

# Delfi: Online Planner Selection for Cost-Optimal Planning

Michael Katz and Shirin Sohrabi and Horst Samulowitz

IBM Research

Yorktown Heights, NY, USA

michael.katz1@ibm.com, ssohrab@us.ibm.com, samulowitz@us.ibm.com

Silvan Sievers

University of Basel

Basel, Switzerland

silvan.sievers@unibas.ch

## Abstract

Cost-optimal planning has not seen many successful approaches that work well across all domains. Some cost-optimal planners excel on some domains, while exhibiting less exciting performance on others. For a particular domain, however, there is often a cost-optimal planner that works extremely well. For that reason, portfolio-based techniques have recently become popular. These either decide offline on a particular resource allocation scheme for a given collection of planners or try to perform an online classification of a given planning task to select a planner to be applied to solving the task at hand.

Our planner *Delfi* is an online portfolio planner. In contrast to existing techniques, *Delfi* exploits deep learning techniques to learn a model that predicts which of the planners in the portfolio can solve a given planning task within the imposed time and memory bounds. *Delfi* uses graphical representations of a planning task which allows exploiting existing tools for image convolution. In this planner abstract, we describe the techniques used to create our portfolio planner.

## Introduction

As planning is known to be computationally hard even for extremely conservative problem formalisms (Bylander 1994), no single planner should be expected to work well on all planning domains, or even on all tasks in a particular domain. As a result, research has not only focused on developing different planning techniques, such as improving search or heuristics, but also on exploiting multiple diverse approaches for solving planning tasks.

One such a approach is to aggregate multiple planners in a portfolio (Seipp et al. 2012; Vallati 2012; Cenamor, de la Rosa, and Fernández 2013; Seipp et al. 2015), which is what we do in this work. Such portfolios are often *sequential* and defined by two decisions: (i) which planner of the available to run next, and (ii) for how long to run it until the next planner is selected. Furthermore, the portfolio-based approaches can be partitioned in those that make those decisions ahead of time, called *offline* portfolios (Helmert et al. 2011; Núñez, Borrajo, and Linares López 2014; Seipp, Sievers, and Hutter 2014a; 2014b; 2014c) and those that make these decisions per given input task, called *online* portfolios (Cenamor, de la Rosa, and Fernández 2014).

Our planner, called *Delfi* for *DEap Learning of PortFolios*, is an online portfolio planner submitted to the Interna-

tional Planning Competition (IPC) 2018. It consists of (a) a collection of cost-optimal planners based on Fast Downward (Helmert 2006), and (b) a module that, given a planning task, selects the planner from the collection for which the confidence that it solves the given planning task is highest. Once selected, the planner is run on the given task for the entire available time. In the remainder of this planner abstract, we describe both components in detail.

## Collection of Cost-Optimal Planners

The large literature on classical planning results in an extensive pool of available planning systems that we could in principle all use. However, there are a few aspects that guided our decision to collect a rather small subset of specific planners. Firstly, the task of integrating the diverse planners within one system able to run them all in the same setting is a big (technical) challenge, and evaluating all of these planners for the training phase of learning the model would be extremely time-consuming. Secondly, portfolio planners always suffer from clearly identifying their components that are primarily responsible for the good performance of the portfolio planner.

Bearing in mind the first aspect, we restricted the pool of planners to those based on Fast Downward (Helmert 2006). This has the additional advantage that we also exploit how far a portfolio exclusively based on a single planning system fares. With respect to the second aspect, we excluded all recent (and state-of-the-art) planners that have not been evaluated in any previous competition. In particular, many of these planners are submitted independently to the IPC 2018. Furthermore, we mainly focused on planners with main components that we co-developed in order to primarily evaluate our own contributions.

These considerations result in a collection of 17 planners for our portfolio planner *Delfi*. With the exception of SymBA\* (Torralba et al. 2014), the winner of the IPC 2014, included as-is in our collection of planners, all planners are based on a recent version of Fast Downward. These 16 planners use A\* search (Hart, Nilsson, and Raphael 1968) and differ in the subsets of the following additional components they use. Please refer to the Appendix for the complete list of planner configurations of our collection, which is identical for both variants of *Delfi*.

- Pruning based on partial order reduction using strong

stubborn sets (Wehrle and Helmert 2014). Delfi uses the implementation of strong stubborn sets available in Fast Downward, which is based on the original implementation of Alkharaji et al. (2012) and Wehrle and Helmert (2012) that has also been used in Metis 2014 (Alkharaji et al. 2014). However, the current implementation has been improved in terms of efficiency since its original development.<sup>1</sup> To support conditional effects, we extended the implementation in the same way as in Metis 2014. We also use the same mechanism that disables pruning after the first 1000 expansions if only 10% or fewer states have been pruned at this point. This component is part of all 16 planners.

