (9), pp. 640-649.

Schroten A., H. P. van Essen, S. J. Aarnink, E. Verhoef, J. Knockaert (2014) Externe en infrastructuurkosten van verkeer, een overzicht voor Nederland in 2010 (in Dutch). Delft, CE Delft.

Sharma S., S. Ukkusuri, T. Mathew (2009) Pareto Optimal Multiobjective Optimization for Robust Transportation Network Design Problem, *Transportation Research Record: Journal of the Transportation Research Board* 2090, pp. 95-104.

Sheffi Y. (1985) *Urban transportation networks: Equilibrium analysis with mathematical programming methods*, Prentice-Hall, Englewood Cliffs, New Jersey.

Sierzchula W. (2015) *Development and Early Adaption of Electric Vehicles*, TU Delft, TRAIL Thesis Series, T2015/2.

Significance, VU University Amsterdam, John Bates Services (2013) Values of time and reliability in passenger and freight transport in The Netherlands, Report for the Ministry of Infrastructure and the Environment.

Soderholm P. (2013) *Environmental Policy and Household Behaviour: Sustainability and Everyday Life*, Routledge.

Sparing D., R. M. P. Goverde (2013) An optimization model for periodic timetable generation with dynamic frequencies. *16th International IEEE Conference on Intelligent Transportation Systems - (ITSC)*, Den Haag.

Spiess H., M. Florian (1989) Optimal strategies: A new assignment model for transit networks, *Transportation Research Part B: Methodological* 23 (2), pp. 83-102.

Sumalee A., S. Shepherd, A. May (2009) Road user charging design: dealing with multi-objectives and constraints, *Transportation* 36 (2), pp. 167-186.

SWOV (2009) De balans opgemaakt, Duurzaam Veilig 1998-2007 (in Dutch). Leidschendam, Stichting Wetenschappelijk Onderzoek Verkeersveiligheid.

Taboada H. A., F. Baheranwala, D. W. Coit, N. Wattanapongsakorn (2007) Practical solutions for multi-objective optimization: An application to system reliability design problems, *Reliability Engineering & System Safety* 92 (3), pp. 314-322.

Tahmasseby S. (2009) *Reliability in urban public transport network assessment and design*, TU Delft, TRAIL Thesis Series, T2009/6.

Tan W. G. Z., L. B. Janssen-Jansen, L. Bertolini (2013) The Role of Incentives in Implementing Successful Transit-Oriented Development Strategies, *Urban Policy and Research* 32 (1), pp. 33-51.

Thon D. (1982) An axiomatization of the Gini coefficient, *Mathematical Social Sciences* 2 (2), pp. 131-143.

Transport for London (2014) Travel in London, Report 7. London, Transport for London.

Tregouët B. (2010) La mobilité des Français, Panorama issu de l'enquête nationale transports et déplacements 2008 (In French). T. L. Jeannic, Commissariat général au développement durable.

Uchida K., A. Sumalee, D. Watling, R. Connors (2007) A Study on Network Design Problems for Multi-modal Networks by Probit-based Stochastic User Equilibrium, *Networks and Spatial Economics* 7 (3), pp. 213-240.

Ukkusuri S. V., G. Patil (2009) Multi-period transportation network design under demand uncertainty, *Transportation Research Part B: Methodological* 43 (6), pp. 625-642.

Ulrich T. (2013) Pareto-Set Analysis: Biobjective Clustering in Decision and Objective Spaces, *Journal of Multi-Criteria Decision Analysis* 20 (5-6), pp. 217-234.

Van Eck G., T. Brands, L. J. J. Wismans, A. J. Pel, R. Van Nes (2014) Model Complexities and Requirements for Multimodal Transport Network Design: Assessment of Classical, State-of-the-Practice, and State-of-the-Research Models, *Transportation Research Record: Journal of the Transportation Research Board* (2429), pp. 178–187.

Van Nes R. (2002) *Design of multimodal transport networks, a hierarchical approach*, TRAIL thesis series, T2002/5.

Van Nes R., V. Marchau, B. Van Wee, I. Hansen (2007) Reliability and Robustness of Multimodal Transport network Analysis and Planning: Towards a New Research Agenda. *3rd INSTR*, Kijkduin, The Netherlands.

Van Ommeren J., D. Wentink (2012) The (hidden) cost of employer parking policies, *International Economic Review* 53 (3), pp. 965-978.

Van Veluwen A., Y. De Vries (2014) Publieksrapportage Rijkswegennet, 3e periode 2013 (in Dutch), Rijkswaterstaat Water Verkeer en Leefomgeving.

Verdus (2014a) "I-PRISM: Innovative Pricing for Sustainable Mobility." Retrieved 15-12-2014, from http://dbr.verdus.nl/768.

Verdus (2014b) "Sustainable accessibility of the Randstad (SAR)." Retrieved 19-11-2014, from http://dbr.verdus.nl/1678.

Verdus (2015a) "Synchronising networks: the modelling, use, governance and design of supernetworks." Retrieved 15-01-2015, from http://dbr.verdus.nl/726.

Verdus (2015b) "TRISTAM: traveller response and information service technology: analysis and modelling." Retrieved 15-01-2015, from http://dbr.verdus.nl/744.

Wang Y. (2014) *Optimal Trajectory Planning and Train Scheduling for Railway Systems*, TU Delft, TRAIL Thesis Series, T2014/7.

Wardman M. (2004) Public transport values of time, *Transport Policy* 11(4), pp. 363-377.

While L., P. Hingston, L. Barone, S. Huband (2006) A faster algorithm for calculating hypervolume, *Evolutionary Computation, IEEE Transactions on* 10 (1), pp. 29-38.

Wilson N. H., A. Nuzzolo (2008) *Schedule-Based Modeling of Transportation Networks*, Springer.

Wismans L., E. V. Berkum, M. Bliemer (2012) Effects of Optimizing Externalities Using Cooperating Dynamic Traffic Management Measures on Network Level, *Journal of Intelligent Transportation Systems* 17 (1), pp. 65-77.

Wismans L., E. V. Berkum, M. Bliemer (2013) Effects of Optimizing Externalities Using Cooperating Dynamic Traffic Management Measures on Network Level, *Journal of Intelligent Transportation Systems* 17(1), pp. 65-77.

Wismans L. J. J. (2012) *Towards Sustainable Dynamic Traffic Management*, University of Twente, TRAIL Thesis Series, T2012/4.

Wismans L. J. J., E. C. Van Berkum, M. C. J. Bliemer (2011) Comparison of Multiobjective Evolutionary Algorithms for Optimization of Externalities by Using Dynamic Traffic Management Measures, *Transportation Research Record: Journal of the Transportation Research Board* 2263 (1), pp. 163-173.

Yang H., M. G. H. Bell (1998) Models and algorithms for road network design: a review and some new developments, *Transport Reviews* 18 (3), pp. 257-278.

Yin Y., S. Lawphongpanich (2006) Internalizing emission externality on road networks, *Transportation Research Part D: Transport and Environment* 11 (4), pp. 292-301.

Zanten M. C. v., A. v. Alphen, J. Wesseling, D. Mooibroek, P. L. Nguyen, H. Groot Wassink, C. Verbeek (2013) Monitoringsrapportage NSL 2013, Stand van zaken Nationaal Samenwerkingsprogramma Luchtkwaliteit (in Dutch), Rijksinstituut voor Volksgezondheid en Milieu (RIVM).

Zavala G. R., A. J. Nebro, F. Luna, C. A. Coello Coello (2014) A survey of multi-objective metaheuristics applied to structural optimization, *Structural and Multidisciplinary Optimization* 49(4), pp. 537-558.

Zhang M. (2013) *A Freight Transport Model for Integrated Network, Service, and Policy Design*, TU Delft, TRAIL Thesis Series, T2013/9.

Zhao F., X. Zeng (2006) Optimization of transit network layout and headway with a combined genetic algorithm and simulated annealing method, *Engineering Optimization* 38 (6), pp. 701-722.

Zhu L., K. Deb, S. Kulkarni (2014) Multi-Scenario Optimization Using Multi-Criterion Methods: A Case Study on Byzantine Agreement Problem. IEEE Congress on Evolutionary Computation. Beijing, China, IEEE Press**:** 2601-2608.

Zitzler E., J. Knowles, L. Thiele (2008) Quality Assessment of Pareto Set Approximations, J. Branke, K. Deb, K. Miettinen and R. Słowiński, *Multiobjective Optimization*, Springer Berlin Heidelberg, 5252, pp. 373-404.

Zitzler E., M. Laumanns, L. Thiele (2001) SPEA2: Improving the strength Pareto evolutionary algorithm. Eurogen. Athens, Greece**:** 95-100.

Zitzler E., L. Thiele (1999) Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach, *Evolutionary Computation, IEEE Transactions on* 3 (4), pp. 257-271.

Zuidgeest M. H. P. (2005) *Sustainable urban transport development, a dynamic optimisation approach*, University of Twente, TRAIL Thesis Series, T2005/3.

Notation

List of symbols

Transportation network and demand

The network is defined by a range of sets and subsets, denoted by capital letters. Embellishments above letters are used to indicate transportation mode (here shown for x as example and based on Sheffi, 1985): plain (x) for car, $\hat{x}$ for PT, $\bar{x}$ for total (PT and car together) and $\tilde{x}$ for freight. Flows through the network can be route based (f) or link based (x). Several indicators δ are used to define the relation between links and routes. Demand q is defined for each origin-destination (OD) pair (index rs). Network attributes (such as travel times) can be link based (index a) or route based (index rs,k). A PT route belongs to a mode chain m, depending on the access and egress modes used in that route.

