

MOPF: A Multi-Objective Planning Framework for Generation of Pareto-Optimal Plan Frontiers using Satisficing Planners

Michal Sroka

Michal.sroka@kcl.ac.uk Department of Informatics, King's College London, UK.

1 Introduction

Current planners focus on finding a single solution to a given problem. The decision maker (DM) has a choice of either finding the solution fast using a satisficing planner or finding a high quality solution using an optimal planner. The quality of the solution is evaluated in terms of a metric function given by the DM. That forces the DM to specify preferences in terms of a single metric function which in many real life problems is hard.

This work is based on the need to relax this constraint and to allow the planning process to be carried out without the need to specify the single metric function. A single plan is sufficient when minimising a single objective function, however, it is very unlikely that a single solution would minimize multiple objective functions at the same time. The solution for a problem with multiple metric functions is no longer a single plan but a set of plans which represents the trade-off between different objectives.

Satisficing planners are doing remarkably well in finding solutions to planning problems efficiently. Although they are not optimal, the speed with which they find a solution makes them very valuable. We explore how we can use them to quickly find solutions in various areas of the search space, which combined together give a high quality solutions set.

In multi objective optimisation we can divide the interaction with the DM into three main categories [Horn, 1997]. First, when DM can articulate its preferences before planning, we can construct an aggregated function of all preferences knowing weights between them. The second category is an interactive articulation of preferences by the DM, this approach is least common and appears when the DM helps to guide the search, keeping in mind that his preferences might change with time. The last category is *a posteriori* articulation of preferences, where the task for the system is to present a good choice of alternative solutions. The last category has received a lot of attention in OR, however, in planning this area is not well explored and this research aims to fill this gap.

In this paper we present a Multi-Objective Planning Framework (MOPF) which is capable of generating sets of qualitatively good plans. A method to evaluate pareto-frontier (PF) of plans is introduced and used to evaluate the planning system. We start by giving a background to this research. Following that we introduce MOPF which works on problems with multiple metrics. In the final sections we present evaluation

of this system, solution it generates and discussion of some multi-objective domains.

2 Background

Planning problems are described using sets of propositions and numeric fluents, most commonly using PDDL [McDermott et al., 1997], a language dedicated to this purpose. In this work we use the PDDL2.1 version [Fox and Long, 2003]. The planning problem describes a start state and the goal condition in a given domain. The planner then outputs a sequence of actions which, if applied, takes us from the start state to a state satisfying the goal. The planning problem also contains objective functions which are the main focus of this research. The objective functions are different from goal states as they do not need to be satisfied in order for the problem to be solved. The plan can be awarded extra value if it satisfies conditions given by objectives or if it minimises cost functions described in objectives.

Metric functions were introduced into PDDL [Fox and Long, 2003] from version 2.1 onwards. An example of a metric specification is

```
(:metric minimize (+ (fuel-used car) (fuel-used truck)))
```

Our goal is to generate a set of plans, where the plans are all good quality, in terms of the objective functions, and significantly different from each other. By different plans, we mean plans which differ in terms of a metric as defined in Definition 1.

Definition 1. *Distance between plans* $|\rho_1 - \rho_2|$ is an euclidean distance in the space described by the metrics.

$$|\rho_1 - \rho_2| = \sqrt{\sum_{i=0}^n (\Theta_i(\rho_2) - \Theta_i(\rho_1))^2}$$

$|\rho_1 - \rho_2|$ is the distance between plans ρ_1 and ρ_2 . Where $\Theta_i(\rho_0)$ is the value of the i^{th} objective function for the plan ρ_0 . Both plans are evaluated using the same set of objective functions.

A different approach to measuring the difference between plans is presented by Nguyen et al [2011], where the distance between plans is defined in terms of actions, causal links or states visited in plan execution. However, if we use the plan distance based on these metrics, the issue of eliminating plans which appear as the same for the DM but are different for the planner is difficult to address. What is more, this work assesses the quality of the solution set based on the distance

between plans and not quality of each of the plans which, even though the planner tries to find good quality solutions, produces sets containing low quality plans.