- Pruning based on structural symmetries (Shleyfman et al. 2015) using DKS (Domshlak, Katz, and Shleyfman 2012) or orbit space search (OSS) (Domshlak, Katz, and Shleyfman 2015). We extended the original implementation of problem description graphs, also called symmetry graphs, which serve as basis for computing symmetries, to support conditional effects. Sievers et al. (2017) recently formally defined this extension in the context of structural symmetries of lifted representations. Out of the 16 planners, 8 use DKS search and the other 8 use OSS, without any other further difference except that merge-and-shrink configurations with OSS need to disable pruning of unreachable states to avoid incorrectly reporting pruned states as dead ends (cf. Sievers et al., 2015, for more details).
- Admissible heuristics:
  - The blind heuristic.
  - The LM-cut heuristic (Helmert and Domshlak 2009). To support conditional effects, we implemented a variant of the LM-cut heuristic that considers effect conditions in the same way as Metis 2014 (Alkharaji et al. 2014) does. However, we refrain from choosing the regular LM-cut heuristic or the variant that supports conditional effects depending on the requirements of the input planning task, and instead always use the latter implementation that comes with a small overhead due to the need for different data structures.
  - The canonical pattern database (PDB) heuristic with hillclimbing to compute pattern collections, also referred to as iPDB in the literature (Haslum et al. 2007). We add a time limit of 900s to the hillclimbing algorithm.
  - The zero-one cost partitioning PDB heuristic with a genetic algorithm to compute pattern collections (Edelkamp 2006).
  - Four variants of the merge-and-shrink heuristic (Dräger, Finkbeiner, and Podelski 2009; Helmert et al. 2014; Sievers 2017). Three of them use the state-of-the-art shrink strategy based on bisimulation (Nissim, Hoffmann, and Helmert 2011) with a size limit of 50000 states on transition systems, always allowing

(perfect) shrinking, called B. The fourth variant uses a greedy variant of B, called G, not imposing any size limit on transition systems, and also always allowing shrinking. All configurations use full pruning (Sievers 2017), i.e., always prune both unreachable and irrelevant states, unless combined with OSS as discussed above, in which case pruning of unreachable states is disabled. We perform exact label reductions based on  $\Theta$ -combinability (Sievers, Wehrle, and Helmert 2014) with a fixed point algorithm using a random order on factors.

Finally, all variants use a time limit of 900s for computing the heuristic, which leads to computing so-called *partial* merge-and-shrink abstractions that do not cover all variables of the task whenever the time limit is hit. In these cases, we pick one of the remaining induced heuristics according to the following rule of thumb: we prefer the heuristic with the largest estimate for the initial state (rationale: better informed heuristic), breaking ties in favor of larger factors (rationale: more fine-grained abstraction), and choose a random heuristic among all remaining candidates of equal preference. For more details on this, we refer to the separate competition entry called Fast Downward Merge-and-Shrink (Sievers 2018).

The remaining difference between the four variants is the merge strategy, which finally results in the following merge-and-shrink configurations:

- \* B-SCCdfp: the state-of-the-art merge strategy based on *strongly connected components* of the causal graph (Sievers, Wehrle, and Helmert 2016), which uses DFP (Sievers, Wehrle, and Helmert 2014) for internal merging.
- \* B-MIASMdfp: the entirely precomputed merge strategy *maximum intermediate abstraction size minimizing* (Fan, Müller, and Holte 2014), which uses DFP as a fallback mechanism.
- \* B-sbMIASM (previously also called DYN-MIASM): the merge strategy *score-based MIASM* (Sievers, Wehrle, and Helmert 2016), which is a simple variant of MIASM.
- \* G-SCCdfp: as SCCdfp, but with the greedy variant of bisimulation-based shrinking.

As mentioned above, each heuristic is used in two planners, once with OSS and once with DKS. For the two PDB-based heuristics that do not support conditional effects natively, we compile away conditional effects by multiplying out all operators, adding copies for each possible scenario of different subsets of satisfied effect conditions and operator preconditions.