	Network definitions
A	Link set A
A_l	Subset of links that is traversed by line l
B	Set of PT vehicle types
D_u^P	Binary variable indicating a park-and-ride facility at stop u
D_u^T	The subset of types of PT lines T that serve stop u
G	Multimodal transportation network
K_{rs}	Set of car routes for OD pair r,s
$\hat{K}_m^{rs}$	Set of PT routes for OD pair r,s on mode chain m
L	Transit line set L
L_u	Set of lines serving stop u
L_{um}^s	Sensible set of lines at stop u to travel to destination s for mode chain m
M	Set of mode chains
M_R	Set of mode chains that start the trip with car access
M_S	Set of mode chains that end the trip with car egress
N	Node set N
R	Origin set
R_U	Set of highly urban origins
S	Destination set

S_U	Set of highly urban destinations
U	Stop set U
U_{rm}^A	Set of candidate stops for access from origin r for mode chain m
U_{sm}^E	Set of candidate stops for egress to destination s for mode chain m
U_u^T	Set of candidate stops for transfer from stop u
$U_{\hat{n}t}^P$	Set of stops to be processed during iteration $\hat{n}t$ of the PT assignment algorithm
$U_{\hat{n}t}^{Tused}$	Set of used stops for transfer from stop u during iteration $\hat{n}t$

Network flows, network attributes and demand

c_k^{rs}	(generalised) costs for car when using route k on OD pair r,s
c_{rs}	(generalised) costs for car when on OD pair r,s
$\hat{c}_{km}^{rs}$	(generalised) costs for PT when using route k in mode chain m on OD pair r,s
$\hat{c}_m^{rs}$	Combined costs for PT mode chain m on OD pair r,s
$\hat{c}_{rs}$	Combined costs for PT on OD pair r,s
$\breve{c}_{lum}^s$	(gen.) costs for PT when using line l at stop u in mode chain m to reach s
$\breve{c}_{um}^{rs}$	(generalised) costs for PT when using stop u in mode chain m to reach s from r
$\breve{c}_{um}^r$	(generalised) access costs for to reach stop u from origin r in mode chain m
d_a	Length of link a (km)
di_k^{rs}	Distance travelled when taking car route k on OD pair r,s
$\hat{di}_{km}^{rs}$	Distance travelled when taking route k in mode chain m on OD pair r,s
$\hat{di}_{kmb}^{rs}$	Distance travelled in vehicle type b taking route k in mode chain m on OD r,s
f_k^{rs}	Route flow for car (vehicles / hour)
$\hat{f}_{km}^{rs}$	Route flow for PT mode chains (passengers / hour)
F_l	Frequency of transit line l
$\hat{n}t$	Number of transfers counter in the PT assignment algorithm
$\hat{n}t_{max}$	Maximum number of transfers in the PT assignment algorithm
$\hat{n}t_{km}^{rs}$	Number of transfers when taking route k in mode chain m on OD pair r,s
p_{lum}^s	Probability (fraction) for boarding line l at stop s to destination s for m
p_{um}^{rs}	Probability (fraction) to choose stop u when travelling from r to s for m
$\bar{q}_{rs}$	Total passenger OD demand (number of trips from origin r to destination s)
q_{rs}	Car OD demand (number of car-only trips from origin r to destination s)
$\hat{q}_{rs}$	PT OD demand (number of trips that include PT from origin r to destination s)
$\hat{q}_{rs}*$	Optimal PT OD demand after solving the lower-level model (UE situation)
$\tilde{q}_{rs}$	Freight OD demand (Passenger Car Equivalents / hour)
$t_a(x_a, \underline{y})$	(generalised) cost function for car links
$\hat{t}_{al}(\underline{y})$	(generalised) cost function for PT links per transit line
$\hat{t}_{a'l'al}(\underline{y})$	(generalised) cost function for pairs of PT links per pair of transit lines
tv_k^{rs}	Travel time when taking car route k on OD pair r,s
$\hat{tv}_{km}^{rs}$	PT in-vehicle time when taking route k in mode chain m on OD pair r,s
$\hat{tv}_{kmb}^{rs}$	PT vehicle type b in-vehicle time when route k in mode chain m on OD pair r,s
$\hat{tw}_{km}^{rs}$	Waiting time when taking route k in mode chain m on OD pair r,s
$\hat{tw}_{kmb}^{rs}$	Waiting time for vehicle type b taking route k in mode chain m on OD pair r,s
tv_a	Car travel time on link a
tv_a^0	Free flow car travel time on link a
tv_{al}	PT In-vehicle travel time on link a in line l, or travel time on link a with access / egress mode l

T_l	PT system type of line l
v_a^0	Free flow speed of link a
x_a	Car link flow of link a (vehicle / hour)
$\tilde{x}_a$	Freight link flow of link a (Passenger Car Equivalents / hour)
$x_a^{\max}$	Capacity of link a (Passenger Car Equivalents / hour)
$\underline{x}$	Car flow vector
$\underline{x}^*$	Optimal car flow vector after solving the lower-level model (UE situation)
x_{al}	PT-vehicle flow on link a for transit line l (vehicle / hour)
$\hat{x}_{al}$	Passenger link flow per transit line (passengers / hour)
$\underline{\hat{x}}$	Passenger flow vector
$\underline{\hat{x}}^*$	Optimal passenger flow vector after solving the lower-level model (UE)
$\delta_{a,k}^{rs}$	Car route, link flow indicator
$\hat{\delta}_{al,km}^{rs}$	PT route, passenger link and line indicator
$\hat{\delta}_{a'l'al,km}^{rs}$	PT route, passenger transfer indicator for pairs of links and lines
$\bar{\delta}_{a,km}^{rs}$	PT route, car link indicator
$\tilde{\delta}_a^{rs}$	Freight (shortest) route, link indicator
Δ_{bl}	PT vehicle type indicator, equals 1 if line l is of vehicle type b, 0 otherwise
ζ_{km}^{rs}	Fraction of travelers on OD pair r,s in PT mode chain m, choosing route k
ψ_m^{rs}	Fraction of PT travelers on OD pair r,s choosing mode chain m

Parameters

C_b	Fare (costs) for using PT of vehicle type b (euros per km)
E^{CO_2}	CO_2 emission factor for cars (grams/(veh*km))
$\hat{E}_b^{CO_2}$	CO_2 emission factor for PT vehicle type b (grams/(veh*km))
O_b	Operating costs for PT vehicle type b (euros per vehicle*hour)
O_{PR}	Operating costs of one park-and-ride space
n	Iteration number in the lower-level model
n_{max}	Number of iterations in the lower-level model
α_1	Weight factor for distance in car generalised cost function
α_2	Weight factor for travel time in car generalised cost function
$\hat{\alpha}_1$	Weight factor for distance in PT generalised cost function
$\hat{\alpha}_2$	Weight factor for in-vehicle time in PT generalised cost function
$\hat{\alpha}_3$	Weight factor for waiting time in PT generalised cost function
$\hat{\alpha}_4$	Weight factor for number of transfers in PT generalised cost function
β_1, β_2	Parameters in BPR function
θ_1	Logit parameter for mode choice
θ_2	Logit parameter PT mode chain choice
θ_3	Logit parameter for stop choice
θ_4	Logit parameter for line choice
Ω	Car occupation (passengers / car)

Optimisation problem

J	Number of optimisation processes
N^j	Number of solutions in Pareto set j
P^j	Pareto set from optimisation process j: all non-dominated solutions in ϕ^j
P^Σ	Superset: all non-dominated solutions in ϕ^Σ

P_{q*}^{j}	Pareto set j, optimised for corresponding demand q
$P_{q*,q}^{j}$	Pareto set j, optimised for demand $q*$ and calculated for demand q
Q	Number of transportation demand scenarios
V	Length of decision vector $\underline{y}$
V'	Length of binary representation of the decision vector
W	Length of objective vector $\underline{Z}$
$\underline{y}$	Decision vector / solution
$\underline{y}^{0}$	Base solution
$\underline{y}'$	Binary representation of the decision vector
$\underline{y}_{i}^{j}$	ith decision vector in pareto set j
Y	Set of feasible decision vectors
$\underline{Z}$	Objective vector upper level
$\underline{\bar{Z}}$	Normalised objective vector upper level
$\underline{Z}^{q}$	Objective vector using transportation demand q
$\underline{z}$	Objective function lower level
$\alpha(\underline{Z})$	Attainment function (used as performance indicator)
κ	Nonzero relation in *AFD* indicator
ρ_{v}^{j}	Fraction of solutions in P^{j} that have a positive value for decision variable y_{v}
ϕ^{j}	Set of solutions calculated during optimisation process j
ϕ^{Σ}	Set of solutions calculated during all J optimisation processes

Solution algorithms and Pareto set analysis

Solution algorithms

H	Number of generations (index h)
H^{j}	Number of generations in optimisation process j
I	Number of restarts (in ε-NSGAII, index i)
ε_{z}	Epsilon value in (in ε-NSGAII) for objective z
v_{a}	Archive size (in NSGAII)
v_{a}^{j}	Archive size (in NSGAII) in optimisation process j
$v_{a,\varepsilon}^{i}$	Size of ε-archive at restart i (in ε-NSGAII)
v_{p}	Population size (in NSGAII)
v_{p}^{0}	Initial population size (in ε-NSGAII)
v_{p}^{i}	Population size after restart i (in ε-NSGAII)
ϖ	Injection scheme parameter (in ε-NSGAII)
π_{h}	Set of archive solutions in generation h
π_{h}^{ε}	Set of ε-archive solutions in generation h
ϑ_{h}	Union of archive and offspring in generation h
σ_{h}	Set of solutions in mating pool in generation h
ς	Number of generations without change to trigger a restart (in ε-NSGAII)
ϕ_{h}	Set of offspring solutions in generation h
φ	Mutation rate