We would like to find plans such that there are no other plans which dominate (Definition 2) plans from our set.

Definition 2. Plan Domination. Plan ρ_0 dominates plan ρ_1 , represented as $\rho_0 \succ \rho_1$, if there exists an objective function, Θ_i , such that $\Theta_i(\rho_0) < \Theta_i(\rho_1)$ and for all other objective functions $\Theta_j(\rho_0) \leq \Theta_j(\rho_1)$.

We would like to generate the set of non-dominated plans using an existing planner. It seems clear that the planner must be able to generate different plans for the same problem with different objective functions. This requires the planner to be metric sensitive [Sroka and Long, 2012], which means that it generates different plans in response to changes in the plan metric. A method for determining whether a planner is metric sensitive is described by Sroka and Long [2012].

3 Pareto-frontiers

Definition 3. Pareto-frontier (PF), Φ , in a solution space Ω is defined as a set of non-dominated plans.

$$\Phi = \{x \subseteq \Omega : \forall y \subseteq \Omega : y \not\prec x\}$$

All plans belonging to a PF are called pareto-optimal. PF contains all most desirable plans solving our given problem, with respect to selected metrics.

In this work, for a given set of metrics, which are used to evaluate solutions, a metric space is constructed where each of the metrics is treated as a dimension of that space. For each plan a value of each metric is calculated and the plan is placed onto the metric space based on its metric values. When we refer to coordinates of pareto points we mean values of the planning metrics for the solution it represents.

3.1 Presenting a pareto-frontier

The presentation of a PF is also a challenge. Once all of the plans are generated and the trade-offs are known, the main concern is how to communicate the alternatives to the DM in a clear way, allowing them to see the trade-offs and make appropriate decisions. There has been a good progress in User Interface representation of solutions presented in the work of Giuliano et al [2007; 2011], where authors are dealing with the multi-objective scheduling problem for observations using space telescopes. The difficulty of presenting the distribution of solutions starts as we introduce more and more dimensions. Dealing with visualisation of up to 3D spaces is not very challenging, however, as the dimensionality increases it is harder to display the results. The approaches taken include projections of the PF onto lower dimensions and presenting them as plots or histograms of objective values, or as explicit values. All this combined in a clear GUI gives the DM a good understanding of the trade-offs involved and allows an informed decision.

3.2 Pareto-frontier operations

This section presents some operations applicable to PF. One purpose of these operations is to be able to present the PF to the DM in the most appealing form, and to be able to abstract unimportant information for even clearer presentation.

The second purpose is to merge information from various experiments to obtain final values.

Merging pareto-frontiers

Merging two frontiers is used to generate one PF of a better quality. This is achieved by adding sets of solutions from both PF and then removing duplicates and dominated solutions. This can be defined as:

Definition 4. Sum of two PF Φ_1 and Φ_2

$$\Phi = \Phi_1 \oplus \Phi_2 = \{x \subseteq \{\Phi_1 \cup \Phi_2\} : \forall y \subseteq \{\Phi_1 \cup \Phi_2\} : y \not\prec x\}$$

Projecting pareto-frontier

In this work we will always project an N dimensional PF on N-1 dimensional space parallel to one of the dimensions. An interesting observation is that when projecting a PF from N to N-1 dimensional space the projection, after removing dominated points, is a valid PF in N-1 dimensional space. This is not obvious from a mathematical point of view as a projection of a function from N to N-1 dimensional space is different to its intersection with that space, and mathematically this intersection is a set of valid solutions on that space. In planning, however, it is not the case as projecting from a higher dimensional spaces is equivalent to abstracting the N-th dimension which allows us to view the solutions only in terms of N-1 other dimensions which makes it easier for the DM to comprehend. Projection to 2D spaces was used to visualise trade-offs in a solution space for a scheduling problem [Johnston and Giuliano, 2011].