- Postprocessing the SAS<sup>+</sup> representation obtained with the translator of Fast Downward (Helmert 2009) by using the implementation of  $h^2$  mutex detection of Alcázar and Torralba (2015). This component is present in 14 out of 16 planners. The two planners that do not exploit  $h^2$  mutexes use the merge-and-shrink configuration B-MIASMdfp (once with OSS, once with DKS) which heavily relies on remaining mutexes in the SAS<sup>+</sup> repre-

<sup>1</sup>See <http://issues.fast-downward.org/issue499> and <http://issues.fast-downward.org/issue628>.




Figure 1: Images constructed from lifted and grounded representations of task pfile01-001.pddl of BARMAN-OPT11.

sensation. Our preliminary experiments showed that using the postprocessing in this case significantly harmed the performance.

### Online Planner Selection

The online planner selection of Delfi is based on a model that predicts for all planners of the portfolio whether they solve a given planning task within the fixed resource and time limits of the competition or not. To learn such a model, we created a collection of tasks to serve as training set, ran all planners in our collection on these tasks to find whether they solve the task or not, and used the resulting data to train a deep neural network. In what follows, we describe how we created the data as well as how we trained the model.

### Data Creation

Our collection of tasks includes all benchmarks of the classical tracks of all IPCs as well as some domains from the learning tracks. We further include the domains BRIEF-CASEWORLD, FERRY, and HANOI from the IPP benchmark collection (Köhler 1999), and the genome edit distance (GEDP) domain (Haslum 2011). We also use domains generated by the conformant-to-classical planning compilation (T0) (Palacios and Geffner 2009) and the finite-state controller synthesis compilation (FSC) (Bonet, Palacios, and Geffner 2009). In addition to existing tasks of these domains, we generated additional ones for some domains where generators were available. Please see the Appendix for a complete list of used domains. To filter out too hard tasks, we removed all tasks from the training set that were not solved by any of our planners.

### Data Representation

To be able to take advantage of existing deep learning tools, we need to represent planning tasks in a way that can be consumed by these tools. In the context of solving other model-based problems, such as SAT and CSP, Loreggia et al. (2016) converted the textual description of input problems to a grayscale image by converting each character to a pixel. Inspired by their ideas, we also chose to represent each task by a grayscale image of a constant size of  $128 \times 128$  pixels.

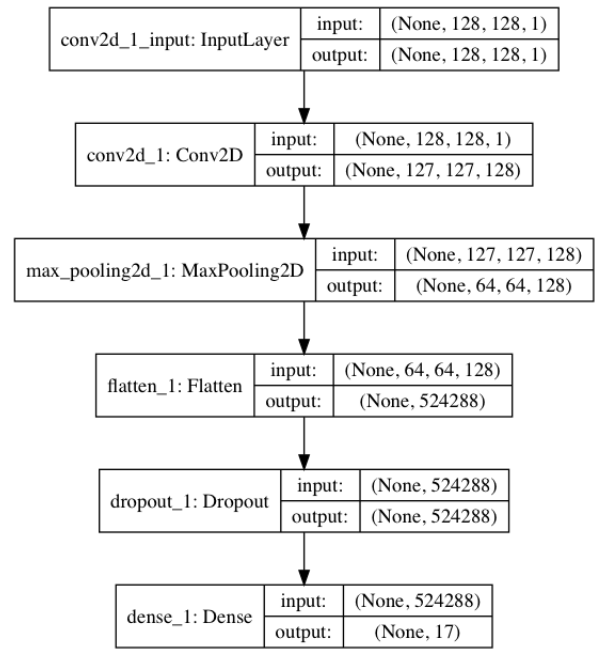


Figure 2: Visualization of the model graph structure.

However, in contrast to Loreggia et al. (2016), we chose to abstract from the textual representation and decided to use a structural representation of planning tasks, namely the *abstract structure* (Sievers et al. 2017), which encodes the PDDL description of the task. We either directly convert this abstract structure to a graph by computing the abstract structure graph as described by Sievers et al. (2017), or we first ground the abstract structure (which corresponds to grounding the planning task) and then turn it into a graph. In the latter case, we technically do not use the abstract structure graph but the conceptually equivalent *problem description graph* (Shleyfman et al. 2015), which usually is used to compute symmetries of a ground task. Finally, we turn the graph into a grayscale image that represents the adjacency matrix of the graph and reduce the grayscale image to the desired constant size.

In some preliminary experiments, we experimented with both ways of creating an image for a planning task and decided to use both. Delfi 1 computes the image from the lifted representation of the task, and Delfi 2 from the grounded one. (This is the only difference of the two variants of our planner.) Figure 1 illustrates the different images we obtain for a task of the BARMAN domain.

