The Pareto set as decision support information in multimodal passenger transportation network design

Ties Brands · Luc J.J. Wismans · Eric C. van Berkum

Abstract Given a range of traffic related sustainability problems, policy makers need to know which measures should be taken to reach their objectives as much as possible. Multi-objective optimisation is useful to support these decisions, because it results in an overview of possibly optimal solutions. This Pareto set can be very large, especially if more than two (mainly opposed) objectives are involved. This is also the case when optimising infrastructure planning in a multimodal passenger transportation network, with accessibility, use of urban space by parking, operating deficit and climate impact as objectives. Methods are presented to derive problem knowledge from the Pareto set. This includes the best values per objective, average trade-off values between pairs of objectives and identification of the min-max solution. These methods make the Pareto set more useful as decision support information: they demonstrate the next step in multi-objective option prioritisation. An insight provided for a case study in the Amsterdam Metropolitan Area in The Netherlands is that 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. Next, improving travel time further can be done cost-efficiently, but reducing CO_2 emissions is expensive when using measures related to multimodal trip making. Finally, increasing frequencies appears to be more effective to improve sustainability than introducing P&R facilities and train stations.

Ties Brands
Centre for Transport Studies, Faculty of Engineering Technology, University of Twente,
P.O. Box 217, 7500AE Enschede, The Netherlands
Goudappel Coffeng Mobility Consultants, Deventer, The Netherlands
Email: tbrands@goudappel.nl

Luc J.J. Wismans
Centre for Transport Studies, Faculty of Engineering Technology, University of Twente,
P.O. Box 217, 7500AE Enschede, The Netherlands
DAT.Mobility, Deventer, The Netherlands
Email: lwismans@dat.nl

Eric C. van Berkum
Centre for Transport Studies, Faculty of Engineering Technology, University of Twente,
P.O. Box 217, 7500AE Enschede, The Netherlands
Email: e.c.vanberkum@utwente.nl

1 Introduction

Highly urbanised regions in the world nowadays face well known sustainability problems in the traffic system, like congestion, use of scarce space in cities by vehicles and the emission of greenhouse gases. This research focusses on the integration of transportation networks of cars, public transport (PT, which includes bus, tram, metro and train) and bicycles as a cost effective solution direction to alleviate these problems.

When the infrastructure related to these measures is planned by decision makers, it is often current practice to design a few alternatives, have these alternatives assessed by a transportation model and choose the best performing alternative. However, the alternative is still likely to have room for improvement. This is the reason for applying optimisation techniques in this context. A multi-objective approach is adopted, because of the complex context of competing sustainability interests, like accessibility, liveability, environmental impact and costs. This results in a multi-objective, multimodal passenger network design problem. An extensive review of passenger network design problems is given in Farahani et al. (2013). In this review a classification of network design studies in literature is given, for example distinguishing between single-objective and multi-objective network design problems and between problems that involve a car network, a PT network or a multimodal network.

The Pareto-optimal set is the outcome of such 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 paper the information needs of decision makers are identified. 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). 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). Such trade-offs are determined in this paper for a case study, 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 multimodal passenger transportation network design problem, which has many discrete variables. In this paper 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. 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. 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.

Following the needs of decision makers to choose one final solution for implementation, the objective of this paper is twofold. The first objective is to enhance existing methods for decision support based on results from multi-objective optimisation, which help decision makers choose a final solution. The second objective is to analyse and understand a case study problem using information in the Pareto set.

The remainder of this paper is structured as follows. In Section 2 the problem is defined and the case study is introduced. In Section 3 the results of the case study are presented, using various methods for decision support. First, the Pareto set is visualised. Second, more general rules are derived concerning relations between objectives and measures. Third, a step-by-step pruning method is applied to choose one final solution for implementation. In Section 4 conclusions are drawn concerning the methods and the case study.