Pareto set analysis

$TO_{w,w'}(\underline{y},\underline{y}')$	Trade-off value for objectives w and w' when comparing solution $\underline{y}$ and $\underline{y}'$
$ATO_{w,w'}(\bar{P})$	Average trade-off value for objectives w and w' in Pareto set P
$SL_{w,v,\nabla_{v}}$	Slope for objective w, decision variable v and step size ∇_{v}

RC_{w,v,∇_v}	Relative change for objective w, decision variable v and step size ∇_v
NRC_{w,v,∇_v}	Normalised relative change for objective w, decision variable v, step size ∇_v
$BCS_{W_C}(P)$	Best compromise solution in Pareto set P for objective set W_C
$RBCS_{W_C}(P)$	Best relative compromise solution in Pareto set P for objective set W_C
$ABCS_{W_C}(P)$	Best absolute compromise solution in Pareto set P for objective set W_C
RW_w	Relative importance of objective w when determining a compromise solution
AW_w	Absolute importance of objective w when determining a compromise solution
	Pruning methods
$d_{i,g}^{NB}$	Distance to the nearest neighbour from solution i in direction g
$e_{ii'}$	Euclidian distance from solution i to i' in the normalised objective space
k	Number of clusters in K-means clustering
Ω_i^g	Convex cone g, related to solution i
χ	Gini index for spread of solutions (used to assess pruning methods)
Δ^r	Minimum spread parameter in the PIT filter
Δ^t	Insignificance parameter in the PIT filter

List of abbreviations

PT	Public Transport
P&R	Park-and-ride
SAR	Sustainable Accessibility of the Randstad
SRMT	Strategy towards Sustainable and Reliable Multimodal Transport in the Randstad
SCBA	Social Cost-Benefit Analysis
MCDM	Multi-Criteria Decision Making
AMA	Amsterdam Metropolitan Area
GA	Genetic Algorithm
EA	Evolutionary Algorithm
NDP	Network Design Problem
OD	Origin-Destination
UE	User Equilibrium
MO	Multi-Objective
MOOP	Multi-Objective Optimisation Problem
NSGAII	Non-dominated Sorting Genetic Algorithm II
MOEA	Multi-Objective Evolutionary Algorithm
SPEA	Strength Pareto Evolutionary Algorithm
TTT	Total Travel Time
ATT	Average Travel Time
USU	Urban Space Used (by parking)
OpD	Operating Deficit
CE	CO_2 Emissions
TUC	Weighted average of TTT, USU and CE
VoT	Value of Time
TOD	Transit-Oriented Development

Summary

Transportation provides many benefits to society, but also has some negative externalities, like climate impact, bad air quality or use of urban space. To make the transportation system more sustainable, its benefits need to be maximised and its negative externalities need to be minimised. This means that all three aspects of sustainability need to be improved: economic, social and environmental sustainability, which is also the goal of the wider research programme Sustainable Accessibility of the Randstad. Because medium-distance to long-distance car trips cause a large share of the externalities of the present transportation system, shifting these trips from car to PT (public transport) can contribute to this aim. Improving the PT network is chosen as solution direction, because it is believed that this can improve social and environmental sustainability and possibly at the same time improve economic sustainability. The polycentric nature of the Randstad area and the expectation that large-scale investments in infrastructure are not feasible in the near future, are main reasons to focus on facilitating multimodal trips. In a multimodal trip a traveller combines at least one private mode (e.g. car, bicycle) with at least one public mode (e.g. train, metro, tram, bus). The transfer between these modes can be facilitated by measures that provide transfer locations, like new park-and-ride facilities and new train stations, and measures that provide an attractive PT network, like increasing frequencies of PT lines. A network approach is needed, because the measures interact with the existing network and with each other. The objective of this research is to determine which (re)designs of the multimodal passenger transportation network in the Randstad contribute best to improving various aspects of sustainability and to provide insight into the extent to which these aspects are improved and into how scores on these aspects and designs relate. To this end the problem is first formulated, after which a solution approach is applied, resulting in an operational modelling framework. Several methods for decision support are used to analyse the optimisation results of a real-life case study, including the long-term robustness of the results. Finally recommendations for further research are made.

Problem definition

A mathematical optimisation approach is adopted, formulated as a multi-objective network design problem (MO-NDP). The problem is defined as a bi-level mathematical optimisation

problem, where the behavioural response of travellers on measures taken is incorporated in the lower-level model. Requirements for this lower-level model are to include multimodal trip making, an accurate calculation of network flows, a low computation time and availability in a software environment. Existing components are combined into a proper transportation model that is fast enough and contains both unimodal and multimodal trip options during mode and route choice. Solving this MO, multimodal NDP yields a Pareto set of network designs. This set contains the best possible designs and the associated sustainability scores, so the resolution of the chosen solution direction is investigated. This insight would not be provided by the traditional approach of choosing a pre-defined set of network solutions formulated by expert judgement, because it is likely that better solutions are overlooked: the behavioural response of travellers is complicated and therefore cannot entirely be included in the expert judgement. The Pareto set also enables identifying the dependencies and trade-offs between objectives.

Solution approach

The MO-NDP is proven to be NP-hard, so heuristics are needed to approximate the solution of the problem. Genetic algorithms (GAs) are well suited to solve MO-NDPs, because these algorithms have a low risk of ending up in a local minimum, do not require the calculation of a gradient and are able to produce a diverse Pareto set. GAs are population-based algorithms, inspired by the process of natural evolution (survival of the fittest). Each new generation of 'children' is created by reproduction from the existing generation of 'parents', as well as by mutation. NSGAII is a popular MO GA that was earlier successfully used to solve MO-NDPs, while ε-NSGAII proved to outperform NSGAII for other types of problems. The performance of both algorithms is compared when applied to the multimodal passenger transportation NDP. Experiments show that ε-NSGAII outperforms NSGAII also for this optimisation problem, especially in the early stage of running the algorithm. This faster convergence is especially useful for cases with objective functions that are highly computationally expensive, like in the case study in this research.

It is inherent in any GA to contain stochastic processes: each run of the algorithm has a random seed as a starting point for the genetic operators. Furthermore, GAs need parameter settings for generation size and mutation rate. Well performing settings are not known before running the algorithm multiple times for a specific case study. Experiments show that decision variable values of solutions in Pareto sets from different runs are substantially different, so the advised decisions are influenced by the stochastic principles of the algorithm. This is related to the flat shapes of the objective functions in the chosen definition of the optimisation problem: similar objective values can be reached by different measures in the multimodal network. As a result the differences in the attained objective values among multiple runs of the algorithm are small: both within runs with the same parameters settings and between runs with different parameter settings. Therefore it can be concluded that the method is suitable to provide a Pareto set that is near-optimal and can be used to derive information for decision support.

Methods for decision support

Several methods to analyse and visualise the Pareto set are developed and applied to make the Pareto set easier to understand and more useful to guide the decision maker to a final solution

for implementation. The first method is a step-by-step pruning method, which sets bounds to one or more objective values or selects decision variables that are politically desirable, until finally one solution is selected for implementation. The second method determines trade-offs between objectives. The third method identifies the min-max solution: for each solution the least scoring objective is leading when selecting the best solution from all Pareto solutions. This is useful in the political context of decision making, where a compromise solution is needed and each political party (represented by an objective) does not accept a bad score. The fourth method systematically reduces the number of solutions, which results in a comprehensive subset of Pareto solutions, which is representative for the entire Pareto set. Application of this method to the case study provides an overview of 15 solutions as a representative subset of the entire Pareto set of 210 solutions. This is still regarded as a reasonable number to be assessed by the interviewed policy officers.

The modelling framework helps to attain better network designs and reveals to what extent concrete measures contribute to achieving policy objectives. This helps policy makers to set realistic goals when deciding on a long-term policy document. Before the optimisation process takes place, the problem has to be defined in conjunction with the decision maker. After the optimisation process the outcomes can be analysed in detail by the methods for decision support. When the methods are included in an interactive decision support tool, the consequences of certain choices are shown directly, which is useful in any practical situation where many options are considered to choose a final option for implementation.

The case study in the Randstad area

The MO-NDP is applied to a realistic case study in the Amsterdam metropolitan area. Based on interviews and workshops with policy officers, for each of the three main aspects of sustainability one most important objective is defined. Total travel time represents accessibility (economic sustainability), urban space used by parking represents social sustainability and CO_2 emissions represent climate impact (environmental sustainability). Finally, the fourth objective is operating deficit to represent cost-efficiency. This is the difference between PT operating costs (including depreciation of investments in park-and-ride facilities) and operating revenues and represents the viewpoint of the regional government, who are responsible for these costs. Investments in rail infrastructure are excluded, because these are paid by the national government. The stakeholders preferred costs to be included as an objective rather than as a budget constraint, to reveal trade-off information, for example to identify the marginal costs needed to reach additional travel time savings. All likely locations for new park-and-ride facilities and new local train stations in the study area are included as decision variables. Possible changes in service lines are included as decision variables as well, in the form of express train status of stations, frequencies of major bus lines and frequencies of local train lines. These measures are small-scale measures compared to the entire multimodal network that is modelled: most PT services, including express train services, and the car network are fixed. This results in relatively small changes in objective values, because these are calculated for the entire network.

