Exploiting Prior Information in Multi-Objective Route Planning

Antony Waldock[1] and David Corne[2]

[1] Advanced Technology Centre, BAE Systems,
Bristol, UK
antony.waldock@baesystems.com
[2] School of MACS, Heriot-Watt University,
Edinburgh, UK
dwcorne@macs.hw.ac.uk

Abstract. Planning a satisfactory route for an autonomous vehicle over a complex unstructured environment with respect to multiple objectives is a time consuming task. However, there are abundant opportunities to speed up the process by exploiting prior information, either from approximations of the current problem instance or from previously solved instances of similar problems. We examine these two approaches for a set of test instances in which the competing objectives are the time taken and likelihood of detection, using real-world data sources (Digital Terrain Elevation Data and Hyperspectral data) to estimate the objectives. Five different instances of the problem are used, and initially we compare three multi-objective optimisation evolutionary algorithms (MOEA) on these instances, without involving prior information. Using the best-performing MOEA, we then evaluate two approaches that exploit prior information; a graph-based approximation method that pre-computes a collection of generic 'coarse-grained' routes between randomly selected zones in the terrain, and a memory-based approach that uses the solutions to previous instances. In both cases the prior information is queried to find previously solved instances (or pseudo-instances, in the graph based approach) that are similar to the instance in hand, and these are then used to seed the optimisation. We find that the memory based approach is most effective, however this is only usable when prior instances are available.

1 Introduction

Route-planning is one of the increasingly many application domains in which a multi-objective optimisation (MOO) approach [1] has been found to have significant advantages over single-objective approaches [2]. In this paper, we further explore multi-objective optimisation algorithms for route-planning of manned and unmanned vehicles in a hostile and unstructured environment, and focus on the question of accelerating the process by exploiting prior information. Speed of optimization can be particularly vital in the application environments of interest to us – broadly speaking, this is because there will often be a need to have a viable route plan within seconds of the decision being made to start to move the vehicle to a given new location. Meanwhile, in the route planning domain, as well as a large number of other interesting application domains, a range of prior information is available and could be used in various ways to bootstrap the optimisation process.

In the case of route planning, the broad geographic area and terrain characteristics where routes are to be planned are known in advance. The solutions to previously solved route planning instances within the same terrain may also be known but it would be infeasible to pre-compute all instances from every possible start and end location (some geographic features in the environment may move between instances). However, it is very appealing to utilise prior information whenever possible. In this paper, we examine two approaches to integrating prior information into a multi-objective optimisation algorithm; (i) solving approximations of the current problem instance, or (ii) information derived from previously solved problems with sufficient similarity.

In the remainder, we briefly cover background material in Section 2 and then introduce multi-objective route-planning. Section 3 then evaluates three MOEAs on our five test instances, without exploiting prior information. In Sections 4 and 5 we then respectively explore two different approaches to include prior information. Section 6 concludes and discusses future work.

2 Background

A multi-objective optimization (MOO) problem is posed as $\arg\min_{\boldsymbol{x}\in X} G_m(\boldsymbol{x})$, where $G_m(\boldsymbol{x})$ is a set of m objective functions and $\boldsymbol{x}$ is defined as a vector of decision variables (or a solution) in the form $\boldsymbol{x} = (x_1, x_2.., x_N)$ from the set of solutions X. The aim is to find the *Pareto* set which contains all the solutions that are not dominated by any other solution. A solution $\boldsymbol{x}$ is said to be non-dominated by $\boldsymbol{y}$, if and only if, $\boldsymbol{x}$ is as good as $\boldsymbol{y}$ in all objectives and $\boldsymbol{x}$ is strictly better then $\boldsymbol{y}$ in at least one objective.

The most effective MOO approaches to date are generally regarded to be multi-objective evolutionary algorithms (MOEAs). Typically, MOEAs (as well as most optimization algorithms) make little or no use of prior information that may be available about the problem at hand. The concept of exploiting prior information is pervasive in artificial intelligence, appearing in several different guises (e.g. case based reasoning (CBR) [3], or, more recently, per-instance tuning [4]), however it is infrequent in the optimization literature, perhaps because appropriate approaches are highly domain-specific. Nevertheless some examples include work in the Genetic Programming community [5, 6] in which populations were seeded with solutions to previous instances, while an approach was recently proposed in [7] which exploits extensive pre-computation of solutions to *potential* instances that may be faced in a given domain. Meanwhile, [8] explores the re-use of the probability models built by an estimation of distribution algorithm (EDA) on previous instances, while seeding with previous solutions is occasionally explored, especially for dynamic optimization [9, 10].