### Model Creation

Our tool of choice for both model creation and training is Keras (Chollet and others 2015) with Tensorflow as a backend. For our model, we employ a simple convolutional neural network (CNN) (LeCun, Bengio, and Hinton 2015) consisting of one convolutional layer, one pooling layer, one dropout layer, and one hidden layer. The main reason to choose a network with few parameters is to reduce the chances of overfitting given the comparably limited amount

of data we created. Figure 2 shows the structure of the CNN.

We model planner performance by a binary feature that indicates whether the planner solves a task within the given time (1800 seconds) and memory (7744 MiB) limits or not. We also experimented with using the actual runtime, hence not predicting whether a planner solves a task or not, but rather predicting the runtime of the planner on the task. However, our preliminary tests indicated that the performance of the network when using the binary feature is comparable to when using the actual runtime. Our conjecture is that this is due to the relatively small amount of training data and due to the fact that the model learned with the binary feature bases the decision for a planner on the confidence that this planner solves the task, which means that it is likely to prefer faster planners. As a result, we decided to use the simpler representation in our model.

Consequently, we trained the CNN by optimizing for binary cross-entropy (Rubinstein 1997) so that each planner has a certain probability assigned to it that indicates how likely it is to solve a problem within the limits. Although our CNN is rather simple, it still features a range of model hyperparameters, which we fine-tuned employing the approach by Diaz et al. (2017). For both lifted and grounded representations, the hyper-parameter optimization found very similar parameters and thus we choose the same parameters in both cases, which are as follows. The convolutional layer filter size is 3, the pooling filter size is 1, and the dropout rate is 0.48. The CNN is optimized using Stochastic Gradient Descent with learning rate 0.1, decay 0.04, momentum 0.95, nesterov set to FALSE and a batch size of 52.

## References

- Alcázar, V., and Torralba, Á. 2015. A reminder about the importance of computing and exploiting invariants in planning. In Brafman, R.; Domshlak, C.; Haslum, P.; and Zilberstein, S., eds., *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling (ICAPS 2015)*, 2–6. AAAI Press.
- Alkhazraji, Y.; Wehrle, M.; Mattmüller, R.; and Helmert, M. 2012. A stubborn set algorithm for optimal planning. In De Raedt, L.; Bessiere, C.; Dubois, D.; Doherty, P.; Frasconi, P.; Heintz, F.; and Lucas, P., eds., *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012)*, 891–892. IOS Press.
- Alkhazraji, Y.; Katz, M.; Mattmüller, R.; Pommerening, F.; Shleyfman, A.; and Wehrle, M. 2014. Metis: Arming Fast Downward with pruning and incremental computation. In *Eighth International Planning Competition (IPC-8): planner abstracts*, 88–92.