It is demonstrated that convergence is almost reached within a reasonable computation time of 3000 function evaluations, showing that the modelling framework is appropriate for this realistically sized case study. The resulting Pareto set is analysed using the mentioned

methods for decision support. The measures related to multimodal trip making are able to achieve a reduction compared to the base situation for all four objectives simultaneously: one specific design shows a reduction of at least 0.4% for total travel time, urban space used and CO_2 emissions, with a 0.4% reduction in costs. For the individual objectives, larger improvements are possible: CO_2 emissions can be improved by 0.45%, total travel time can be improved by 0.79% and urban space used by 1.5%, but not without deteriorating at least one other objective. In absolute terms every AM peak almost 4000 hours of travel time, more than 2000 parked cars or more than 12 tons of CO_2 emissions (equivalent to the daily direct CO_2 emissions of more than 500 Dutch households) can be saved. Total travel time, urban space used and CO_2 emissions are mainly in line with each other and are opposed to operating deficit. However, no single solution exists that is optimal for all three objectives. An average trade-off value of 6.90 euro / hour is found, which means that on average one hour of travel time is saved by spending an additional €6.90 on operating PT. This value is smaller than the value of time used in Dutch cost-benefit analyses, so from the viewpoint of the regional government cost-efficient measures exist within the selection of measures considered. It should be noted that investments in infrastructure (paid by the national government) are not included in the analysis, so whether the measures are cost-efficient from a societal point of view is unsure. To save one kilo of CO_2 emissions, the regional government on average has to spend an additional €1.10. This is much more than alternative measures to reduce CO_2 emissions of the transportation system, like tighter CO_2 emission standards for cars (€0.11 per kilo) or stimulating electric cars (€0.35 per kilo). The combined effects of additional expenses on the objectives are as follows: each additional €1.00 of operating deficit on average results in a marginal gain of approximately 0.1 hours of total travel time, 0.07 fewer parking spaces used and 0.5 kilos fewer CO_2 emissions. The min-max solution for all four objectives is in the best 30% of the range covered by each objective, so with equal importance for the normalised objective functions it is possible to satisfy all four objectives to a large extent simultaneously.

Opening park-and-ride facilities leads to lower total travel time in general, but also to more CO_2 emissions, because not only former car trips are attracted, but also former PT-only trips. Furthermore, increasing train frequencies and, to a smaller extent, bus frequencies in general reduces total travel time, urban space used and CO_2 emissions at the cost of increasing operating deficit. Increasing frequencies appears to have a larger effect on the sustainability objectives than opening new park-and-ride facilities and new train stations, indicating that multimodal trips can be stimulated better by improving the quality of the PT network itself than by introducing additional transfer locations. Another reason to prefer frequency increases over new transfer locations is that frequency increases are easier to reverse, making these measures easier to adapt to changing circumstances in the future. In Pareto-optimal network designs with low frequencies on local train lines more stations acquire express train status to compensate for the low frequency at these stations. High frequencies occur simultaneously on train and on bus lines in Pareto-optimal network designs: in this case study buses and trains strengthen each other. Finally, from all 37 possible measures in the case study only one measure (a possible location for a new train station) is identified as 'always regret' (occurring in none of the Pareto solutions). No 'no regret' measures (occurring in all Pareto solutions) are found. This indicates that whether a measure performs optimally highly depends on the objectives that are preferred.

This research shows that measures related to multimodal trip making can contribute to reaching sustainability goals by making travel alternatives that include PT more attractive. On the other hand, the resolution of this type of measures for e.g. CO_2 emissions is limited, because the high-emission alternative (i.e. the car-only trip) remains attractive for the traveller. Therefore the impact of network-wide measures that make the car alternative less attractive is demonstrated. The same transportation model is used for this demonstration, despite this model excluding behavioural responses that become relevant for this kind of measures like elastic transportation demand. The results show that road pricing achieves much larger improvements in environmental and social sustainability, but total travel time increases, because former car users switch to slower PT alternatives. Furthermore, the alternative spatial planning concept transit-oriented development is tested, where all new residents and jobs in the region until 2030 are developed in the vicinity of train stations. This leads to considerable improvements in environmental sustainability and also to a small decrease of delay times.

Long-term robustness

The transportation network is designed for a long time period, so demand forecasting is necessary to make a proper design. It is relevant to know to what extent resulting network designs are robust for uncertainty in this demand forecast. Therefore Pareto sets are constructed for the base demand forecast for 2030 and for two other future demand scenarios, i.e. the demand forecast for 2020 as lower boundary scenario for PT use and transit-oriented development as upper boundary scenario for PT use.

It turns out that the network solutions that are optimal for the 2030 demand scenario, perform also quite well if demand is lower or has changed due to transit-oriented development. The loss in objective function values is small and the values of decision variables are not strongly influenced by transportation demand. An important reason for this is that all three tested demand scenarios are realistic forecasts with a similar structure: in the NDP those measures that influence the largest flows generally are most effective. A distinction between measures that are easy to reverse or hard to reverse is useful here in case a measure appears to be performing badly due to changing circumstances. The impact of omitting the objective function CO_2 emissions from the analysis is small, because this objective is largely in line with the other objectives urban space used and total travel time.

Recommendations for further research

The lower-level transportation model may be extended in several ways to improve the accuracy of the network flows further. First, the techniques used to model combined mode and route choice may be improved to handle the behavioural complexity of multimodal trip making better. Second, the attributes in the generalised costs calculation may be extended, for example to include comfort level or crowding. The third possible approach is completely different and includes (daily) activity patterns in the models, taking aspects like vehicle availability and combinations of multiple trips in a tour into account. Although it is expected that a more elaborate lower-level model leads to better optimisation results, this is not known. Simpler models might in the end lead to better optimisation results, because a lower computation time for the lower-level model enables more function evaluations during optimisation. More research is needed to indicate whether the conceptual limitations of the used model are negligible at the full network scale (leading to the same optimal network

designs) or systematically bias the assignment results (leading to different network designs being optimal).

For the solution algorithm ε-NSGAII more similar tests are needed on different real-life design problems, because this algorithm has not often been applied and performs differently for each case study. It is also of interest whether other heuristics can be found with better performance than ε-NSGAII, given the fast developments in the field of genetic computing. Research for more case studies is needed to be able to draw more general conclusions concerning the effects of stochastic processes inherent in GAs on the outcome of one single run. For future MO optimisation studies the methods applied in this research can be used directly to present other Pareto results in a more comprehensive way. However, more research is needed to present the complicated results of MO optimisation in an easy-to-understand way, because now these results are difficult to explain to decision makers.

In general a vast number of stakeholders are involved in operating PT. The modelling framework in this research is applied using the interests of a single stakeholder, but for society as a whole it would be best to optimise the PT network as one integrated network, taking into account the effects for all relevant governments and other parties in the region. This relates to the used costs function, which only included the costs relevant for the regional government, neglecting costs to be paid by the national government. Furthermore, in the current definition the four objectives are defined as network-wide values. A relevant question is whether the benefits of the network designs are distributed over society in a fair way (equity), for example geographically, but also over users of various modes of transportation. The effects of measures, i.e. the way the objective values are determined, may change due to technological developments, like hydrogen cars or electric cars. The availability of new modes can also be included, like the introduction of the electric bicycle. This may lead to different network designs, and it is relevant to known to what extent these developments influence the optimisation results.

The solution direction presented in this research is only one step on the long road to a sustainable transportation system in the Randstad. If car trips are made less attractive and at the same time enough capacity and quality are provided in the PT network (including the connections between the PT network and the bicycle and car networks), improving environmental sustainability may be possible without harming economic sustainability. A favourable extension of the NDP is therefore to include other types of measures in the multimodal NDP, like new rail infrastructure or road pricing. In that case the behavioural response model should be extended, leaving the assumption of a fixed total transportation demand. Furthermore, combinations with other measures or developments can be made, for example technological developments to improve environmental performance of vehicles, like the introduction of electric cars, or a different spatial planning.