2 Problem definition

2.1 Multi-objective optimisation

The chosen multi-objective approach enables coping with the complex context of competing sustainability interests, like accessibility, liveability, environmental impact and costs. These w objectives Z_w are not translated into a single objective by using weights for each objective, because the weights as well as the normalisation of the different objectives are arbitrary. Furthermore, tradeoffs between objectives can only be achieved by studying the Pareto optimal set (Coello et al., 2006). The entire decision space Y is searched during optimisation. Mathematically, the concept of Pareto optimality is as follows. Assuming two decision vectors $\underline{y}_i, \underline{y}_{i\prime} \in Y$, then $\underline{y}_i$ is said to (weakly) dominate (or cover) $\underline{y}_{i\prime}$ iff $Z_w(\underline{y}_i) \leq Z_w(\underline{y}_{i\prime}) \forall w$ (also written as $\underline{y}_i \preccurlyeq \underline{y}_{i\prime}$). All solutions that are not weakly dominated by another known solution are possibly optimal for the decision maker: these non-dominated solutions form the Pareto-optimal set P.

2.2 Bi-level problem

The transportation network design problem is solved as a bi-level optimisation problem (see also for example dell'Olio et al., 2006). The upper level represents a

network authority that wants to optimise system objectives. In the lower level the behavioural response to a network design is modelled, where the travellers minimise their own generalised costs in the multimodal network. To this end, a combined modal split / assignment model is applied in the lower level (see Brands et al., 2014a). The resulting equilibrium is a constraint for the upper level problem.

2.3 Study area

The case study area covers the Amsterdam Metropolitan Area in The Netherlands (Fig. 1). 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 trains, regional trains and intercity trains. Bicycles can be parked at most bus stops and at all train stations. A selection of PT stops facilitates park-and-ride transfers. Origins and destinations are aggregated into 102 transportation zones.

Fig. 1 Map of the study area, showing origins / destinations, railways, roads

2.3.1 Decision variables
In the network of the study area, 37 decision variables are defined related to transfer facilities or to PT facilities. For every potential network development, a decision variable is defined in advance. Opening / closure of train stations, intercity status of train stations and opening / closure of park and ride (P&R) facilities are represented by binary variables. For PT line frequency, a discrete set of choice options is predefined, depending on the expected load for that transit line. The characteristics of links, 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 resulting feasible region Y contains approximately $7*10^{13}$ possible decision vectors.

2.3.2 Objective functions
The values of the objective functions are calculated based on loads and costs in the network for one hour in the AM peak. Four objectives are considered, namely total travel time (TTT, as a measure for accessibility), number of car trips to urban zones (as a measure for use of urban space for parking, USU), CO_2 emissions (CE, as a measure for climate impact) and PT operating deficit (OpD, as a measure for cost efficiency). All objectives are to be minimised.

2.4 Solution method

The optimisation problem is solved using the evolutionary algorithm ε-NSGAII (Kollat and Reed, 2006). This method was earlier shown to outperform the well-known predecessor of the algorithm NSGAII (Deb et al., 2002) when applied to the same case study in Brands et al. (2014b). This better performance especially occurs when limited number of function evaluations are applied. This is useful, because the lower-level transportation model in the case study has a high computation time.

3 Results

The optimisation algorithm produces a Pareto set as a result. From the Pareto set more general problem knowledge is derived. This helps to better understand the network design problem in a multimodal context and to finally choose one solution from the Pareto set for implementation in a multi-objective option prioritisation process.

In total 2384 solutions were evaluated during the execution of the algorithm. From these solutions 210 were Pareto optimal (i.e. non-dominated). This Pareto set is an approximation of the true Pareto set, since it would take too much computation time to calculate all solutions and thus the true Pareto set is not known.

3.1 Visualisation of the Pareto set

The scatter plot shown in Fig. 2 is a common way to visualise a Pareto set, especially to show trade-offs between objectives. By using different colours for the dots in a scatter plot, one additional objective is included in the scatter plot (see Fig. 2). 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.