- Bonet, B.; Palacios, H.; and Geffner, H. 2009. Automatic derivation of memoryless policies and finite-state controllers using classical planners. In Gerevini, A.; Howe, A.; Cesta, A.; and Refanidis, I., eds., *Proceedings of the Nineteenth International Conference on Automated Planning and Scheduling (ICAPS 2009)*, 34–41. AAAI Press.
- Bylander, T. 1994. The computational complexity of propositional STRIPS planning. *Artificial Intelligence* 69(1–2):165–204.
- Cenamor, I.; de la Rosa, T.; and Fernández, F. 2013. Learning predictive models to configure planning portfolios. In *ICAPS 2013 Workshop on Planning and Learning*, 14–22.
- Cenamor, I.; de la Rosa, T.; and Fernández, F. 2014. IBaCoP and IBaCoPB planner. In *Eighth International Planning Competition (IPC-8): planner abstracts*, 35–38.
- Chollet, F., et al. 2015. <https://keras.io>.
- Diaz, G. I.; Fokoue-Nkoutche, A.; Nannicini, G.; and Samulowitz, H. 2017. An effective algorithm for hyperparameter optimization of neural networks. *IBM Journal of Research and Development* 61(4).
- Domshlak, C.; Katz, M.; and Shleyfman, A. 2012. Enhanced symmetry breaking in cost-optimal planning as forward search. In McCluskey, L.; Williams, B.; Silva, J. R.; and Bonet, B., eds., *Proceedings of the Twenty-Second International Conference on Automated Planning and Scheduling (ICAPS 2012)*, 343–347. AAAI Press.
- Domshlak, C.; Katz, M.; and Shleyfman, A. 2015. Symmetry breaking in deterministic planning as forward search: Orbit space search algorithm. Technical Report IS/IE-2015-03, Technion.
- Dräger, K.; Finkbeiner, B.; and Podelski, A. 2009. Directed model checking with distance-preserving abstractions. *International Journal on Software Tools for Technology Transfer* 11(1):27–37.
- Edelkamp, S. 2006. Automated creation of pattern database search heuristics. In *Proceedings of the 4th Workshop on Model Checking and Artificial Intelligence (MoChArt 2006)*, 35–50.
- Fan, G.; Müller, M.; and Holte, R. 2014. Non-linear merging strategies for merge-and-shrink based on variable interactions. In Edelkamp, S., and Barták, R., eds., *Proceedings of the Seventh Annual Symposium on Combinatorial Search (SoCS 2014)*, 53–61. AAAI Press.
- Hart, P. E.; Nilsson, N. J.; and Raphael, B. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics* 4(2):100–107.
- Haslum, P.; Botea, A.; Helmert, M.; Bonet, B.; and Koenig, S. 2007. Domain-independent construction of pattern database heuristics for cost-optimal planning. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence (AAAI 2007)*, 1007–1012. AAAI Press.
- Haslum, P. 2011. Computing genome edit distances using domain-independent planning. In *ICAPS 2011 Scheduling and Planning Applications workshop*, 45–51.
- Helmert, M., and Domshlak, C. 2009. Landmarks, critical paths and abstractions: What’s the difference anyway? In Gerevini, A.; Howe, A.; Cesta, A.; and Refanidis, I., eds., *Proceedings of the Nineteenth International Conference on Automated Planning and Scheduling (ICAPS 2009)*, 162–169. AAAI Press.
- Helmert, M.; Röger, G.; Seipp, J.; Karpas, E.; Hoffmann, J.; Keyder, E.; Nissim, R.; Richter, S.; and Westphal, M. 2011. Fast Downward Stone Soup. In *IPC 2011 planner abstracts*, 38–45.