Samenvatting

Vervoer heeft vele voordelen voor de samenleving, maar ook negatieve externe effecten, zoals bijdrage aan klimaatverandering, een negatief effect op de luchtkwaliteit en ruimtebeslag in steden. Om tot een duurzamer transportsysteem te komen, is het nodig de voordelen te maximaliseren en de externe effecten te minimaliseren. Met andere woorden, alle drie de hoofdaspecten van duurzaamheid moeten verbeterd worden: economische en sociale duurzaamheid en het milieu. Dit is ook de doelstelling van het bredere onderzoeksprogramma Duurzame Bereikbaarheid van de Randstad. Een groot deel van de externe effecten van het huidige transportsysteem wordt veroorzaakt door middellange tot lange autoritten, waardoor het vervangen van deze ritten door ritten per openbaar vervoer (OV) kan bijdragen aan dit doel. Als oplossingsrichting wordt het verbeteren van het OV-netwerk gekozen, omdat dit positieve invloed kan hebben op sociale duurzaamheid en het milieu en mogelijk tegelijkertijd de economische duurzaamheid kan verbeteren. De polycentrische aard van de Randstad en de verwachting dat grootschalige investeringen in infrastructuur in de nabije toekomst niet haalbaar zijn, zijn belangrijke redenen om te focussen op multimodale ritten. In een multimodale rit combineert een reiziger ten minste één private vervoerwijze (bijv. auto, fiets) met ten minste één collectieve vervoerwijze (bijv. trein, metro, tram, bus). De overstap tussen deze vervoerwijzen kan worden vergemakkelijkt door maatregelen die voor overstaplocaties zorgen, bijvoorbeeld nieuwe parkeer- en reis voorzieningen (P&R) en nieuwe treinstations, of door maatregelen die zorgen voor een aantrekkelijk OV-netwerk, zoals het verhogen van frequenties van OV-lijnen. Omdat de maatregelen een wisselwerking hebben met het bestaande netwerk en met elkaar, is een netwerkbenadering nodig. Het doel van dit onderzoek is om te bepalen welke (her)ontwerpen van het multimodale personenvervoersnetwerk in de Randstad het beste bijdragen aan het verbeteren van verschillende aspecten van duurzaamheid en om inzicht te krijgen in welke mate deze aspecten verbeterd kunnen worden en hoe de scores op deze aspecten samenhangen met de netwerkontwerpen. Nadat het probleem is gedefinieerd, wordt een oplosmethode toegepast, wat leidt tot een operationeel modelraamwerk. De optimalisatieresultaten van een casestudie worden geanalyseerd door verschillende methoden voor beslisondersteuning toe te passen. Ook wordt aandacht besteed aan de robuustheid van de resultaten op de lange termijn. Tenslotte worden aanbevelingen gedaan voor verder onderzoek.

Probleemdefinitie

Het probleem wordt aangepakt als een wiskundig optimalisatieprobleem, geformuleerd als een netwerkontwerpprobleem met meerdere doelstellingen (multi-objective network design problem, MO-NDP). Het wordt gedefinieerd als een twee-niveau wiskundig optimalisatieprobleem, waar de gedragsreactie van reizigers op de genomen maatregelen is opgenomen in het lage niveau. Dit model moet om kunnen gaan met multimodale ritten, moet de stromen in het netwerk nauwkeurig berekenen, een lage rekentijd hebben en beschikbaar zijn in een softwareomgeving. Bestaande modelcomponenten zijn gecombineerd tot een adequaat verkeersmodel dat snel genoeg rekent en zowel unimodale als multimodale keuzeopties biedt tijdens vervoerwijze- en routekeuze. Het resultaat van het oplossen van dit multimodale MO-NDP leidt tot een Pareto-set van netwerkontwerpen. Deze verzameling bevat de best mogelijke ontwerpen, inclusief bijbehorende duurzaamheidsscores: het oplossend vermogen van dit type maatregelen wordt onderzocht. De traditionele benadering, het onderzoeken van enkele door experts voorgedefinieerde netwerkoplossingen, biedt dit inzicht niet, omdat betere oplossingen mogelijk over het hoofd worden gezien. De gedragsreactie van reizigers is namelijk ingewikkeld, waardoor het niet mogelijk is deze volledig mee te nemen in expert judgement. De Pareto-set maakt het daarnaast mogelijk om afhankelijkheden en afwegingen tussen doelen te bepalen.

Oplossingsmethode

Het is bekend dat het MO-NDP NP-hard is, en dus zijn heuristieken nodig om te oplossing te benaderen. Genetische algoritmes (GA's) zijn geschikt voor het oplossingen van MO-NDP's, omdat dit type algoritmes een laag risico heeft om in een lokaal minimum te eindigen, het niet nodig is om een gradiënt uit te rekenen en omdat ze in staat zijn een diverse Pareto-set te produceren. GA's werken met een populatie, geïnspireerd op het proces van natuurlijke evolutie ('survival of the fittest'). Elke nieuwe generatie 'kinderen' wordt gecreëerd uit een bestaande generatie 'ouders' door reproductie en mutatie. Een populair MO GA is NSGA-II, dat eerder succesvol is toegepast om MO-NDP's op te lossen. Ook is bewezen dat ε-NSGAII beter presteert dan NSGAII voor andere typen problemen. De prestatie van beide algoritmes wordt vergeleken voor toepassing op het multimodale passagiers-NDP. De experimenten laten zien dat ε-NSGAII ook voor dit optimalisatieprobleem beter presteert dan NSGAII, vooral in een vroeg stadium van het algoritme. Deze snellere convergentie is vooral nuttig voor gevallen waar langdurige berekeningen nodig zijn om de doelfunctiewaarden uit te rekenen, zoals in de casestudie in dit onderzoek.

Stochastische processen zijn eigen aan GA's: elke keer dat het algoritme wordt gedraaid, is een andere random seed nodig om de genetische operatoren op te starten. Ook hebben GA's parameters die moeten worden ingesteld, zoals omvang van een generatie en mate van mutatie. Welke parameters goed werken voor een specifieke casestudie is niet bekend voordat het algoritme meerdere keren is getest. De experimenten laten zien dat de waarden van beslisvariabelen van oplossingen in Pareto-sets uit afzonderlijke runs van het algoritme behoorlijk verschillen. De geadviseerde beslissingen worden dus beïnvloed door de stochastische processen in het algoritme. Dit heeft te maken met de platte vorm van de doelfuncties in de gekozen definitie van het optimalisatieprobleem: vergelijkbare doelfunctiewaarden kunnen worden bereikt door verschillende maatregelen. De verschillen in

de bereikte doelfunctiewaarden zijn daardoor klein, zowel voor runs met de zelfde parameterwaarden als voor runs met verschillende parameterwaarden. Hieruit kan worden geconcludeerd dat de methode geschikt is om een Pareto-set te leveren die dicht bij het optimum zit, dat gebruikt kan worden om beslisondersteunende informatie uit af te leiden.

Beslisondersteunende methoden

Er zijn verschillende methoden om de Pareto-set te analyseren en te visualiseren ontwikkeld en toegepast. Dit helpt bij het begrijpen van de Pareto-set en ondersteunt de beslisser bij het kiezen van één te implementeren oplossing op basis van de Pareto-set. De eerste methode reduceert de set stap voor stap door grenzen te stellen aan één of meer doelwaarden en / of door politiek gewenste beslisvariabelen te selecteren, totdat er één oplossing overblijft om te implementeren. De tweede methode bepaalt afwegingen tussen doelen in de vorm van trade-off waarden. De derde methode bepaalt de min-max-oplossing, in welk geval voor elke oplossing het slechtst scorende doel bepalend is voor de kwaliteit van die oplossing. Dit is handig indien men in de politieke context van het beslisproces op zoek is naar een compromis, waarin elke politieke partij zijn eigen doel nastreeft en geen slechte score voor dat doel accepteert. De vierde methode reduceert het aantal oplossingen systematisch, waardoor een overzichtelijke deelverzameling van Pareto-oplossingen ontstaat, die zo veel mogelijk representatief is voor de hele Pareto-set. Toepassing van deze methode op de casestudie resulteert in een set van 15 oplossingen die de hele Pareto-set van 210 oplossingen vertegenwoordigt, wat door de betrokken beleidsambtenaren als te overzien wordt beoordeeld.

Het modelraamwerk draagt bij aan het vinden van betere netwerkontwerpen en het laat zien in welke mate concrete maatregelen bijdragen aan het bereiken van beleidsdoelen. Dit helpt beleidsmakers bij het opstellen van langetermijnbeleidsdocumenten. Voorafgaand aan het optimalisatieproces moet het probleem in overleg met de beleidsmakers worden gedefinieerd. Naderhand kunnen de uitkomsten in detail worden geanalyseerd met de genoemde beslisondersteunende methoden. Indien deze methoden in een interactieve tool worden opgenomen, kunnen de consequenties van bepaalde keuzes direct worden getoond, wat nuttig is in elke praktijksituatie waarin veel opties worden overwogen om uiteindelijk één maatregel te kiezen die wordt geïmplementeerd.

De casestudie in de Randstad.

Het MO-NDP is toegepast op een realistische casestudie in de stadsregio Amsterdam. Voor elk van de drie hoofdaspecten van duurzaamheid is één belangrijkste doel geformuleerd, op basis van interviews en workshops die gehouden zijn met beleidsambtenaren. Totale reistijd staat voor bereikbaarheid (economische duurzaamheid), ruimtebeslag door parkeren staat voor sociale duurzaamheid en CO_2-uitstoot staat voor invloed op klimaatverandering (milieu). Het vierde doel is beperken van de benodigde exploitatiebijdrage (kostenefficiëntie), wat bestaat uit het verschil tussen exploitatiekosten van het OV (inclusief afschrijving van investeringen in P&R) en de opbrengsten uit kaartverkoop. Dit komt overeen met het oogpunt van de regionale overheden, die verantwoordelijk zijn voor deze kosten. Investeringen in spoorinfrastructuur worden niet meegenomen, omdat die door de nationale overheid worden betaald. De betrokken beleidsambtenaren wilden liever dat het beperken van kosten als doel wordt meegenomen dan als randvoorwaarde, omdat trade-off informatie nuttig is,

bijvoorbeeld om de marginale kosten te bepalen die nodig zijn om extra reistijdbesparingen te realiseren. Alle potentieel succesvolle locaties voor nieuwe P&R-voorzieningen en voor Sprinter-stations in het studiegebied zijn opgenomen als beslisvariabele. Ook mogelijke veranderingen in OV-lijnen zijn basisvariabelen, in de vorm van Intercity-status van stations, frequenties van belangrijke buslijnen en frequenties van Sprinter-treinen. Deze maatregelen zijn kleinschalig ten opzichte van het totale multimodale netwerk in het studiegebied. Een meerderheid van de OV-lijnen, waaronder alle Intercity-lijnen, en het hele autonetwerk veranderen namelijk niet. Hierdoor zijn de veranderingen in de doelwaarden relatief klein, omdat deze voor het netwerk in het hele studiegebied worden berekend.