We examine two approaches with which to integrate prior information in route planning, in the scenario that instances will occur with previously unknown start and end locations, but within a known geographic area (e.g. a 5km by 5km square). The first approach is to prepare in advance Pareto optimal but coarse grained route solutions for a large collection of potential start and end locations within the region. For any such start/end pair, we abstract the search space as a directed graph, and then use a multi-objective extension to traditional A* called NAMOA* [11] to find Pareto optimal coarse-grained routes. which, in turn, seed the population of a MOEA solving the instance at hand. The second approach uses solutions to previously solved similar route planning instances to seed the initial population for the new instance.

We consider a route planning scenario where a route is required that minimises a set of competing objectives such as the fuel used, the time taken, the distance travelled, or the likelihood of being detected by observers. We build on the route planning problem defined in [2] and are informed by previous studies of motion planning for autonomous vehicles [12, 13]. Route planning is the construction of a route that navigates between two geographic locations. The start and end location are defined by a latitude, longitude and heading from true north. For convenience, we encode a route in relative polar coordinates where α_i is the heading relative to the next way point and r is the distance to travel in this direction. To evaluate a route, the objective functions used here are the time taken and likelihood of detection, as defined in [2], and the route is divided into 30 segments. The objective functions are calculated using Digital Terrain Elevation Data (DTED) [3] and the NASA LandSat Multispectral data. The Multispectral data is combined with a classifier to infer the terrain type and hence the maximum speed allowed on that portion of the route segment. To evaluate the performance of different MOEAs, five instances, P_1 to P_5, were generated. The definition of the routes and java code for the objective functions is available at http://code.google.com/p/multi-objective-route-planning/.

3 Comparison of Multi Objective Optimisation Algorithms

First, we compare three different MOEAs, MOEA/D, SMPSO, and NSGA-II, on our five problem instances, without using prior information. Multi Objective Evolutionary Algorithms Based on Decomposition (MOEA/D)[14] was selected to represent the current state of the art MOEA. Speed-constrained Multi-objective PSO (SMPSO)[15] is used to provide a baseline algorithm from the Particle Swarm Optimisation (PSO) community, while the Non-dominated Sorting Genetic Algorithm II (NSGAII)[16] is also tested as a commonly used effective benchmark MOEA.

The implementations of MOEA/D, SMPSO and NSGAII have been taken from JMetal[4] and the results presented are averaged over 50 independent runs

[3] EarthExplorer (http://earthexplorer.usgs.gov)
[4] http://jmetal.sourceforge.net/

which were limited to 200,000 evaluations per run. NSGAII and SMPSO had a population size of 100 and MOEA/D a population size of 600. Comparisons are quantified using the Inverted Generational Distance (IGD), as defined in [17]. The 'known' Pareto optimal front is calculated by combining the solutions generated from all experiments presented in addition to one million randomly generated samples. Results are summarised in Table 1.

Prob	P_1			P_2			P_3			P_4			P_5		
Alg	S	N	M	S	N	M	S	N	M	S	N	M	S	N	M
Mean	61.0	464.7	33.8	152.4	817.9	47.9	130.8	345.5	29.8	107.2	407.5	47.7	54.1	205.7	5.8
Std	12.0	150.9	28.2	41.6	193.4	36.1	25.5	141.2	19.2	33.6	76.3	25.5	9.8	56.9	14.8

Table 1. The IGD values (M. = mean and Std = standard deviation) for (M)=MOEA/D, (S)=SMPSO and (N)=NSGAII on Problems 1 to 5

The results in Table 1 follow preliminary experiments which hand-optimised the parameters of each of the algorithms. For these five problem instances, MOEA/D clearly produces solutions closest, in terms of IGD, to the reference Pareto optimal front. Hence, MOEA/D is used in the remainder of this paper.

4 Graph-Based Approximation

In the route planning scenario of interest, before an instance of the problem arises, we know the broad geographical region in which the start and end locations will be. We describe here a way to exploit that prior information, based on *a priori* finding coarse-grained solutions to many potential instances, based on all possible pairs of start and end locations over a 50m by 50m mesh. For each such pair, a solution tree is generated by using the encoding outlined in Section 2, but with each bearing (α_i) restricted to a discrete set, e.g. $[-30, 0, +30]$. Once the maximum number of segments has been reached, the current location is joined to the end location using a single straight line segment. The Pareto set of solutions on this tree is then extracted by using NAMOA* (of which more below). Given a new instance of the route planning problem, we then find the pre-solved coarse-grained instance whose start and end locations best match the new instance, and use the Pareto set found by NAMOA* to seed the MOEA/D population.