- Helmert, M.; Haslum, P.; Hoffmann, J.; and Nissim, R. 2014. Merge-and-shrink abstraction: A method for generating lower bounds in factored state spaces. *Journal of the ACM* 61(3):16:1–63.
- Helmert, M. 2006. The Fast Downward planning system. *Journal of Artificial Intelligence Research* 26:191–246.
- Helmert, M. 2009. Concise finite-domain representations for PDDL planning tasks. *Artificial Intelligence* 173:503–535.
- Köhler, J. 1999. Handling of conditional effects and negative goals in IPP. Technical Report 128, University of Freiburg, Department of Computer Science.
- LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *Nature* 521(7553):436–444.
- Loreggia, A.; Malitsky, Y.; Samulowitz, H.; and Saraswat, V. A. 2016. Deep learning for algorithm portfolios. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI 2016)*, 1280–1286. AAAI Press.
- Nissim, R.; Hoffmann, J.; and Helmert, M. 2011. Computing perfect heuristics in polynomial time: On bisimulation and merge-and-shrink abstraction in optimal planning. In Walsh, T., ed., *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011)*, 1983–1990. AAAI Press.
- Núñez, S.; Borrajo, D.; and Linares López, C. 2014. Miplan and dpmplan. In *Eighth International Planning Competition (IPC-8): planner abstracts*.
- Palacios, H., and Geffner, H. 2009. Compiling uncertainty away in conformant planning problems with bounded width. *Journal of Artificial Intelligence Research* 35:623–675.
- Rubinstein, R. Y. 1997. Optimization of computer simulation models with rare events. *European Journal of Operational Research* 99(1):89–112.
- Seipp, J.; Braun, M.; Garimort, J.; and Helmert, M. 2012. Learning portfolios of automatically tuned planners. In McCluskey, L.; Williams, B.; Silva, J. R.; and Bonet, B., eds., *Proceedings of the Twenty-Second International Conference on Automated Planning and Scheduling (ICAPS 2012)*, 368–372. AAAI Press.
- Seipp, J.; Sievers, S.; Helmert, M.; and Hutter, F. 2015. Automatic configuration of sequential planning portfolios. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI 2015)*, 3364–3370. AAAI Press.
- Seipp, J.; Sievers, S.; and Hutter, F. 2014a. Fast Downward Cedalion. In *Eighth International Planning Competition (IPC-8): planner abstracts*, 17–27.
- Seipp, J.; Sievers, S.; and Hutter, F. 2014b. Fast Downward Cedalion. In *Eighth International Planning Competition (IPC-8) Planning and Learning Part: planner abstracts*.
- Seipp, J.; Sievers, S.; and Hutter, F. 2014c. Fast Downward SMAC. In *Eighth International Planning Competition (IPC-8) Planning and Learning Part: planner abstracts*.
- Shleyfman, A.; Katz, M.; Helmert, M.; Sievers, S.; and Wehrle, M. 2015. Heuristics and symmetries in classical planning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI 2015)*, 3371–3377. AAAI Press.
- Sievers, S.; Wehrle, M.; Helmert, M.; and Katz, M. 2015. An empirical case study on symmetry handling in cost-optimal planning as heuristic search. In Hölldobler, S.; Krötzsch, M.; Peñaloza-Nyssen, R.; and Rudolph, S., eds., *Proceedings of the 38th Annual German Conference on Artificial Intelligence (KI 2015)*, volume 9324 of *Lecture Notes in Artificial Intelligence*, 151–165. Springer-Verlag.
- Sievers, S.; Röger, G.; Wehrle, M.; and Katz, M. 2017. Structural symmetries of the lifted representation of classical planning tasks. In *ICAPS 2017 Workshop on Heuristics and Search for Domain-independent Planning*, 67–74.
- Sievers, S.; Wehrle, M.; and Helmert, M. 2014. Generalized label reduction for merge-and-shrink heuristics. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI 2014)*, 2358–2366. AAAI Press.
- Sievers, S.; Wehrle, M.; and Helmert, M. 2016. An analysis of merge strategies for merge-and-shrink heuristics. In *Proceedings of the Twenty-Sixth International Conference on Automated Planning and Scheduling (ICAPS 2016)*, 294–298. AAAI Press.
- Sievers, S. 2017. *Merge-and-shrink Abstractions for Classical Planning: Theory, Strategies, and Implementation*. Ph.D. Dissertation, University of Basel.
- Sievers, S. 2018. Fast downward merge-and-shrink. In *Ninth International Planning Competition (IPC-9): planner abstracts*.
- Torralba, Á.; Alcázar, V.; Borrajo, D.; Kissmann, P.; and Edelkamp, S. 2014. SymbA\*: A symbolic bidirectional A\* planner. In *Eighth International Planning Competition (IPC-8): planner abstracts*, 105–109.
- Vallati, M. 2012. A guide to portfolio-based planning. In Sombattheera, C.; Loi, N. K.; Wankar, R.; and Quan, T. T., eds., *Proceedings of the 6th International Workshop on Multi-disciplinary Trends in Artificial Intelligence (MIWAI 2012)*, volume 7694, 57–68. Springer.
- Wehrle, M., and Helmert, M. 2012. About partial order reduction in planning and computer aided verification. In McCluskey, L.; Williams, B.; Silva, J. R.; and Bonet, B., eds., *Proceedings of the Twenty-Second International Conference on Automated Planning and Scheduling (ICAPS 2012)*, 297–305. AAAI Press.
- Wehrle, M., and Helmert, M. 2014. Efficient stubborn sets: Generalized algorithms and selection strategies. In *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling (ICAPS 2014)*, 323–331. AAAI Press.