Er wordt convergentie bereikt binnen een redelijke rekentijd van ongeveer 3000 functie-evaluaties, wat laat zien dat het modelraamwerk geschikt is voor casestudies van realistische omvang. Het resultaat, de Pareto-set, wordt geanalyseerd met behulp van de genoemde methoden voor beslisondersteuning. De gekozen maatregelen (gerelateerd aan multimodale ritten) zijn in staat een reductie te realiseren ten opzichte van het huidige netwerk voor alle vier de doelen tegelijkertijd. Eén specifiek netwerkontwerp laat een reductie zien van ten minste 0,4% voor totale reistijd, ruimtebeslag in steden en CO_2-uitstoot, tegen 0,4% lagere kosten. Voor individuele doelen zijn grotere verbeteringen mogelijk: de CO_2-uitstoot kan met 0,45% worden verminderd, de totale reistijd met 0,79% en het ruimtebeslag in steden met 1,5%, maar dan gaat er wel ten minste één ander doel op achteruit. Absoluut gezien kan er elke ochtendspits bijna 4000 uur reistijd, ruimtebeslag van 2000 geparkeerde auto's of meer dan 12 ton CO_2-uitstoot (gelijk aan de dagelijkse directe CO_2-uitstoot van meer dan 500 Nederlandse huishoudens) worden bespaard. Totale reistijd, ruimtebeslag en CO_2-uitstoot kunnen over het algemeen tegelijk worden verminderd, doorgaans tegen een minder gunstig exploitatiesaldo. Toch is het niet mogelijk een oplossing te vinden die optimaal is voor zowel reistijd, ruimtebeslag als CO_2-uitstoot. De gemiddelde trade-off waarde is €6,90 per uur gebleken, wat betekent dat gemiddeld één uur reistijd bespaard kan worden door €6,90 extra te besteden aan OV-exploitatie. Deze waarde is kleiner dan de tijdwaardering die in Nederlandse kosten-baten-analyses wordt gebruikt, dus vanuit het oogpunt van de regionale overheid bestaan er kosteneffectieve maatregelen gerelateerd multimodale ritten. Hierbij moet worden opgemerkt dat investeringen in infrastructuur niet zijn meegenomen in de analyse, omdat deze betaald worden door de nationale overheid. Hierdoor is het niet zeker of de maatregelen maatschappelijk kosteneffectief zijn. Om de CO_2-uitstoot met één kilo terug te dringen, blijkt dat €1,10 nodig is. Dit is veel meer dan alternatieve maatregelen uit de literatuur om de CO_2-uitstoot van het transportsysteem te verminderen, zoals strengere eisen aan CO_2-uitstoot voor personenauto's (€0,11 per kilo) of het stimuleren van elektrische auto's (€0,35 per kilo). Het gecombineerde effect van extra uitgaven op reistijd, ruimtebeslag en CO_2-uitstoot is als volgt: elke extra euro exploitatiebijdrage leidt gemiddeld tot ongeveer 6 minuten minder totale reistijd, 0,07 minder geparkeerde auto en 0,5 kilo minder CO_2-uitstoot. De min-max-oplossing voor alle vier de doelen valt binnen de beste 30% van het bereik van elk van de doelen, dus indien de genormaliseerde doelwaarden gelijk worden gewaardeerd, is het mogelijk om aan alle vier de doelstellingen in grote mate te voldoen.

In het algemeen leidt het openen van P&R-voorzieningen tot een lagere totale reistijd, maar ook tot meer CO_2-uitstoot. Dit komt doordat niet alleen voormalige autoritten worden vervangen door P&R-ritten, maar ook voormalige volledige OV-ritten. Verder heeft het verhogen van treinfrequenties, en in mindere mate van busfrequenties, in het algemeen een

positieve invloed op zowel totale reistijd, ruimtebeslag als CO_2-uitstoot, maar is wel een hogere exploitatiebijdrage nodig. Het verhogen van frequenties blijkt een groter effect te hebben op de duurzaamheidsdoelstellingen dan het openen van nieuwe P&R-voorzieningen en treinstations. Dit duidt erop dat multimodale ritten beter gestimuleerd kunnen worden door de kwaliteit van het OV-systeem zelf te verbeteren, dan door het toevoegen van extra overstaplocaties. Een andere reden om het verhogen van frequenties te verkiezen boven het toevoegen van overstapvoorzieningen is dat frequentieverhogingen makkelijker zijn terug te draaien, waardoor deze maatregelen zich beter kunnen aanpassen aan eventuele toekomstige veranderingen. Indien een Pareto-optimaal netwerkontwerp een lage frequentie heeft op een bepaalde Sprinter-lijn, komt het vaak voor dat stations op die lijn Intercity-status krijgen, om op deze stations te compenseren voor de lage frequentie. Op trein- en op buslijnen komen tegelijkertijd hoge frequenties voor in Pareto-optimale ontwerpen, wat erop duidt dat treinen en bussen elkaar in deze casestudie aanvullen. Tenslotte komt slechts één van de 37 mogelijke maatregelen in de casestudie nooit voor in de Pareto-set. Dit mogelijk nieuwe station zal dus nooit geopend hoeven te worden bij het nastreven van de vier bestudeerde doelen. Andersom zijn er geen maatregelen gevonden die altijd genomen moeten worden, welk van de vier doelen men ook belangrijk vindt. Dit laat zien dat het feit of een maatregel optimaal presteert sterk afhangt van welk doel men nastreeft.

Dit onderzoek laat zien dat maatregelen die gericht zijn op het stimuleren van multimodale ritten kunnen bijdragen aan duurzaamheidsdoelstellingen door OV-reisalternatieven aantrekkelijker te maken. Aan de andere kant is het oplossend vermogen van dit type maatregelen beperkt, vooral voor CO_2-uitstoot, omdat het reisalternatief met hoge CO_2-uitstoot (in dit geval de autorit) aantrekkelijk blijft voor de reiziger. Daarom is de invloed van netwerkbrede oplossingen die de auto minder aantrekkelijk maken verkend. Hiervoor is het zelfde verkeersmodel gebruikt, ook al houdt dit model geen rekening met gedragsveranderingen die eigenlijk wel relevant worden voor dit type maatregelen, zoals veranderende verkeersvraag. De verkenning laat zien dat een vaste kilometerheffing een veel grotere reductie van het ruimtebeslag in steden en van de CO_2-uitstoot laat zien. Wel neemt de totale reistijd toe, omdat voormalige autogebruikers nu een langzamer OV-alternatief kiezen. Ook is het effect van alternatieve ruimtelijke planning verkend, het zogenaamde 'transit-oriented development'. Hierbij wordt ervan uitgegaan dat alle nieuwe inwoners en banen in de regio tot 2030 ontwikkeld worden in de buurt van treinstations. Dit leidt tot een aanzienlijke verbetering van de CO_2-uitstoot en ook tot een kleine afname van de vertragingstijd.

Robuustheid op de lange termijn

Een vervoersnetwerk wordt ontworpen voor een lange tijdsperiode en dus is het nodig om de vraag naar vervoer te voorspellen. Het is belangrijk om te weten in welke mate de netwerkontwerpen robuust zijn voor onzekerheden in deze vervoersprognose. Daarom worden er naast Pareto-sets voor de standaardprognose voor 2030 ook Pareto-sets gemaakt voor twee andere prognoses: de prognose voor 2020 als een ondergrens van de bandbreedte voor OV-gebruik en een prognose voor transit-oriented development als een bovengrens van de bandbreedte.

Het blijkt dat netwerkontwerpen die optimaal zijn voor de 2030 vraagprognose ook behoorlijk goed presteren indien de vraag lager is of veranderd is door transit-oriented development. Het verlies in doelwaarde is klein en de waarden van de beslisvariabelen worden niet sterk beïnvloed door veranderende vervoersvraag. Een belangrijke reden daarvoor is dat alle drie de gebruikte vraagprognoses realistisch zijn met een vergelijkbare structuur: in het netwerk-ontwerpprobleem zijn die maatregelen die de grootste stromen beïnvloeden doorgaans het meest effectief, en deze grote stromen bevinden zich in alle scenario's ongeveer op dezelfde relaties. Voor het geval dat een maatregel slecht blijkt te presteren door veranderende omstandigheden is het goed onderscheid te maken tussen maatregelen die makkelijk of moeilijk terug te draaien zijn. Tenslotte blijkt dat het loslaten van CO_2-uitstoot als doelstelling geen grote invloed heeft op het resultaat, omdat dit doel grotendeels met dezelfde maatregelen bereikt kan worden als de doelen ruimtebeslag en totale reistijd.