Figure 1 (a) shows the final IGD value for the five different route planning problems when MOEA/D is initialised randomly and with the non-dominated solutions found by solving two configurations of the graph-based approximation. The results clearly shown that initialising the MOEA/D population with solutions generated from an initial graph-based approximation has an improvement in the final IGD value. With 3 segments and 3 bearings (3,3) at each node, the final IGD is statistically different (according to a two tailed, paired T test with confidence level 0.99) for 1 of the problems (P_4) and with 5 segments and 5 bearings (5,5), the final IGD value is statistically different for all 5 problems.

Figure 1 (b) shows the number of route evaluations required to reach the 110% of the maximum final IGD value when using MOEA/D with a randomly

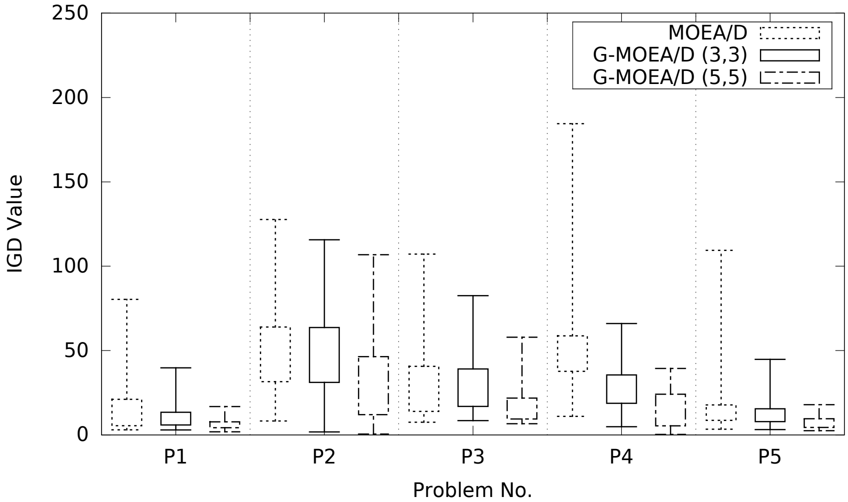

(a) The final IGD value

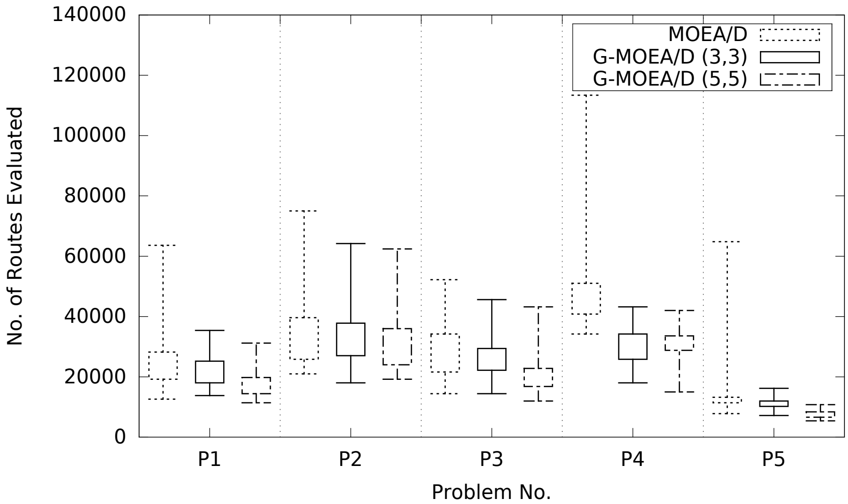

(b) No. of Routes Evaluated

Fig. 1. The IGD value on the last optimisation iteration and the number of evaluations to reach 110% of the maximum IGD value (as defined by MOEA/D) for MOEA/D and Graph-Based MOEA/D with (3,3) and (5,5) segments and bearings for Problems 1 to 5

initialised population. The results show that the number of evaluations required is statistically reduced on all route planning problems except for problem 2 (with 5 segments and 5 bearings). The results show that, on average, using a graph-based approximation can reduce the number of routes evaluated by 5,318 for 3 segments and 3 bearings and 8,582 routes for 5 segments and 5 bearings where NAMOA* only evaluates, on average, 11 and 143 routes for each of these configurations.