Combining metrics of a pareto-frontier

When combining two metrics into one weighted sum, this is equivalent to mathematical projection of N dimensional space onto an N-1 dimensional space, described by all N-2 metrics plus one extra metric for combined metrics. For example, when combining first and second metric from the space (m1, m2, m3, m4) in the following way: new_metric = $\alpha \cdot (m1) + \beta \cdot (m2)$. The resulting frontier is equivalent to its projection onto the plane parallel to the line: $\alpha \cdot (m1) + \beta \cdot (m2) = 0$ and perpendicular to the other dimensions. Computationally we could simply convert it as $[\alpha \cdot (m1) + \beta \cdot (m2), m3, m4]$.

Remarks about pareto-frontier transitions

The above transformations of the PF show that we can obtain different PF and look at the solution space from various angles simply by performing mathematical actions on the existing PF. This can be done without the need for re-planning and is a very robust way of presenting the frontier to the DM and allowing the DM to explore various aspects of the solution space. This gives the DM a better understanding of the trade-offs available and the domain as a whole.

3.3 Evaluating a pareto-frontier

Related Work

An approach to generate sets of plans spanned across the search space, described by Nguyen et al [2009; 2011], is based on the assumption that the DM cannot explicitly specify preferences. The output of the planning process is a dDISTANTkSET set of plans. dDISTANTkSET contains k plans and each plan is distant from all others by a minimum

of the distance d . The distance is defined based on 1) Actions 2) Causal links or 3) states visited.

The work presented by Nguyen et al [2009; 2011] adapts LPG [Gerevini, Saetti, and Serina, 2004] to generate sets of plans. The change is to use Integrated Convex Preference (ICP) measure inside its heuristic instead of the standard execution cost. These sets of plans do not intend to be optimal as the focus is on generating variety of different solutions in terms of their plan distance measure. Therefore it is a different method to the one presented in this research which aims for all plans to be of a good quality.

The output set is evaluated based on ICP which is defined in Definition 5. For each plan from the set ρ^* a weighted sum of time t_{p_i} and cost c_{p_i} of that plan is calculated. $w_0 = 0$ and $w_k = 1$. For more details please refer to [Nguyen et al., 2011] section 3.2.

Definition 5.

$$ICP(\rho^*) = \sum_{i=1}^k \int_{w_{i-1}}^{w_i} h(w)(w \times t_{p_i} + (1 - w) \times c_{p_i})dw$$

The ICP score was integrated into LPG and was used to drive its search towards plans different from the ones already found in the planning process. Using ICP in that form forces the planner to find multiple different solutions.

The use of ICP to drive the planner to find different solutions gives a good range of different plans. However, this is achieved at the cost of quality. The reason why the quality of the solution found is deteriorated by increasing its qualitative difference is that it is much easier for the planner to find qualitatively different solutions further from the pareto-optimal set than it is on the pareto-optimal set where sometimes solutions are close to each other.

Another possible approach to evaluating pareto frontiers is to calculate an average distance from each of the pareto-optimal solutions to a reference point, which is usually set to be 0 on each axis. An extension of this approach is to divide the utility space into sectors and compare the parts of frontier lying in respective sectors of the space [Ziadloo and Ghamary, 2009]. The distribution of the frontier is also calculated with regards to the sectors. This approach, although easy to calculate, is difficult to apply to frontiers where each of the frontier occupies different sectors of the space. In this cases the method does not give us any information. Also for cases where one frontier would strongly dominate the other in areas where the other contain no points.

Hyper-volume Indicator (HVI) is a measure used to assess D-dimensional pareto frontiers in computational problems. It calculates the hyper-volume of a dominated space bounded by a nadir reference point which is dominated by all points. The calculation of the Hyper-volume is equivalent to calculating Klee’s measure which was first formed by Victor Klee in his article [Klee, 1977]. This problem was then solved by John Bentley in 1977 however his notes are unpublished. There have been many improvements to the algorithm [Fonseca, C.M. and Paquete, L. and Lopez-Ibanez, M., 2006].