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), 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.
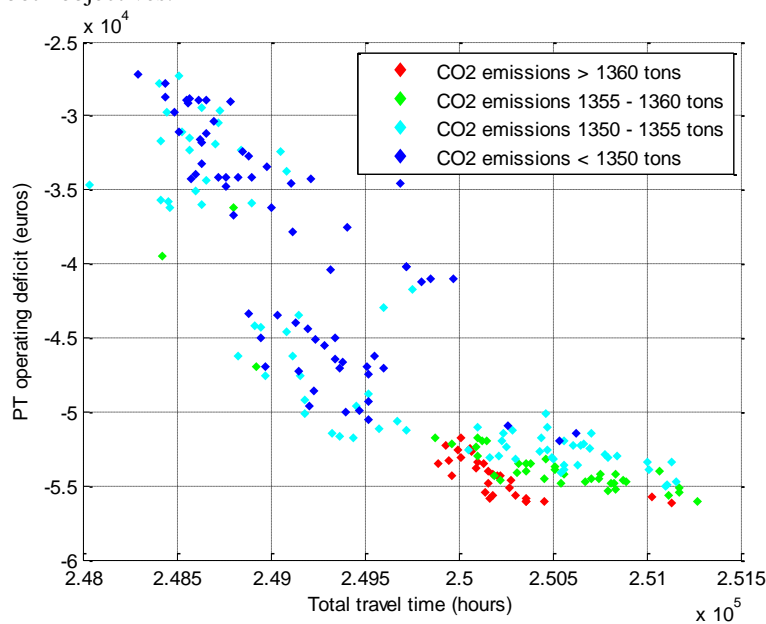

Fig. 2 A 2D scatterplot for total travel time and operating deficit, with a distinction between CE categories (also called decision map)

A parallel coordinate plot (see Fig. 3) 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. In a parallel coordinate plot trade-offs between objectives cannot be directly observed like in a scatter plot, but 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.
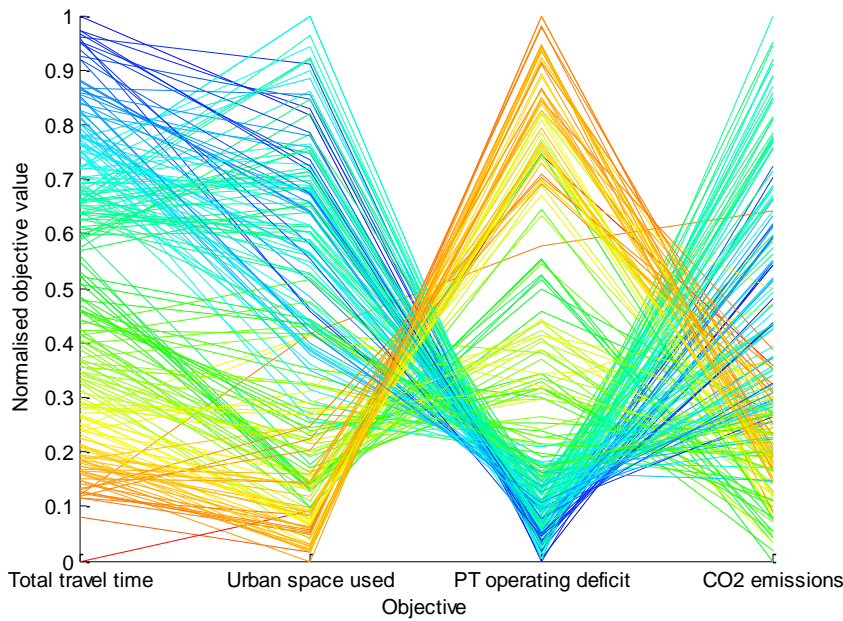
Fig. 3 Parallel coordinate plot of the Pareto set: one line represents one solution, where the normalised values of the 4 objective values are plotted in the 4 columns. The lines are coloured using their value for total travel time

3.2 Best value per objective

The extent to which the formulated objective values can be improved indicates to what extent a more sustainable transportation network can be established by the selected measures. Therefore the minimum values per objective are compared with the objective values in the base solution (see Table 1). The base solution is the most likely transportation network that will be developed for 2030 when all known plans are realised. The relative value for OpD is 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, that can be both positive and negative).