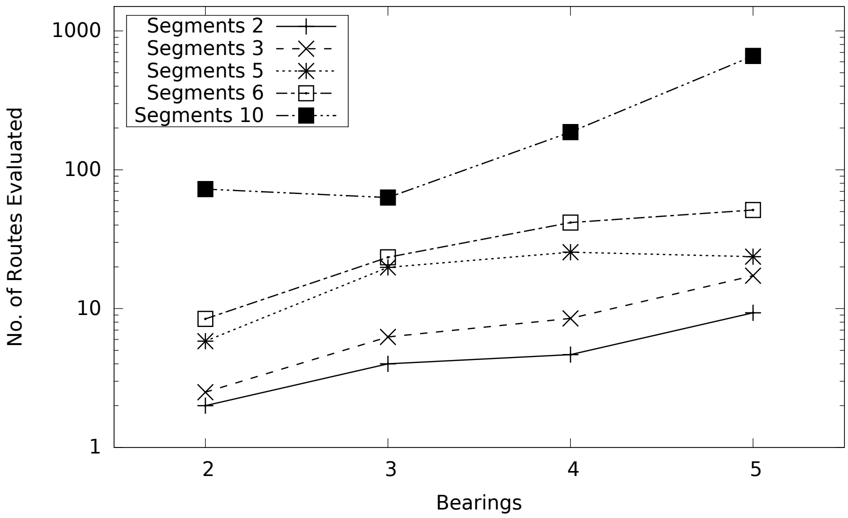

Fig. 2. The No. of Routes Evaluated for P_1 with different graph configurations

Although, the exponential complexity of NAMOA* is clearly shown in Figure 2 (Note the log scale) where the number of routes evaluated significantly increases (for a small reduction in the IGD). The number of routes evaluated are calculated by totalling the number of edges evaluated during the operation of the NAMOA* algorithm. One reason for this is that a suitable tree-pruning heuristic is unavailable for the likelihood of detection and hence a lower bound must be assumed (likelihood of detection is zero) therefore reducing the removal of dominated branches (or partially explored routes) in the graph.

This section has shown that seeding MOEA/D with solutions generated from an approximation of the problem, solved using NAMOA*, has a significant impact on the IGD value but NAMOA* can only be run for very crude approximations of the problem before the number of evaluations required quickly becomes infeasible. The next section evaluates a different seeding strategy that is based on storing solutions between problem instances.

5 Memory-Based

In this section, a memory-based approach is examined, where previous solutions are reused from previously solved instances of similar problems. The approach is only appropriate when the problems are not evolving rapidly and when information from previously solved problems have some bearing on the current problem being solved i.e. the shape and distribution of the routes. The non-dominated solutions generated by MOEA/D are stored in a k-d tree [18] which is referenced by the latitude and longitude of the start and end locations. All of the solutions from previous problems are currently stored in the k-d tree but only a subset of the K closest neighbours, up to half the total population, are used to initialise the population of MOEA/D when solving a new instance.

To evaluate the approach, it is necessary to generate a sequence of problem instances. A sequence of ten instances of route planning problems, $P_{i,j}$, where i is the problem number and j is the index in the sequence, are generated by randomly selecting a start and end location in the area of the route planning problems P_1 to P_5. Once a route planning problem in the sequence has been solved, the non-dominated solutions are added to the k-d tree and the next route planning problem is tackled. At the beginning of the next route planning problem, the closest neighbours to this route planning problem are extracted from the k-d tree and the non-dominated solutions for these problems added to the initial population of MOEA/D. The remaining population used in MOEA/D is randomly initialised within the input parameter range.

Figure 3 (a) shows the final IGD value for MOEA/D and the Memory-Based MOEA/D (neighbours=3) on P_1 over the sequence of ten route planning instances. The IGD value for each iteration is generated over 50 runs of MOEA/D and the same sequence is used for each run. The IGD results show that initialising MOEA/D with solutions from previously similar instances results in a statistically better (using a two tailed, paired with 0.99 probability) set of solutions for five of the ten instances (4,6, 8,9 and 10). Figure 3 (b) shows the number of evaluations required to reach 110% of the maximum final IGD value as found by MOEA/D. The results show that for some instances the number of evaluations is significantly lower (4,6,8,9 and 10) with on average a reduction in the number of routes evaluated for these four instances is 13,369. The results show that MOEA/D initialised using solutions from previously solved problems has the potential, on some instances, to reduce the number of evaluations required to produce a reasonable approximation of the Pareto optimal front.