Proposition of a new metric

In this research we use a similar measure to the HVI Figure 1, our measure non-dominated hyper-volume indicator (NDHVI), defined in Definition 6. An example 2D and 3D frontier and its NDHVI is represented in Figure 2. NDHVI

is the opposite to the hyper-volume indicator but can be calculated in an analogous way. This measure is a very good basis for evaluation of PF and their comparison between each other. The reason why this method is preferred over the HVI is the choice of bounds. In our experiments we can evaluate the frontier once all planning is finished and we know what is the worst possible plan. We will use a plan dominated by all plans generated as a bound for the NDHVI, as opposed to the point dominated by all pareto-optimal plans which is used in HVI. This guarantees that if we improve on our frontier, by finding a new non dominated point, the score for the new, and better, frontier is lower than the score for the original one. This property does not hold for the HVI.

The PF with smaller NDHVI is considered better. The reason why we believe that this metric gives good quality evaluation is that the only place in the space where we can find solutions which would add to the PF is the non-dominated space. This means that if a solution does not improve the metric, it is not pareto-optimal and the converse is also true; if we find a solution which, if added to the frontier, would improve on its δ score it means that this is a pareto-optimal plan.

Definition 6. For an N dimensional PF Φ a measure of its quality is NDHVI, $\delta(\Phi)$ and is equal to the N dimensional hyper-volume of the non-dominated space.




Figure 1: Example of hyper-volume indicator.




Figure 2: a) 2D PF and its δ space. b) 3D PF and its δ space.

3.4 Estimated evaluation

The actual quality as defined in Definition 6 is hard to calculate due to its irregular shape. We propose an estimated value for the PF which is calculated as follows. For a given PF, calculate the Delaunay triangulation [Delaunay, 1934] on the set of points. For each triangle calculate a volume under this triangle projected on a space $(x, y, \max Z)$, where $\max Z$ is a maximum value of Z coordinate among the three vertexes of the triangle. This volume is visualised on Figure 3 b). This approximation will never overestimate the actual value as defined in Definition 6 and therefore gives a good guideline to

the quality of our solution.




Figure 3: a) 3D Pareto-frontier with its triangulation b) Pareto-frontier and its expanded triangulation estimation of the δ space.

Evaluation of this measure is presented in the results section together with the NDHVI results for comparison.

4 Comparison of PF and PF-generators

In the previous section we have demonstrated how to evaluate PF here we show a method to compare PF and PF generators to assess which system generates better frontiers.

4.1 Comparing two frontiers

Here we describe a method to qualitatively compare two frontiers of plans. Value of δ is used as a quality measure for PF. Lets consider the following: $\Delta_1 = \frac{\delta_{\Phi_1} - \delta_{\Phi_{1+2}}}{\delta_{\Phi_{1+2}}}$ Where δ_{Φ_1} is the NDHVI of the first frontier. $\delta_{\Phi_{1+2}}$ is the NDHVI of the frontier $\Phi_{1+2} = \Phi_1 \oplus \Phi_2$. The same score is calculated for the second frontier: $\Delta_2 = \frac{\delta_{\Phi_2} - \delta_{\Phi_{1+2}}}{\delta_{\Phi_{1+2}}}$ This two scores can be then compared to determine which frontier lies closer to Φ_{1+2} , which we assume is our reference frontier.

4.2 Comparing a set of frontiers

When comparing a set of frontiers we use the same measure Δ . As the reference frontier we take a result of merging all frontiers from the set of frontiers being compared. Then their Δ values are compared and we can sort frontiers with respect to their quality.

4.3 Comparing pareto-frontier generators

In order to compare PF generators, G1 and G2, we calculate the value for Δ for each of the generators for each domain. If, as in our case, the planners are stochastic, we calculate the mean $\hat{\Delta}$ and standard deviation σ_{Δ} for each domain and for each generator. We will also calculate Δ_{FPF} for Φ_{FPF} for each domain. Comparison of $\hat{\Delta}$ tells us what improvement in quality we can expect to get after a single run of each system with confidence σ_{Δ} . Whereas Δ_{FPF} gives us an indication on the difference of quality of frontiers if the planners are given more time. Therefore both of the numbers are important for more precise comparison of planning systems.