	TTT (hours)	USU (# of cars)	OpD (euros)	CE (tons)
Base situation	250007	68720	-36594	1352
Minimum	248034	67659	-56203	1346
Improvement w.r.t. base	-0.79%	-1.54%	-6.19%	-0.45%

Table 1 Minimum values found per objective, compared to the base situation

In Fig. 4 in entire Pareto set is visualised in a parallel coordinate plot, where the optimal solutions per objective are highlighted in blue. The base solution is shown in black. 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. One 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).
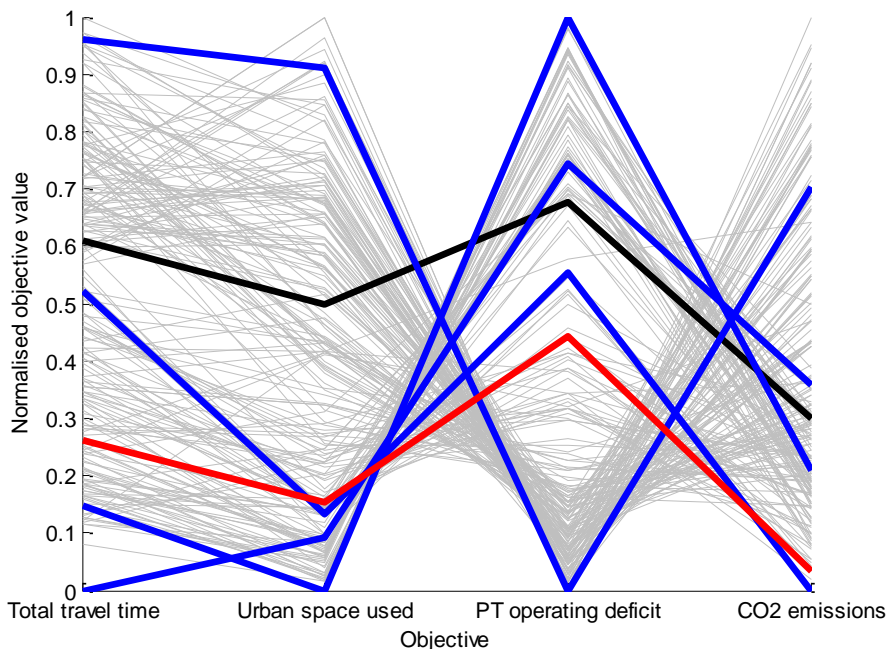

Fig. 4 Parallel coordinate plot, including the optimal solutions per objective (blue), the base solution (black), one solution that gives considerable improvement with respect to the base solution in all four objectives (red) and all other Pareto optimal solutions (grey).

3.3 Interdependencies between objectives

3.3.1 Correlation between objectives

In Table 2 a 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.

	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

Table 2 Correlation between pairs of objectives

3.3.2 Trade-off values between pairs of objectives

A trade-off value indicates the extent to which a deterioration for an objective has to be accepted, if a decision maker likes to improve anther objective by moving to another solution (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. By introducing the average trade-off value, one value is provided for trade-off between a pair of objectives, along the entire (known) Pareto front (see Fig. 5). In Eq. 1 average trade-off is formally defined.

$$ATO_{w,w'}(P) = \frac{Z_w(\underset{\underline{y} \in P}{\operatorname{argmin}} Z_{w'}(\underline{y})) - \underset{\underline{y} \in P}{\min} Z_w(\underline{y})}{Z_w(\underset{\underline{y} \in P}{\operatorname{argmin}} Z_w(\underline{y})) - \underset{\underline{y} \in P}{\min} Z_{w'}(\underline{y})} \qquad (1)$$
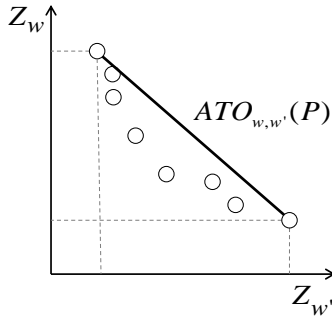

Fig. 5 The average trade-off value, based on the extreme solutions for a pair of objectives

The trade-off results for the case study are shown in Table 3. 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 (which is slightly lower than the value of time in The Netherlands). It is likely that more efficient investments than the average are possible by selecting the right solutions. Note that the other two objectives (very likely) also change values and result in additional benefits or costs, which are disregarded in this table.