## Appendix

### Collection of Planner Configurations

The following are the configurations for the 16 Fast Downward based planners.

1. `--symmetries 'sym=structural_symmetries(search_symmetries=dks)'`  
`--search 'astar(blind,symmetries=sym,`  
`pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`
2. `--symmetries 'sym=structural_symmetries(search_symmetries=dks)'`  
`--search 'astar(ceilmcut,symmetries=sym,`  
`pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`
3. `--symmetries 'sym=structural_symmetries(search_symmetries=dks)'`  
`--search 'astar(`  
`merge_and_shrink(shrink_strategy=shrink_bisimulation(greedy=false),`  
`merge_strategy=merge_sccs(order_of_sccs=topological,merge_selector=score_based_filtering(scoring_functions=[goal_relevance,dfp,total_order(atomic_before_product=false,atomic_ts_order=reverse_level,product_ts_order=new_to_old)]),label_reduction=exact(before_shrinking=true,before_merging=false),max_states=50000,`  
`threshold_before_merge=1,max_time=900),`  
`symmetries=sym,pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`
4. `--symmetries 'sym=structural_symmetries(search_symmetries=dks)'`  
`--search 'astar(`  
`merge_and_shrink(shrink_strategy=shrink_bisimulation(greedy=false),`  
`merge_strategy=merge_stateless(merge_selector=score_based_filtering(scoring_functions=[sf_miasm(`  
`shrink_strategy=shrink_bisimulation,max_states=50000),total_order(atomic_before_product=true,`  
`atomic_ts_order=reverse_level,product_ts_order=old_to_new)]),label_reduction=exact(before_shrinking=true,`  
`before_merging=false),max_states=50000,threshold_before_merge=1,max_time=900),`  
`symmetries=sym,pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`
5. `--symmetries 'sym=structural_symmetries(search_symmetries=dks)'`  
`--search 'astar(`  
`merge_and_shrink(shrink_strategy=shrink_bisimulation(greedy=false),merge_strategy=merge_precomputed(`  
`merge_tree=miasm(abstraction=miasm_merge_and_shrink(),fallback_merge_selector=score_based_filtering(`  
`scoring_functions=[goal_relevance,dfp,total_order(atomic_ts_order=reverse_level,product_ts_order=new_to_old,atomic_before_product=false)])),label_reduction=exact(before_shrinking=true,before_merging=false),`  
`max_states=50000,threshold_before_merge=1,max_time=900),`  
`symmetries=sym,pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`
6. `--symmetries 'sym=structural_symmetries(search_symmetries=dks)'`  
`--search 'astar(`  
`merge_and_shrink(shrink_strategy=shrink_bisimulation(greedy=true),merge_strategy=merge_sccs(order_of_sccs=`  
`topological,merge_selector=score_based_filtering(scoring_functions=[goal_relevance,dfp,`  
`total_order(atomic_before_product=false,atomic_ts_order=level,product_ts_order=random)])),`  
`label_reduction=exact(before_shrinking=true,before_merging=false),`  
`max_states=infinity,threshold_before_merge=1,max_time=900),`  
`symmetries=sym,pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`
7. `--symmetries 'sym=structural_symmetries(search_symmetries=dks)'`  
`--search 'astar(cpdbs(patterns=hillclimbing(max_time=900),transform=multiply_out_conditional_effects),`  
`symmetries=sym,pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`
8. `--symmetries 'sym=structural_symmetries(search_symmetries=dks)'`  
`--search 'astar(zopdbs(patterns=genetic(pdb_max_size=50000,num_collections=5,num_episodes=30,`  
`mutation_probability=0.01),transform=multiply_out_conditional_effects),symmetries=sym,`  
`pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`
9. `--symmetries 'sym=structural_symmetries(search_symmetries=oss)'`  
`--search 'astar(blind,symmetries=sym,pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`
10. `--symmetries 'sym=structural_symmetries(search_symmetries=oss)'`  
`--search 'astar(ceilmcut,symmetries=sym,pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`
11. `--symmetries 'sym=structural_symmetries(search_symmetries=oss)'`  
`--search 'astar(`  
`merge_and_shrink(shrink_strategy=shrink_bisimulation(greedy=false),merge_strategy=merge_sccs(order_of_sccs=`  
`topological,merge_selector=score_based_filtering(scoring_functions=[goal_relevance,dfp,`  
`total_order(atomic_before_product=false,atomic_ts_order=reverse_level,product_ts_order=new_to_old)]),`  
`label_reduction=exact(before_shrinking=true,before_merging=false),max_states=50000,threshold_before_merge=1,`  
`max_time=900,prune_unreachable_states=false),`  
`symmetries=sym,pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`