For example, the reduction in the number of evaluations can be graphically seen in Figure 4 (a) where the memory-based MOEA/D (shown in dashed) is compared with a randomly initialised MOEA/D (shown in black). The reason for this earlier reduction in IGD value can be seen by comparing the solutions extracted from the memory at the start of the optimisation. The solutions extracted from the memory provide a reasonable approximation of the Pareto optimal front in the first few iterations of the algorithm. A comparison of the initial and optimised routes can be seen in Figure 4 (b) where the initial routes provide a broad spectrum of possible routes with previously successful shapes without having to exhaustively search all possible combinations as with the previous graph-based approach.

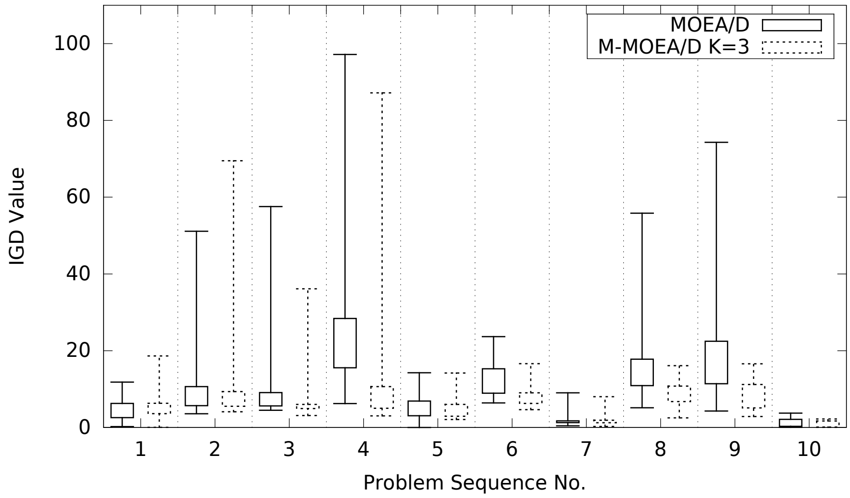

(a) IGD Value

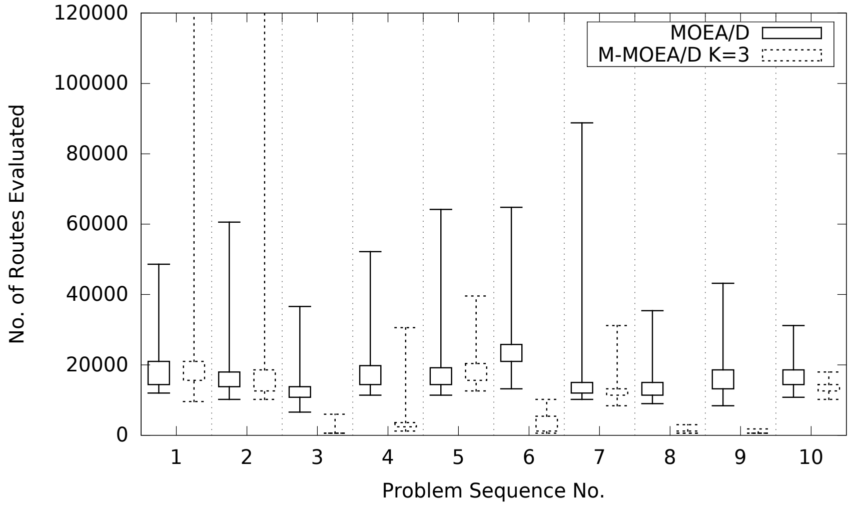

(b) Evaluations

Fig. 3. The IGD Value on the last iteration and the number of evaluations to reach 110% of the maximum final IGD value (as defined by MOEA/D) for MOEA/D and Memory-Based MOEA/D for Problem Sequence 1