5 MOPF: Multi-Objective Planning

Framework for pareto-frontier generation

MOPF is capable of generating sets of plans, in the form of PF. An input to the system is a PDDL domain and a problem file containing multiple objective functions. The framework can work with any metric sensitive planner. Its output

Table 1: Table presents Standard Deviation, based on 30 runs, of a PF generated in each of the iterations. Each cells represents standard deviation divided by an average PF score and the FPF size in brackets.

P	Bread	Production	Sugar	Trader
1	0.05 [9]	0.08 [82]	0.18 [2]	0.02 [9]
2	0.09 [16]	0.19 [38]	0.01 [1]	0.04 [9]
3	0.08 [8]	0.06 [101]	0 [1]	0.03 [9]
4	0.06 [8]	0.12 [67]	0.04 [6]	0.02 [7]
5	0.08 [21]	0.14 [35]	0.11 [11]	0.01 [9]
6	0.1 [21]	0 [6]	0.09 [19]	0.02 [12]
7	0.05 [15]	0.14 [22]	0 [1]	0.03 [4]
8	0.07 [41]	0 [6]	0.09 [1]	0.02 [7]
9	0.11 [17]	0.08 [22]	0.1 [15]	0.01 [10]
10	0.07 [21]	0.06 [80]	0.13 [17]	0.02 [5]
11	0.07 [28]	0.09 [76]	0.07 [1]	0.04 [6]
12	0.08 [55]	0.07 [55]	0.07 [3]	0.04 [4]
13	0.1 [29]	0 [1]	0.06 [1]	0.02 [9]
14	0.09 [56]	0.09 [6]	0.07 [2]	0.03 [7]
15	0.09 [63]	0.07 [43]	0.07 [23]	0.02 [9]
16	0.07 [59]	0.09 [44]	0.09 [22]	0.02 [13]
17	0.09 [68]	0.07 [79]	0.02 [1]	0.03 [7]
18	0.07 [42]	0.25 [15]	0.06 [8]	0.03 [6]
19	0.09 [86]	0.14 [6]	0.06 [5]	0.03 [7]
20	0.07 [95]	0.08 [134]	0.06 [38]	0.04 [10]
Avg	0.08 [37.9]	0.09 [45.9]	0.07 [8.9]	0.03 [7.95]

is a list of plans which are pareto-optimal: none of the plans is dominated by any other found in the experiment. These pareto-optimal plans represent a trade-off between the objective functions specified inside the problem file. The system also produces a graphical representation of the frontier in the space based on the objectives. In the next section LPG and LPRPG, which are used within the framework, are described. Following that a description of how our system works is given.

5.1 Planners used within the framework

LPG-TD

The first configuration of our planning system is based on LPG-TD [Gerevini, Saetti, and Serina, 2004] which is a local search, stochastic planner. In this work we use it with -n 2 option, which was experimentally chosen to be the best.

LPRPG-Stochastic

LPRPG [Coles et al., 2008] uses relaxed planning graph (RPG) heuristics combined with linear programming (LP) methods. In this work we have extended LPRPG with stochastic decisions between states which are evaluated as very similar. Stochasticity is introduced as it can substitute to some extend for metric sensitivity.

5.2 Pareto-frontier generator

In order to generate the frontier, the planning system forces the planner to search in different areas of the search space (with regards to the specified objectives). Directing the planner to various areas of the search space is achieved by combining the objectives into a single metric function and manipulating its weights. The weights used are evenly distributed across an n -simplex of objectives.

Table 2: Table presents results of comparison between LPRPG and LPG. Δ_1 and Δ_2 represents results for LPRPG and LPG respectively. Δ^* is the score for FPF.