Aanbevelingen voor toekomstig onderzoek

Het verkeersmodel in het lage niveau van het optimalisatieprobleem kan op verschillende manieren worden uitgebreid om de stromen in het netwerk nauwkeuriger te kunnen berekenen. Ten eerste kunnen de technieken om gecombineerde vervoerwijze- en routekeuze te modelleren verder worden verbeterd door de gedragsmatig ingewikkelde multimodale rit beter te beschrijven. Ten tweede kunnen de attributen van de gegeneraliseerde kostenfunctie uitgebreid worden door bijvoorbeeld comfort of drukte mee te nemen. Een derde mogelijkheid is het meenemen van (dagelijkse) activiteitenpatronen in de modellen, wat het mogelijk maakt rekening te houden met aspecten als voertuigbeschikbaarheid en de invloed van meerdere ritten binnen een tour op elkaar. Het wordt verwacht dat een uitgebreider model in het lage niveau tot betere optimalisatieresultaten leidt, maar dit is niet bekend. Het kan ook zijn dat eenvoudigere modellen uiteindelijk tot betere optimalisatieresultaten leiden, omdat minder rekentijd voor het model in het lage niveau betekent dat er meer functie-evaluaties kunnen plaatsvinden. Er is meer onderzoek nodig om uit te zoeken of de conceptuele beperkingen van het gebruikte model beperkt invloed hebben op netwerkniveau (en dus leiden tot de zelfde optimale netwerkontwerpen) of de resultaten systematisch beïnvloeden (en dus leiden tot andere optimale netwerken).

Voor de oplossingsmethode ε-NSGAII is het nodig meer vergelijkbare experimenten uit te voeren op andere casestudies uit de praktijk, omdat het algoritme nog niet vaak is toegepast en anders kan presteren voor elke casestudie. Ook zou het kunnen dat er andere heuristieken bestaan die beter presteren dan ε-NSGAII, gegeven de snelle ontwikkelingen in het onderzoeksgebied van GA's. Om meer algemene conclusies te kunnen trekken over de effecten van de stochastische processen inherent aan GA's op de uitkomst van één enkele run, is het nodig om meer casestudies te onderzoeken. In toekomstige optimalisatiestudies met meerdere doelstellingen is het direct mogelijk om de in dit onderzoek toegepaste methodes toe te passen om de optimalisatieresultaten overzichtelijk te presenteren. Meer onderzoek is echter nodig om de ingewikkelde resultaten van optimalisatie met meerdere doelstellingen eenvoudig te presenteren, omdat deze resultaten nu vaak ingewikkeld zijn om aan beslissers uit te leggen.

Er zijn vele stakeholders betrokkenen bij het exploiteren van OV. Het modelraamwerk in dit onderzoek is toegepast voor de belangen van één enkele stakeholder. Voor de gehele

maatschappij zou het beter zijn om het OV-netwerk te optimaliseren als één geïntegreerd netwerk, rekening houdend met de effecten voor alle relevante overheden en andere stakeholders in de regio. Dit heeft te maken met de gebruikte kostenfunctie, waarin alleen de kosten voor de regionale overheid zijn meegenomen en de kosten voor de nationale overheid genegeerd. Daarnaast zijn de vier doelstellingen nu geformuleerd voor het totale netwerk, terwijl het relevant is of de positieve effecten van de netwerkontwerpen op een eerlijke manier over de samenleving worden verdeeld. Dit kan bijvoorbeeld gaan over geografische verdeling, maar ook over de verdeling over gebruikers van verschillende vervoerwijzen. Door technologische ontwikkelingen kunnen de effecten van maatregelen veranderen, waardoor doelfuncties anders gedefinieerd worden, bijvoorbeeld door elektrische auto's of waterstofauto's. Ook met de beschikbaarheid van nieuwe vervoerwijzen, zoals de elektrische fiets, kan rekening gehouden worden in de modellen. Dit kan leiden tot andere netwerkontwerpen: het is relevant om te weten in welke mate deze ontwikkelingen het optimalisatieresultaat beïnvloeden.

De oplossingsrichting die in dit onderzoek is behandeld, is slechts een stap op de lange weg die nog te gaan is voordat er een duurzaam vervoerssysteem in de Randstad bestaat. Indien autoritten minder aantrekkelijk worden gemaakt en tegelijkertijd genoeg capaciteit en kwaliteit wordt geboden in het OV-netwerk (inclusief de verbindingen tussen het OV-netwerk en de auto- en fietsnetwerken), zouden er flinke verbeteringen voor milieu gerelateerde duurzaamheid mogelijk zijn zonder economische duurzaamheid in de weg te staan. Daarom zou uitbreiding van het multimodale netwerkontwerpprobleem met andere typen maatregelen nuttig zijn, bijvoorbeeld nieuwe spoorinfrastructuur of kilometerheffing. In dat geval is het nodig om het model dat de gedragsreactie van reizigers bepaalt uit te breiden, waardoor de aanname van een vaste vervoersvraag niet meer nodig is. Verder kunnen combinaties gezocht worden met ander type maatregelen, bijvoorbeeld technologische ontwikkelingen om de milieuprestatie van voertuigen te verbeteren, zoals het introduceren van elektrische auto's, of veranderingen in de planning van ruimtelijke ontwikkelingen.

About the author

Ties Brands was born on the 28th of April, 1983. He grew up in Gemert, Noord-Brabant, The Netherlands. After secondary school, he moved to Enschede to study at the University of Twente. He obtained his BSc degree (Civil Engineering) in 2004 and his two MSc degrees (Civil Engineering and Applied Mathematics) in 2008. His MSc thesis was on optimisation of dynamic road pricing and was conducted at the Dutch consultancy company Goudappel Coffeng in Deventer. After graduation, Ties started working at Goudappel Coffeng as a public transport consultant, specialised in public transport modelling and data analysis. In 2010 he started, parallel to his work as a consultant, a PhD project at the University of Twente, Centre for Transport Studies. He still works at Goudappel Coffeng, mainly on projects related to public transportation, for example network planning studies, ridership predictions, cost-benefit analyses and data analyses.

Author's publications

Journal publications

Van Oort, N., T. Brands and E. de Romph (in press) Short term ridership prediction in public transport by processing smart card data. *Transportation Research Record.*

Van Oort, N., D. Sparing, T. Brands and R.M.P. Goverde (in press) Efficiency and quality improvement in public transport by applying big data: the Dutch example. *Public Tranport.*

Van Oort, N., T. Brands, E. de Romph and J.A. Flores (2015) Unreliability effects in public transport modelling, *International Journal of Transportation* 3(1), pp.113-130.

Brands T. and E. C. van Berkum (2014) Performance of a genetic algorithm for solving the multi-objective, multimodal transportation network design problem, *International Journal of Transportation* 2(1), pp. 1-20.

Brands T., E. de Romph, T. Veitch and J. Cook (2014) Modelling public transport route choice with multiple access and egress modes. *Transportation Research Procedia* 1, pp. 12-23.

Eck, G. van, T. Brands, L.J.J. Wismans, A.J. Pel, and R. van Nes (2014) Model complexities and requirements for multimodal transport network design: Assessment of classical, state-of-the-practice, and state-of-the-research models. *Transportation research record* 2429, pp. 178-187.

Oort, N. van, R. van Leusden and T. Brands (2014) Verkeersmodellen verrijken met onbetrouwbaarheid OV vanuit een reizigersperspectief (in Dutch). *Tijdschrift vervoerswetenschap* 50(3), pp. 3-18.

Wismans, L.J.J., Brands, T., Berkum, E.C. van & Bliemer, M.C.J. (2014) Pruning and ranking the Pareto optimal set, application for the dynamic multi-objective network design problem. *Journal of advanced transportation* 48(6), pp. 588-607.

Peer reviewed conference proceedings

Brands T., L. J. J. Wismans and E. C. van Berkum (2015) The Pareto set as decision support information in multimodal passenger transportation network design. *13th Conference on Advanced Systems in Public Transport*, 19-23 July 2015, Rotterdam, The Netherlands.

Oort, N. van, M. Drost, T. Brands and M. Yap (2015) Data-driven public transport ridership prediction approach including comfort aspects. *13th Conference on Advanced Systems in Public Transport*, 19-23 July 2015, Rotterdam, The Netherlands.

Brands T., L. J. J. Wismans and E. C. van Berkum (2014) Multi-objective Transportation Network Design: Accellerating Search by Applying ε-NSGAII. *IEEE Congress on Evolutionary Computation*, Beijing, China, IEEE Press, pp. 405-412.

Brands T., L. J. J. Wismans and E. C. van Berkum (2014) Multi-Objective Optimization of Multimodal Passenger Transportation Networks: Coping with Demand Uncertainty. *OPT-i, 1st International Conference on Engineering and Applied Sciences Optimization*. M. Papadrakakis, M. G. Karlaftis and N. D. Lagaros. Kos Island, Greece, pp. 547-561.

Oort, N. van, T. Brands, E. de Romph and J.A. Aceves Flores (2014) Incorporating unreliability of transit in transport demand models: theoretical and practical approach. *Proceedings of 93th Annual Meeting of the Transportation Research Board*, January 12-16 2014, Washington DC.

Brands, T. (2014) Using results form multi-objective optimization as decision support information in multimodal passenger transportation network design. *Proceedings of the 2nd Trail PhD conference*, November 13th, 2014, Delft, The Netherlands.

Ohazulike, A. E. and T. Brands (2013) Multi-objective optimization of traffic externalities using tolls: a comparison of genetic algorithm with game theoretical approach. *IEEE Congress on Evolutionary Computation*, Cancún, Mexico, pp. 2465-2472.

Brands, T. and E.C. van Berkum (2013) Multi-objective optimization of multimodal transportation networks: interpretation of the Pareto set. *Proceedings of the BIVEC-GIBET Transport Research Day*, 30-31th of May, 2013.