12. `--symmetries 'sym=structural_symmetries(search_symmetries=oss)'`  
`--search 'astar(`  
`merge_and_shrink(shrink_strategy=shrink_bisimulation(greedy=false),merge_strategy=merge_stateless(merge_selector=`  
`score_based_filtering(scoring_functions=[sf_miasm(shrink_strategy=shrink_bisimulation,max_states=50000),`  
`total_order(atomic_before_product=true,atomic_ts_order=reverse_level,product_ts_order=old_to_new)])),`  
`label_reduction=exact(before_shrinking=true,before_merging=false),`  
`max_states=50000,threshold_before_merge=1,max_time=900,prune_unreachable_states=false),`  
`symmetries=sym,pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`
13. `--symmetries 'sym=structural_symmetries(search_symmetries=oss)'`  
`--search 'astar(`  
`merge_and_shrink(shrink_strategy=shrink_bisimulation(greedy=false),merge_strategy=merge_precomputed(merge_tree=`  
`miasm(abstraction=miasm_merge_and_shrink(),fallback_merge_selector=score_based_filtering(scoring_functions=`  
`[goal_relevance,dfp,total_order(atomic_ts_order=reverse_level,product_ts_order=new_to_old,`  
`atomic_before_product=false)])),label_reduction=exact(before_shrinking=true,before_merging=false),`  
`max_states=50000,threshold_before_merge=1,max_time=900,prune_unreachable_states=false),`  
`symmetries=sym,pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`
14. `--symmetries 'sym=structural_symmetries(search_symmetries=oss)'`  
`--search 'astar(`  
`merge_and_shrink(shrink_strategy=shrink_bisimulation(greedy=true),merge_strategy=merge_sccs(order_of_sccs=`  
`topological,merge_selector=score_based_filtering(scoring_functions=[goal_relevance,dfp,`  
`total_order(atomic_before_product=false,atomic_ts_order=level,product_ts_order=random)])),`  
`label_reduction=exact(before_shrinking=true,before_merging=false),max_states=infinity,`  
`threshold_before_merge=1,max_time=900,prune_unreachable_states=false),`  
`symmetries=sym,pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`
15. `--symmetries 'sym=structural_symmetries(search_symmetries=oss)'`  
`--search 'astar(cpdb(patterns=hillclimbing(max_time=900),transform=multiply_out_conditional_effects),`  
`symmetries=sym,pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`
16. `--symmetries 'sym=structural_symmetries(search_symmetries=oss)'`  
`--search 'astar(zopdb(patterns=genetic(pdb_max_size=50000,num_collections=5,num_episodes=30,`  
`mutation_probability=0.01),transform=multiply_out_conditional_effects),symmetries=sym,`  
`pruning=stubborn_sets_simple(minimum_pruning_ratio=0.01),num_por_probes=1000)'`

## Domains of the Training Set

The following lists contain all benchmark domains we used for training, named as in the repository under <https://bitbucket.org/SilvanS/ipc2018-benchmarks>. Domains with the prefix `ss` are either additional domains not contained in the original repository under <https://bitbucket.org/aibasel/downward-benchmarks> or copies of already present domains containing additional tasks that we generated.

STRIPS domains:

```
['airport', 'barman-opt11-strips', 'barman-opt14-strips', 'blocks',
'childsnaack-opt14-strips', 'depot', 'driverlog', 'elevators-opt08-strips',
'elevators-opt11-strips', 'floortile-opt11-strips',
'floortile-opt14-strips', 'freecell', 'ged-opt14-strips', 'grid',
'gripper', 'hiking-opt14-strips', 'logistics00', 'logistics98', 'miconic',
'movie', 'mprime', 'mystery', 'nomystery-opt11-strips',
'openstacks-opt08-strips', 'openstacks-opt11-strips',
'openstacks-opt14-strips', 'openstacks-strips', 'parcprinter-08-strips',
'parcprinter-opt11-strips', 'parking-opt11-strips', 'parking-opt14-strips',
'pathways-noneg', 'pegsol-08-strips', 'pegsol-opt11-strips',
'pipesworld-notankage', 'pipesworld-tankage', 'psr-small', 'rovers',
'satellite', 'scanalyzer-08-strips', 'scanalyzer-opt11-strips',
'sokoban-opt08-strips', 'sokoban-opt11-strips', 'storage',
'tetris-opt14-strips', 'tidybot-opt11-strips', 'tidybot-opt14-strips',
'tpp', 'transport-opt08-strips', 'transport-opt11-strips',
'transport-opt14-strips', 'trucks-strips', 'visitall-opt11-strips',
'visitall-opt14-strips', 'woodworking-opt08-strips',
'woodworking-opt11-strips', 'zenotravel', 'ss_barman', 'ss_ferry',
'ss_goldminer', 'ss_grid', 'ss_hanoi', 'ss_hiking', 'ss_npuzzle',
'ss_spanner',]
```

Domains with conditional effects:

```
['briefcaseworld', 'cavediving-14-adl', 'citycar-opt14-adl', 'fsc-blocks',  
'fsc-grid-a1', 'fsc-grid-a2', 'fsc-grid-r', 'fsc-hall', 'fsc-visualmarker',  
'gedp-ds2ndp', 'miconic-simpleadl', 't0-adder', 't0-coins', 't0-comm',  
't0-grid-dispose', 't0-grid-push', 't0-grid-trash', 't0-sortnet',  
't0-sortnet-alt', 't0-uts', 'ss_briefcaseworld', 'ss_cavediving',  
'ss_citycar', 'ss_maintenance', 'ss_maintenance_large', 'ss_schedule',]
```