P	$\hat{\Delta}_1$	$\hat{\Delta}_2$	$\hat{\Delta}_1 - \hat{\Delta}_2$ [σ_{Δ}]	Δ_1^*	Δ_2^*	$\Delta_1^* - \Delta_2^*$
1	1.27	1	0.27 [20.00%]	1.49	1	0.49
2	1.16	1	0.16 [27.00%]	1.51	1	0.51
3	1.36	1	0.36 [9.00%]	1.51	1	0.51
4	1.14	1.02	0.13 [27.00%]	1.34	1.01	0.34
5	1.05	1.02	0.04 [56.00%]	1.16	1.01	0.15
6	1.06	1	0.06 [12.00%]	1.06	1	0.06
7	1.12	1	0.12 [29.00%]	1.26	1	0.26
8	1.06	1	0.06 [53.00%]	1.04	1	0.04
9	1.14	1	0.14 [18.00%]	1.29	1	0.29
10	1.43	1	0.42 [14.00%]	1.45	1	0.45
11	1.47	1	0.46 [40.00%]	1.44	1	0.44
12	1.33	1	0.33 [7.00%]	1.44	1	0.44
13	2	1	1 [0.00%]	2	1	1
14	1.51	1	0.51 [7.00%]	1.82	1	0.82
15	1.26	1	0.26 [9.00%]	1.32	1	0.32
16	1.25	1	0.25 [13.00%]	1.31	1	0.31
17	1.23	1	0.23 [9.00%]	1.33	1	0.33
18	1.1	1	0.1 [40.00%]	1.28	1	0.28
19	1.71	1	0.71 [6.00%]	1.76	1	0.76
20	1.19	1	0.19 [14.00%]	1.28	1	0.28
Avg	1.29	1	0.29 [20.00%]	1.4	1	0.4

For each of the weight settings we run a planner to find a solution. Then all plans are evaluated with respect to the objectives and this information is stored. After removing dominated plans, the end result is a set of non-dominated plans in terms of these objectives.

LPG and LPRPG are not optimal planners and therefore the PF found is not guaranteed to be the set of the best non-dominated solutions possible for the domain. However from our experiments this frontier is usually a good approximation of the PF. By re-running the experiment multiple times and merging all PF (as in Definition 4) we can estimate the Final Pareto-Frontier (FPF) as defined in Definition 7. This method converges to an FPF in about twenty iterations as shown in experiments presented in Figure 4 and Figure 5 which are described in the experimental section.

Definition 7. For a set of PF Φ_i where $i=1..N$; we define the Final Pareto-Frontier (FPF) Φ_{FPF} such that $\Phi_{FPF} = \sum_{i=1}^N \Phi_i$

This is equivalent to taking the union of the sets of solutions from all of Φ_i and removing the dominated solutions.

6 Experiments and results

In our experiments we use domains Bread, Production, Trader and Sugar [Radzi, 2011], where the length of the plan does not correspond with the solution quality. This forces the planner to search for better solutions among longer plans.

6.1 Convergence to FPF

In this section we show that the FPF is a good approximation of the optimal PF by showing that its score converges to a certain value. This value depends on the planner used within the framework. The planning system in a single run generates

a PF Φ_1 . Then the system is run again and the second estimate of the PF Φ_2 is merged with Φ_1 according to Definition 4. The same process is repeated and at each stage we obtain a

new estimate of the FPF $\Phi_{FPF}^j = \sum_{i=1}^j \Phi_i$ where the sum is as in Definition 4. The distance from the final FPF is calculated as $\Delta_j = \frac{\delta(\Phi_{FPF}^j) - \delta(\Phi_{FPF})}{\delta(\Phi_{FPF})}$. The values for Δ_j generated in each step for domains Bread, Sugar and Production are presented in Figure 4 and Figure 5.

Standard deviation of the distance measure

The aim of this experiment is to check whether the frontiers generated in a single run of the system are of stable quality. With every run of the framework PF generated was clearly different, we could merge them and achieve an FPF which is, in most domains, 30% to 40% better than in a single run.

Table 1 presents aggregated data for 30 runs for each problem file. This very low deviations of the quality of frontiers generated in each of the runs of the framework confirm that even though the framework generates different frontiers, they are of similar quality.