$w \setminus w'$	TTT (hours)	USU (# of cars)	OpD (euros)	CE (kilos)
TTT (hours)		-2.44	-0.144	-0.232
USU (# of cars)	-0.411		-0.067	-0.0661
OpD (euros)	-6.94	-14.9		-1.12
CE (kilos)	-4.31	-15.1	-0.89	

Table 3 Average trade-off values between the four objectives

3.4 Relations between decision variables and objectives

3.4.1 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 4), 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.

	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 4 Correlation between values per type of decision variable and objective values

	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

Table 5 Correlation between values per type of decision variable and objective values

In Table 5 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.

3.4.2 Percentage of active decision variables in Pareto set

In Fig. 6 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) 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.

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%, so 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.
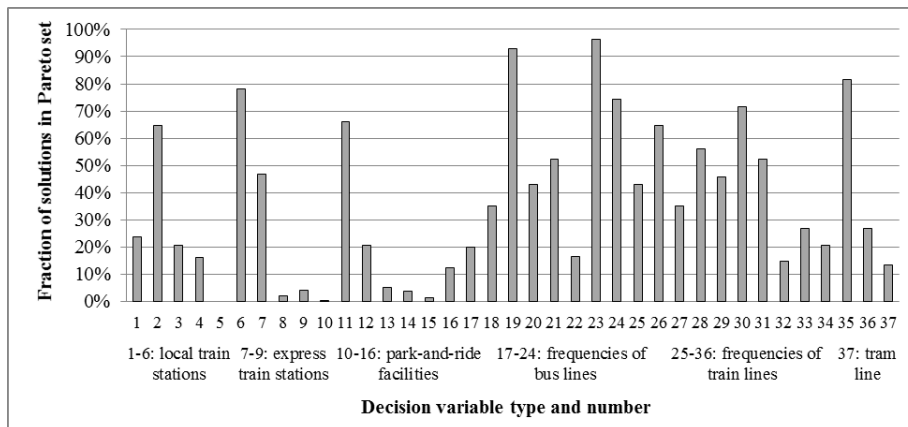

Fig. 6 Effectiveness of decision variables to reach Pareto-optimality: fraction of solutions in the Pareto set that have a nonzero value for each decision variable

3.5 Step-by-step pruning

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. 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 who 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.

3.5.1 Using values of objective values

Starting from all solutions in the upper left corner of Fig. 7, 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 (lower right corner of Fig. 7).
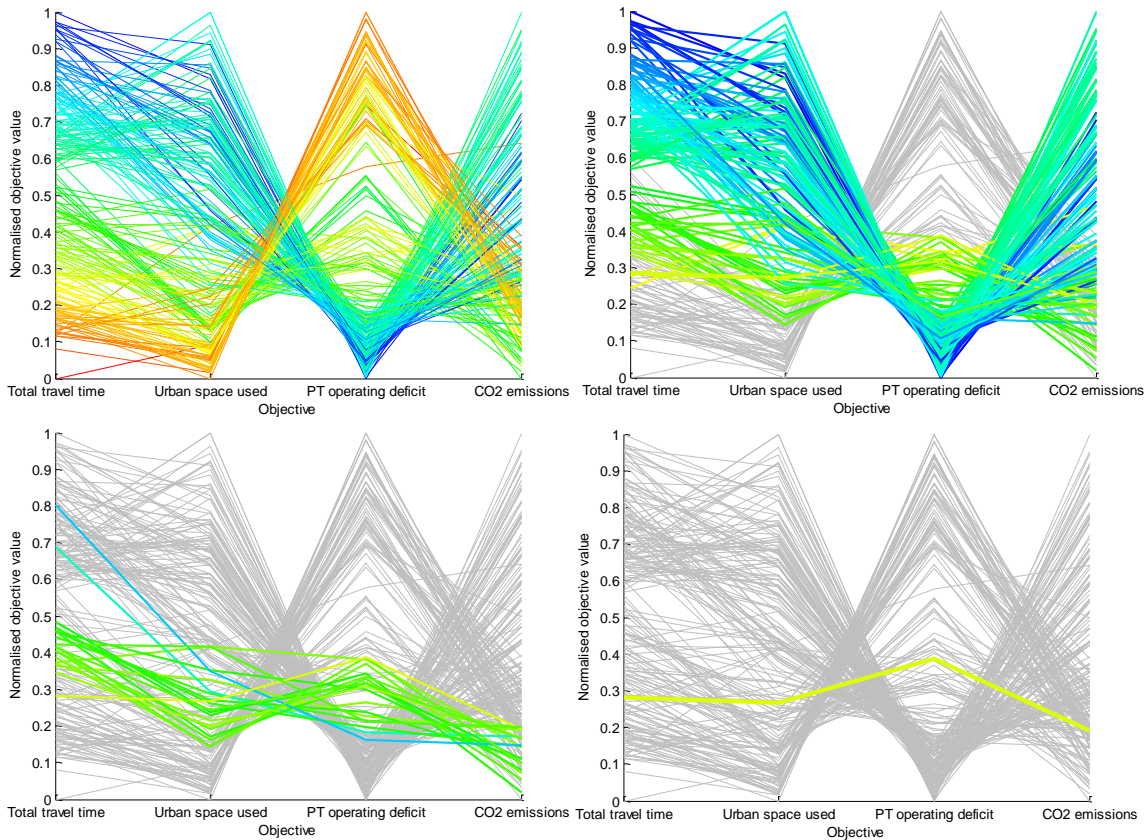

