Accelerating the Search for Optimal Dynamic Traffic Management

improving the Pareto optimal set of Dynamic Traffic Management measures that minimise externalities using function approximations

Niels Fikse

Graduation Date:
January 12, 2011

Graduation committee:
Prof. E.C. van Berkum
Dr. T. Thomas
Dr. M.C.J. Bliemer (Goudappel Coffeng BV)
Ir. L.J.J. Wismans (Goudappel Coffeng BV)

Organisation:
Goudappel Coffeng BV

In the past decades traffic demand has been increasing nearly continuously, which has provided governments all over the world with significant challenges. In the Netherlands constructing new roads is, due to various reasons, not longer considered to be the solution, the focus is now more on efficient use of existing infrastructure.

One of the instruments that is frequently used to increase the efficiency of infrastructure is Dynamic Traffic Management (DTM). In DTM we use different measures such as directing traffic through traffic lights, adding or removing lanes and variable speed limits to provide road users with the `best possible' infrastructure. It is however difficult to determine what is `best', especially now environmental and safety issues are becoming more and more important. The best possible set of measures from a travel time perspective, may very well result in very high CO$_2$ emissions, annoyance due to excessive noise and many fatalities.

It is therefore that research is being done on determining a set of possible DTM applications that can be considered the best solutions. Here `best' means that these solutions are not outperformed by any other solution on all objectives. Unfortunately finding all solutions in this set is impossible, it would easily take millennia to find them. Science has therefore resorted to finding only a part of this set (but a representative one) using heuristics such as Genetic Algorithms. However finding a part of this set using this method still takes months, which is unacceptable in the traffic and transport consultancy business. It is here where our research takes off.

Main goal of our research is therefore to accelerate the search for this set of best solutions (also known as Pareto optimal set). In our research we focus solely on accelerations that can be obtained by using approximation techniques, which is why our research goal is defined as *`accelerating the search for the Pareto optimal set found by multiobjective genetic algorithms for multiobjective network design problems, in which externalities are the objectives and DTM measures the decision variables, using function approximations'*.

It is therefore that we performed a literature study into approximation techniques, from which we derived three main techniques: the Response Surface Method (RSM), the Radial Basis Function (RBF) and Kriging/DACE. Because all of the approximation techniques have parameters that can be set, we were able to develop 148 different variants. In order to be able to determine which variant would provide the best results, we chose two simple road networks which could be used for testing and selected a set of quality measures from literature.

We found that variants that score very good on one quality measure, do not necessarily perform well on another. Furthermore we found that selecting the right parameters can significantly influence the results of the approximation techniques. However eventually we can conclude that the Kriging/DACE approach without optimising the power in the cost function is always amongst the best performing approaches. Benefit of the Kriging/DACE approach is that it does not only provide estimated objective values, but also the corresponding estimated errors. Another solution which performs reasonably well, and best on one quality measure, is the RSM approach with only cubic squared interaction terms. Main benefits of the latter approach are that it is easy to understand (it is the basis of the Least Squares Method) and that the approach is extremely fast (it can determine objective values in less than a second). It is therefore that we selected these two approaches as possible approximation methods for the remainder of the research.

We also performed a literature study into how Genetic Algorithms (and NSGA-II in particular) can be accelerated. It became clear that many of the approaches are quite complicated and/or require further optimisation, which would lead to high computational effort. We therefore selected two approaches which could easily be integrated into the original NSGA-II algorithm. The first is the Inexact Pre Evaluation (IPE) which is a deterministic approach and evaluates only those solutions which are, based on the approximated objective values, part of the Pareto optimal set. The second is the Probability of Improvement (PoI) approach, which is stochastic and determines for each solution the probability that it improves the Pareto optimal set. Next it only evaluates the *n* best solutions or the solutions with a probability higher than *x*%.

We combined the two approximation methods (RSM and DACE) and the two acceleration approaches (IPE and PoI) into three different Approximation Method Assisted NSGA-II (AMAN) algorithms. The fourth combination was impossible since PoI requires the expected error for each objective value and RSM is not able to provide this information. In order to determine which of the three approaches is best, we performed a literature study to find performance measures which can be used to compare Pareto fronts, and applied the approaches to the two test networks mentioned earlier. Unfortunately we only had time for a single run, which makes that the results are not indisputable.

We found that the results between the different AMANs (when compared with the original NSGA-II algorithm) do not point towards a single `best' approach. In fact, an approach that scores well on one performance measure can easily score quite bad on another. However based on the combined results over the two test networks, we find that PoI-DACE provides the most promising results. Not only did it provide results that were comparable to the results of the original NSGA-II algorithm, it also provided those results in only 50\% of the time that was needed by the NSGA-II algorithm. It is therefore that we selected this approach to be used in the last phase of this research.

In the last phase we tested the PoI-DACE algorithm on the (more realistic) case of Almelo. In this network we had seven controlled traffic lights and two sections of motorway with variable speed limits. In order to determine the performance of the PoI-DACE approach (in comparison with the original NSGA-II algorithm) we used the performance measures which were also used for comparing the AMANs on the test networks. Due to the fact that performing a run for both the NSGA-II and the AMAN algorithm takes about three weeks, we were, again, only able to perform a single run.

The results of the analysis were quite promising. The area that was dominated by the NSGA-II, but not by the AMAN was only 3\% of the total area dominated by the NSGA-II algorithm. Furthermore we found that the spread of solutions over the Pareto front was better and that a reduction of 30\% in calculation time is realisable. Unfortunately we also found that the influence of stochasticity (there are a lot of random processes involved in NSGA-II), is significant. In order to reduce the uncertainty in these conclusions, we would have to perform dozens, if not hundreds, of runs.

We furthermore tried to interpret the Pareto optimal set that was found from a traffic and transport engineering perspective, which appeared to be a difficult task. Using grouped data and a multitude of boxplots we could, for some of the DTM measures, determine a relation between the settings and the resulting objective values. Unfortunately we were not able to find correlation effects between different DTM measures, something that might be caused by a lack of data.

Based on the results on the different test networks and the Almelo case we find that it is highly likely that the proposed AMAN (and probably also the other AMANs) can achieve a Pareto front that is comparable to the one found by NSGA-II. Besides PoI-DACE is able to do so with a reduction in calculation time of 30\%. We therefore can state that we can indeed accelerate the search for the Pareto optimal set by applying approximation techniques.

It does however seem wise to do some further research. Especially the performance of AMANs can be disputed, since only a single run has been performed. In order to provide reliable results at least dozens of runs should be performed before we can conclude, statistically, that a specific AMAN is equal to the original NSGA-II algorithm.

We also recommend that the behaviour of the PoI approach, or more specifically the change of approximated values and errors over time, is studied. We were unable to apply a `better than *x*% policy' because it appeared that after a few iterations all solutions were accepted.

Finally we suggest that more time and effort is spend in analysing the resulting Pareto front. Unfortunately we were unable to detect important relationships between DTM measures, however that might be possible if sufficient data and time is available.