6.2 Comparison of LPG and LPRPG within the framework

This section presents the comparison of LPG and LPRPG planners based on the method described in Section 4.3 and Production domain. To evaluate the relative difference of quality of this two planners Δ values for FPF were calculated as well as Δ values for frontiers generated in each of the runs. The results are presented in Table 2.

7 Conclusion

In this work we have presented a range of available techniques for evaluating a PF which is used in various areas of research. Modifications to current methods, and a new method of evaluating frontiers were presented. Results from Figure 4 a) and b) and Figure 5 a) and b) show that even though the new, triangulation based, method is easier to calculate it gives comparable results.

We have introduced a novel framework which for a given multi-objective domain generates a PF of plans. The framework is generic and can be used with any metric-sensitive satisficing planner. The evaluation shows that the framework can generate stable quality solutions.

A method of evaluating such systems was presented and used to evaluate the framework in configuration with LPG and modified LPRPG planners. The comparison showed that the configuration with LPG generates 29% better quality solutions in a single run and when given extra time it converges to a 40% better FPF within our framework.

As shown, the domains, mainly Bread and Production, demonstrate visible trade-offs between metrics and allow the planner to achieve the goal in variety of paths which are different in terms of the objectives. This makes them good multi-objective benchmarks.

There is much more scope for research and experimentation with the framework and developing LPRPG to exhibit

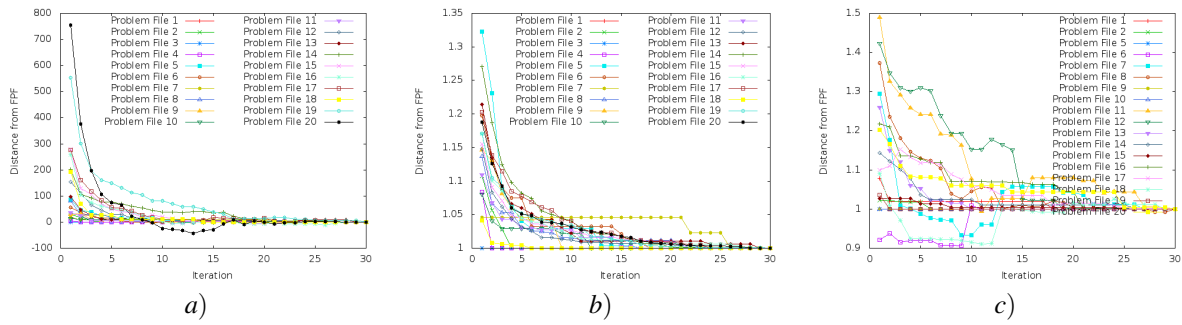


Figure 4: LPG Convergence to FPF for domain Bread using a) estimated NDHVI. b) NDHVI. c) Sugar domain using NDHVI.)

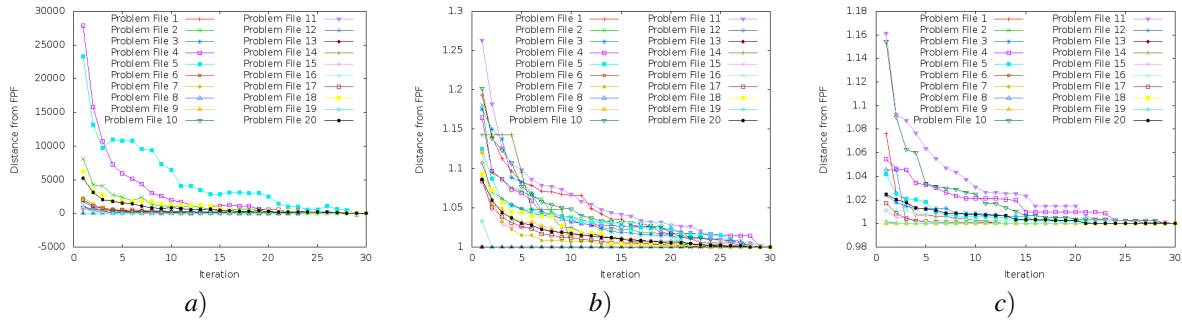


Figure 5: Convergence to FPF for domain Production using a) LPG and estimated NDHVI. b) LPG and NDHVI. c) LPRPG and NDHVI.)

more metric sensitivity due to its heuristic rather than stochastic behaviour. We are also planning to expand the portfolio of available multi-objective domains which would offer a clear trade-off between resources and paths of achieving the goal.