Fig. 7 Step-by-step reduction from all Pareto solutions to one solution to be implemented by setting bounds to objective values

3.5.2 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 Fig. 8. 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 Fig. 8).
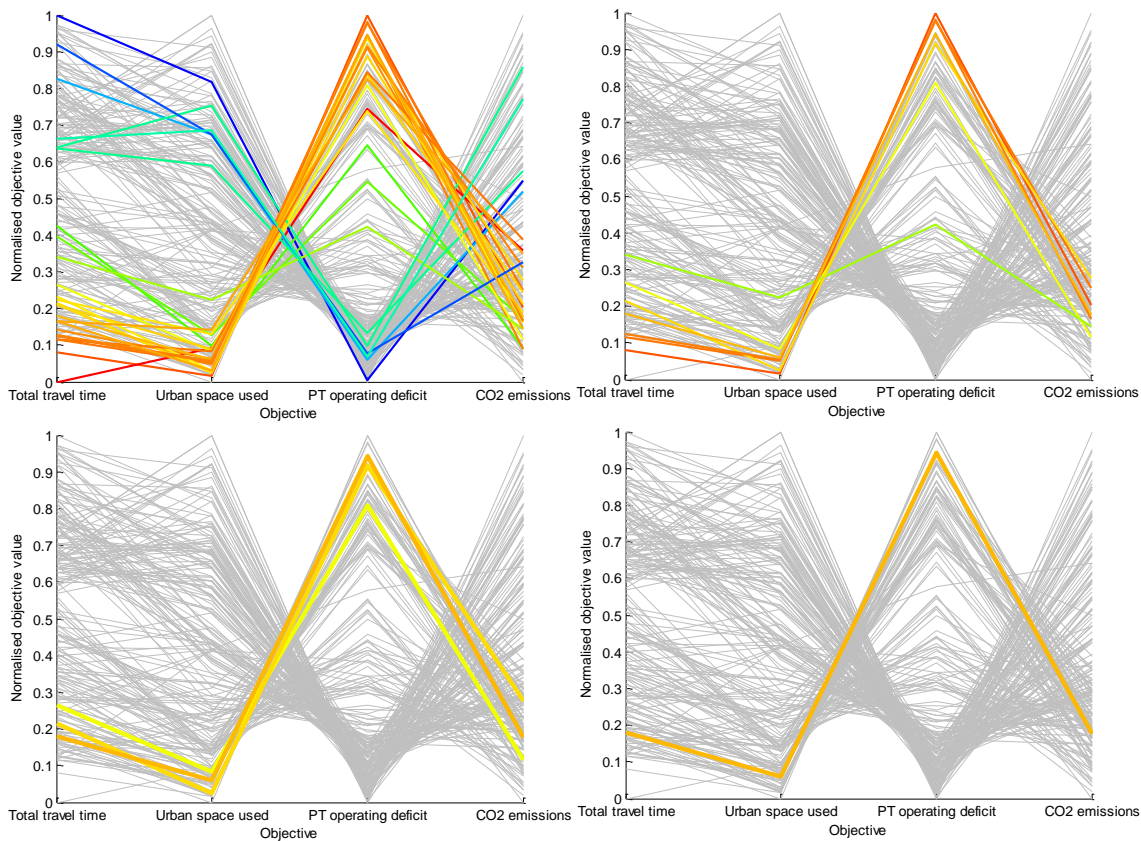