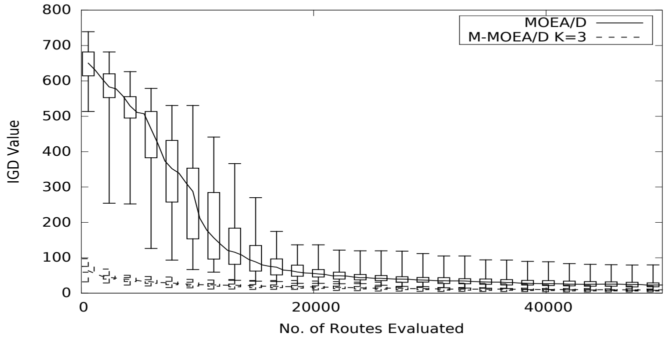

(a) Comparison of IGD Value during the optimisation

(b) Geo-spatial representation of the routes for Problem $P_{1,9}$

Fig. 4. Analysis of Problem $P_{1,9}$ - ©Google and GeoEye

6 Conclusions and Future Work

This paper has presented an examination of two approaches to using prior information, graph-based and memory-based, for MOEA/D when applied to a route planning problem over an unstructured environment. The experimental results have shown that both approaches enable MOEA/D to generate a set of solutions closer to the known Pareto optimal front in fewer iterations than a traditional random initialisation. Solving a graph-based approximation has been shown to produce routes closer to the known Pareto optimal front for the five route planning problems but is only feasible for small graphs because the number of evaluations required increases exponentially. Using a memory-based approach has also been shown to generate routes closer to the known Pareto optimal front for a sequence of problem instances but this is highly dependent on whether a sufficiently similar problem has been solved previously and whether the environment has evolved since solving that problem. The approaches examined in this paper are applicable to a wide range of multi-objective optimisation applications.

Future work will concentrate on how to apply these techniques to dynamic environments where the problem is evolving either during or between route planning problems and to analyse methods of storing and extracting subsets of solutions based on the similarity of the instances.

7 Acknowledgements

This project was funded by the Systems Eng. for Autonomous Systems (SEAS) Defence Tech. Centre (DTC).

References

1. Deb, K.: Multi-Objective Optimization Using EAs. John Wiley and sons Inc (2002)
2. Waldock, A., Corne, D.: Multi-objective optimisation applied to route planning. In: Genetic and Evolutionary Computation Conference (GECCO) 2011. Volume 1. (2011) 1827–1834
3. Leake, D., ed.: Case-based Reasoning: Experiences, Lessons and Future Directions. AAAI Press, Menlo Park, CA (1996)
4. Hutter, F., Hoos, H., Stutzle, T.: Automatic algorithm configuration based on local search. In: Proc. AAAI, MIT Press (2007) 1152–1157
5. Arcuri, A., White, D.R., Clark, J., Yao, X.: Multi-objective improvement of software using co-evolution and smart seeding. In: Proceedings of the 7th International Conference on Simulated Evolution and Learning. SEAL '08, Berlin, Heidelberg, Springer-Verlag (2008) 61–70
6. Westerberg, C.H., Levine, J. In: Investigations of Different Seeding Strategies in a Genetic Planner. Volume 2037. Springer-Verlag (2001) 505–514
7. Corne, D., Reynolds, A.: Optimization by precomputation. In: Proc. NABIC 2011, IEEE CIS Press (2011)
8. Hauschild, M., Pelikan, M.: Intelligent bias of network structures in the hierarchical boa. In: Proc. GECCO 2009, ACM Press (2009)
9. Branke, J.: Memory enhanced evolutionary algorithms for changing optimization problems. In: Proc. CEC 2009. (1999) 1875–1882
10. Potter, M., Wiegand, R., Blumenthal, H., Sofge, D.: Effects of experience bias when seeding with prior results. In: Proc. CEC 2005. Volume 3. (2005) 2730 – 2737 Vol. 3
11. Mandow, L., de-la Cruz, J.L.P.: Multiobjective a* search with consistent heuristics. J. ACM **57** (2010)
12. Pivtoraiko, M., Kelly, A.: Efficient constrained path planning via search in state lattices. In: The 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space. (2005)
13. Lavalle, S.M.: Rapidly-exploring random trees: A new tool for path planning. Technical report (1998)
14. Li, H., Zhang, Q.: Multiobjective optimization problems with complicated pareto sets, moea/d and nsgaii. IEEE Trans. Evolutionary Computation (2008)
15. Nebro, A., Durillo, J., García-Nieto, J., Coello Coello, C., Luna, F., Alba, E.: Smpso: A new pso-based metaheuristic for multi-objective optimization. In: 2009 IEEE Symp. on MCDM), IEEE (2009) 66–73
16. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A fast elitist multi-objective genetic algorithm: Nsga-ii. IEEE Trans on Evol. Comp. **6** (2000) 182–197
17. Zhang, Q., Liu, W., Li, H.: The performance of a new version of moea/d on cec09 unconstrained mop test instances. In: Proc CEC 2009. (2009) 203–208
18. Bentley, J.L.: Multidimensional binary search trees used for associative searching. Commun. ACM **18** (1975) 509–517