Oort, N. van, D. Sparing, T. Brands and R. M. P. Goverde (2013) Optimizing public transport planning and operations using automatic vehicle location data: the Dutch example. *Proceedings of the 3rd International Conference on Models and Technologies for Intelligent Transport Systems*, Dresden, Germany, pp. 291-300.

Brands, T. (2013) Multi-objective, multimodal transportation network design: robustness analysis. *Proceedings of the 1ˢᵗ Trail PhD conference*, November 14ᵗʰ, Delft, The Netherlands.

Brands, T. and E.C. van Berkum (2012) Multi-objective optimization of multimodal transportation networks. *Proceedings of the 1st European Symposium on Quantitative Methods in Transportation Systems*, 4-7th of September 2012, Lausanne, Switserland.

Brands, T. and E.C. van Berkum (2012) Multi-objective optimization of multimodal transportation networks: interpretation of the Pareto set from a case study in Amsterdam. *Proceedings of the 12ᵗʰ Trail congress*, October, 2012.

Brands, T. & G. van Eck (2010) Multimodal network design and assessment: proposal for a dynamic multi objective approach. *Proceedings of the 11th Trail Congress: Connecting People - Integrating Expertise*, 23-24 November, 2010.

Brands, T., E. C. van Berkum, and D. H. van Amelsfort (2009) Optimal toll design in dynamic traffic networks using a pattern search approximation algorithm, *88th annual meeting of the Transportation Research Board*, Washington D.C.

Non-peer reviewed conference proceedings

Van Oort, N., M. Drost, T. Brands (2014) Betere OV prognoses met anonieme OV chipkaartdata (in Dutch), *Colloquium vervoersplanologisch speurwerk,* Eindhoven.

Oort, N. van, Leusden, R. & Brands, T. (2013) Betrouwbare modellen door onbetrouwbaar OV (in Dutch), *Colloquium Vervoersplanologisch Speurwerk*, 21-22 November 2013, Rotterdam.

Wismans, L.J.J., Brands, T. and E.C. van Berkum (2011) Decision support by pruning and ranking the Pareto optimal set. *Proceedings of OR2011*, Zürich, Switzerland, 30ᵗʰ of August – 2ⁿᵈ of September, 2011.

Brands, T. and E.C. van Berkum (2011) Multi-objective, multimodal network design, *Proceedings of IFORS2011*, Melbourne, Australia, 10ᵗʰ – 15ᵗʰ of July 2011.

Brands, T. and E.C. van Berkum (2010) Optimalisatie van het multimodale vervoersnetwerk in de Randstad, rekening houdend met meerdere doelstellingen (in Dutch), *Colloquium vervoersplanologisch speurwerk*, Roermond, 25-26 november 2010.

Brands, T., D.H. Van Amelsfort and E.C. van Berkum (2008) Optimising differentiated tolls on large scale networks, by using an intellegent search algorithm, *European transport conference*, Noordwijkerhout, The Netherlands.

Presentations (non-peer reviewed)

Brands, Ties, Gijs van Eck and Daniel Sparing (2013) Optimalisatie van het multimodale vervoersnetwerk in de Randstad, rekening houdend met meerdere doelstellingen (in Dutch), *presented at KIVI NIRIA year congress 2013*, November 6th, Eindhoven, The Netherlands

Brands, Ties (2014) Modellen als hulpmiddel bij het ontwerpen van een optimaal multimodaal verkeersnetwerk (In Dutch), *presented at PLATOS colloquium*, 5th of March, Den Haag, The Netherlands.

Brands, T. and E.C. van Berkum (2014) Sustainable multimodal passenger transportation networks: scientific results. *Presented at the Verdus congress 'Vitality in City and Region'*, 16th of June, Rotterdam, The Netherlands.

TRAIL Thesis Series

The following list contains the most recent dissertations in the TRAIL Thesis Series. For a complete overview of more than 100 titles see the TRAIL website: www.rsTRAIL.nl.

The TRAIL Thesis Series is a series of the Netherlands TRAIL Research School on transport, infrastructure and logistics.

Brands, T., *Multi-objective Optimisation of Multimodal Passenger Transportation Networks,* T2015/15, October 2015, TRAIL Thesis Series, the Netherlands

Ardiç, Özgül, *Road Pricing Policy Process: The interplay between policy actors, the media and public,* T2015/14, September 2015, TRAIL Thesis Series, the Netherlands

Xin, J., *Control and Coordination for Automated Container Terminals,* T2015/13, September 2015, TRAIL Thesis Series, the Netherlands

Anand, N., *An Agent Based Modelling Approach for Multi-Stakeholder Analysis of City Logistics Solutions,* T2015/12, September 2015, TRAIL Thesis Series, the Netherlands

Hurk, E. van der, *Passengers, Information, and Disruptions,* T2015/11, June 2015, TRAIL Thesis Series, the Netherlands

Davydenko, I., *Logistics Chains in Freight Transport Modelling,* T2015/10, May 2015, TRAIL Thesis Series, the Netherlands

Schakel, W., *Development, Simulation and Evaluation of In-car Advice on Headway, Speed and Lane,* T2015/9, May 2015, TRAIL Thesis Series, the Netherlands

Dorsser, J.C.M. van, *Very Long Term Development of the Dutch Inland Waterway Transport System: Policy analysis, transport projections, shipping scenarios, and a new perspective on economic growth and future discounting,* T2015/8, May 2015, TRAIL Thesis Series, the Netherlands

Hajiahmadi, M., *Optimal and Robust Switching Control Strategies: Theory, and applications in traffic management,* T2015/7, April 2015, TRAIL Thesis Series, the Netherlands

Wang, Y., *On-line Distributed Prediction and Control for a Large-scale Traffic Network*, T2015/6, March 2015, TRAIL Thesis Series, the Netherlands

Vreeswijk, J.D., *The Dynamics of User Perception, Decision Making and Route Choice,* T2015/5, February 2015, TRAIL Thesis Series, the Netherlands

Lu, R., *The Effects of Information and Communication Technologies on Accessibility,* T2015/4, February 2015, TRAIL Thesis Series, the Netherlands

Ramos, G. de, *Dynamic Route Choice Modelling of the Effects of Travel Information using RP Data,* T2015/3, February 2015, TRAIL Thesis Series, the Netherlands

Sierzchula, W.S., *Development and Early Adoption of Electric Vehicles: Understanding the tempest,* T2015/2, January 2015, TRAIL Thesis Series, the Netherlands

Vianen, T. van, *Simulation-integrated Design of Dry Bulk Terminals,* T2015/1, January 2015, TRAIL Thesis Series, the Netherlands

Risto, M., *Cooperative In-Vehicle Advice: A study into drivers' ability and willingness to follow tactical driver advice,* T2014/10, December 2014, TRAIL Thesis Series, the Netherlands

Djukic, T., *Dynamic OD Demand Estimation and Prediction for Dynamic Traffic Management,* T2014/9, November 2014, TRAIL Thesis Series, the Netherlands

Chen, C., *Task Complexity and Time Pressure: Impacts on activity-travel choices,* T2014/8, November 2014, TRAIL Thesis Series, the Netherlands

Wang, Y., *Optimal Trajectory Planning and Train Scheduling for Railway Systems,* T2014/7, November 2014, TRAIL Thesis Series, the Netherlands

Wang, M., *Generic Model Predictive Control Framework for Advanced Driver Assistance Systems,* T2014/6, October 2014, TRAIL Thesis Series, the Netherlands

Kecman, P., *Models for Predictive Railway Traffic Management,* T2014/5, October 2014, TRAIL Thesis Series, the Netherlands

Davarynejad, M., *Deploying Evolutionary Metaheuristics for Global Optimization,* T2014/4, June 2014, TRAIL Thesis Series, the Netherlands

Li, J., *Characteristics of Chinese Driver Behavior*, T2014/3, June 2014, TRAIL Thesis Series, the Netherlands

Mouter, N., *Cost-Benefit Analysis in Practice: A study of the way Cost-Benefit Analysis is perceived by key actors in the Dutch appraisal practice for spatial-infrastructure projects*, T2014/2, June 2014, TRAIL Thesis Series, the Netherlands

Ohazulike, A., *Road Pricing mechanism: A game theoretic and multi-level approach,* T2014/1, January 2014, TRAIL Thesis Series, the Netherlands

Cranenburgh, S. van, *Vacation Travel Behaviour in a Very Different Future,* T2013/12, November 2013, TRAIL Thesis Series, the Netherlands

Samsura, D.A.A., *Games and the City: Applying game-theoretical approaches to land and property development analysis,* T2013/11, November 2013, TRAIL Thesis Series, the Netherlands

Huijts, N., *Sustainable Energy Technology Acceptance: A psychological perspective*, T2013/10, September 2013, TRAIL Thesis Series, the Netherlands

TRAIL

Summary

The resolution of measures related to multimodal trip making and corresponding network designs are identified by solving a multi-objective optimisation problem. Methods to analyse the resulting Pareto set provide insight into how total travel time, CO_2 emissions, urban space used by parking and costs relate. In the Randstad case study increasing frequencies appears to be more effective to improve sustainability than introducing P&R facilities and train stations.

About the Author

Ties Brands carried out his PhD research from 2010 to 2015 at the University of Twente, Centre for Transport Studies. In parallel, he has worked as a public transportation consultant at Goudappel Coffeng since 2008.

ERASMUS UNIVERSITY ROTTERDAM

Radboud University Nijmegen

TU Delft
Delft University of Technology

UNIVERSITY OF TWENTE.

TU/e
Technische Universiteit Eindhoven
University of Technology

university of groningen
faculty of economics and business