References

- [Coles et al., 2008] Coles, A.; Fox, M.; Long, D.; and Smith, A. 2008. A hybrid relaxed planning graph heuristic for numeric planning domains. *ICAPS*.
- [Delaunay, 1934] Delaunay, B. 1934. Sur la sphère vide. *Izvestia Akademii Nauk SSSR, Otdelenie Matematicheskikh i Estestvennykh Nauk* 7:793800.
- [Fonseca, C.M. and Paquete, L. and Lopez-Ibanez, M., 2006] Fonseca, C.M. and Paquete, L. and Lopez-Ibanez, M. 2006. An Improved Dimension-Sweep Algorithm for the Hypervolume Indicator. In *Evolutionary Computation, 2006. CEC 2006. IEEE Congress on*, volume , 1157 – 1163.
- [Fox and Long, 2003] Fox, M., and Long, D. 2003. PDDL 2.1: An Extension to pddl for Expressing Temporal Planning Domains. *Journal of Artificial Intelligence Research* 20:61–124.
- [Gerevini, Saetti, and Serina, 2004] Gerevini, A.; Saetti, A.; and Serina, I. 2004. LPG-TD: a Fully Automated Planner for PDDL2.2 Domains. *ICAPS-04*.
- [Giuliano, Rager, and Ferdous, 2007] Giuliano, M. E.; Rager, R.; and Ferdous, N. 2007. Towards a Heuristic for Scheduling the James Webb Space Telescope. *ICAPS*.
- [Horn, 1997] Horn, J. 1997. F1.9 multicriterion decision making. *Handbook of Evolutionary Computation* 97/1.
- [Johnston and Giuliano, 2011] Johnston, M. D., and Giuliano, M. 2011. Multi-Objective Scheduling for the Cluster II Constellation.
- [Klee, 1977] Klee, V. 1977. Can the Measure of $\ln 1 [a_i, b_i]$ be Computed in Less Than $O(n \log n)$ Steps? *The American Mathematical Monthly* 84:284–285.
- [McDermott et al., 1997] McDermott, D.; Ghallab, M.; Howe, A.; Knoblock, C.; Ram, A.; Veloso, M.; Weld, D.; and Wilkins, D. 1997. PDDL The Planning Domain Definition Language.
- [Nguyen et al., 2009] Nguyen, T. A.; Do, M. B.; Kambhampati, S.; and Srivastava, B. 2009. Planning with Partial preference Models. *International Joint Conferences on Artificial Intelligence IJCAI* 1:1.
- [Nguyen et al., 2011] Nguyen, T.; Do, M.; Gerevini, A.; Serina, I.; Srivastava, B.; and S.Kambahampati. 2011. Planning with Partial preference Models. *Technical Report*.
- [Radzi, 2011] Radzi, N. H. M. 2011. *Multi-Objective Planning using Linear Programming*. Ph.D. Dissertation, University of Strathclyde.
- [Sroka and Long, 2012] Sroka, M., and Long, D. 2012. Exploring Metric Sensitivity of Planners for Generation of Pareto Frontiers. In *STAIRS*, volume 241 of *Frontiers in Artificial Intelligence and Applications*, 306–317. IOS Press.
- [Ziadloo and Ghamsary, 2009] Ziadloo, M., and Ghamsary, S. S. 2009. A Framework to Evaluate Multi-Objective Optimization Algorithms in Multi-Agent Negotiations. In *CIMSA 2009*, Frontiers in Artificial Intelligence and Applications.