Fig. 8 Step-by-step reduction to one solution to be implemented by selecting values for decision variables

The next step is to select the solutions that include the train station in the village of Halfweg (between Amsterdam and Haarlem), which results in 3 remaining solutions with 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.

3.6 Choosing one solution as a compromise

In the end, one solution has to be chosen from the Pareto set for implementation. 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. 2, where $\overline{Z}_w$ represents the normalised value for objective w and W_c is the compromise subset of objectives, that 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 2.2)

$$BCS_{W_C}(P) = \underset{\underline{y} \in P}{\operatorname{argmin}}(\max_{w \in W_C} \overline{Z}_w(\underline{y})) \qquad (2)$$

In Fig. 9 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.
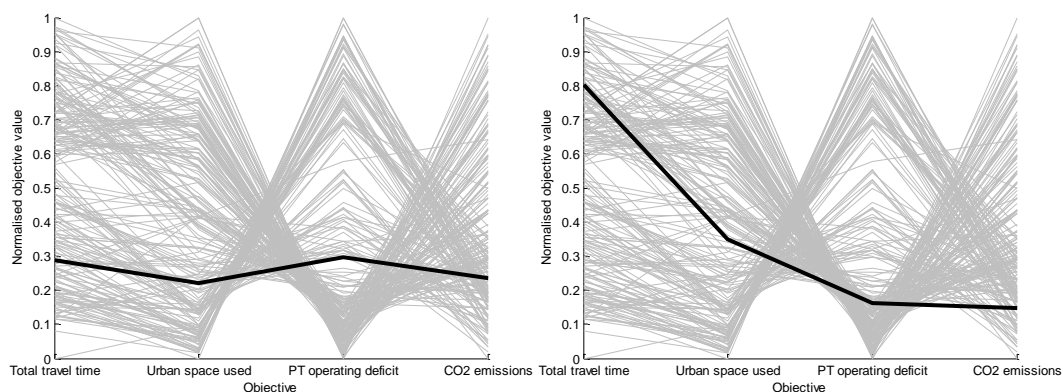

Fig. 9 Min-max solution for all four objectives (left) and for OpD and CE (right)

4 Conclusions

In this paper several methods were applied to make the Pareto set resulting from a case study 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.

4.1 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 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 second 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, 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 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.

All developed methods in principle are ready to be applied to any practical choice situation. The methods presented in this paper make the Pareto set more useful as decision support information: they demonstrate the next step in multi-objective option prioritisation. 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 make it possible to interactively change the settings of the methods, providing a direct feedback loop between the modeller and the decision maker, for example in a workshop setting.

4.2 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.

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.

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. 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. Increasing frequencies appears to be more effective to improve sustainability than introducing P&R facilities and train stations.

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. For future research it is recommended 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.

References

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), 1061-1071

Brands T., E. de Romph, T. Veitch, J. Cook (2014a) Modelling Public Transport Route Choice, with Multiple Access and Egress Modes, Transportation Research Procedia 1(1), 12-23.

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

Coello C., S. de Computación, C. Zacatenco (2006) Twenty years of evolutionary multi-objective optimization: A historical view of the field, IEEE Computational Intelligence Magazine 1(1), 28-36.

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), 182-197.

Deb K. (2003) Unveiling innovative design principles by means of multiple conflicting objectives, Engineering Optimization 35 (5), 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, 42-56

dell'Olio L., J. 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, 23-31.

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), 281-302.

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

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, 55-71

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), 792-807.

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, 241-255

Wismans L. J. J. (2012) Towards Sustainable Dynamic Traffic Management, University of Twente, TRAIL Thesis Series, T2012/4.

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), 65-77