Good Old Mercury Planner

Michael Katz
IBM Research
Yorktown Heights, NY, USA
michael.katz1@ibm.com

Jörg Hoffmann
Saarland University
Saarbrücken, Germany
hoffmann@cs.uni-saarland.de

Abstract

Mercury is a sequential satisficing planner that favorably competed in the International Planning Competition (IPC) 2014. *Mercury* planner is based mainly on the red-black planning heuristic. Red-black planning is a systematic approach to partial delete relaxation, taking into account *some* of the delete effects: Red variables take the relaxed (value-accumulating) semantics, while black variables take the regular semantics. *Mercury* planner exploits a powerful tractable fragment requiring the *black causal graph* – the projection of the causal graph onto the black variables – to be a DAG. Further, it applies techniques aimed at making red-black plans executable, short-cutting the search. As in 2014, Mercury planner is entered into sequential satisficing and agile tracks of the competition.

Planner structure

Mercury planner (Katz and Hoffmann 2014a) is a sequential satisficing planner that is implemented in the Fast Downward planning system (Helmert 2006). The planner is submitted for participation in the International Planning Competition (IPC) 2018.

Satisficing Track

The variant that competes in the *satisficing track* performs multiple iterations of heuristic search, starting with a fast and inaccurate greedy best-first search with deferred heuristic evaluation. Once a solution is found, next iterations run weighted A^* with deferred heuristic evaluation, gradually decreasing the weight parameter, similarly to the famous LAMA planning system (Richter and Westphal 2010). The cost of the best plan found so far is used in following iterations for search space pruning. Also similarly to LAMA, each search iteration alternates between four queues, two per heuristic, with all successors and successors reached by preferred operators only. The heuristics are the landmark count heuristic (Porteous, Sebastia, and Hoffmann 2001), and the red-black planning heuristic (Katz, Hoffmann, and Domshlak 2013b; 2013a; Katz and Hoffmann 2013; 2014b; Domshlak, Hoffmann, and Katz 2015). For red-black heuristic, which is based on FF (Hoffmann and Nebel 2001), the preferred operators are obtained as the preferred operators of FF heuristic.

Agile Track

The variant that competes in the *agile track* performs a single iteration of a greedy best-first search with deferred heuristic evaluation, alternating between two queues ordered by the red-black planning heuristic. These queues are filled with all successors and successors reached by preferred operators defined by the red-black planning heuristic.

Red-Black Planning Heuristic

In order to describe the configuration of the red-black planning heuristic, we need to specify how a red-black task is constructed (which variables are chosen to be red and which black), also known as painting strategy, as well as how the red-black task is solved. For red-black task construction the variables are ordered by causal graph level and iteratively painted red until the black causal graph becomes a DAG (Domshlak, Hoffmann, and Katz 2015). For solving the red-black task, the following algorithm is used: The algorithm receives a red-black planning task, as well as a set of red facts that is sufficient for reaching the red-black goals. Such a set is typically obtained from a relaxed solution to the task. Then, it iteratively (i) selects an action that can achieve some previously unachieved fact from that set, (ii) achieves its preconditions, and (iii) applies the action. Finally, when all the facts in the set are achieved, it achieves the goal of the task. There are two optimizations applied to ehnance red-black plan applicability: selecting the next action in (i) preferring actions such that achieving their black preconditions does not involve deleting facts from the set above, and selecting the sequences of actions in (ii), preferring those that are executable in the current state (Katz and Hoffmann 2014a).

Supported Features

As in the previous competition, a support for conditional effects is currently required. *Mercury* planner supports conditional effects by compiling them away. This was done by multiplying them out in the translation step. On one hand, this can lead to an exponential blow-up in the task representation size. On the other hand, it does not split up an operator application into a sequence of operator applications. Our decision was based on the speculation that the latter option could potentially decrease red-black plan applicability, one of the main advantages of the current red-black heuristics.

In order to be able to take advantage of the larger memory resource available to the participants of the current competition, the planner is built with the support for 64bit enabled.

References

Domshlak, C.; Hoffmann, J.; and Katz, M. 2015. Red-black planning: A new systematic approach to partial delete relaxation. *Artificial Intelligence* 221:73–114.

Helmert, M. 2006. The Fast Downward planning system. *Journal of Artificial Intelligence Research* 26:191–246.

Hoffmann, J., and Nebel, B. 2001. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research* 14:253–302.

Katz, M., and Hoffmann, J. 2013. Red-black relaxed plan heuristics reloaded. In Helmert, M., and Röger, G., eds., *Proceedings of the Sixth Annual Symposium on Combinatorial Search (SoCS 2013)*, 105–113. AAAI Press.

Katz, M., and Hoffmann, J. 2014a. Mercury planner: Pushing the limits of partial delete relaxation. In *Eighth International Planning Competition (IPC-8): planner abstracts*, 43–47.

Katz, M., and Hoffmann, J. 2014b. Pushing the limits of partial delete relaxation: Red-black DAG heuristics. In *ICAPS 2014 Workshop on Heuristics and Search for Domain-independent Planning (HSDIP)*, 40–44.

Katz, M.; Hoffmann, J.; and Domshlak, C. 2013a. Red-black relaxed plan heuristics. In desJardins, M., and Littman, M. L., eds., *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI 2013)*, 489–495. AAAI Press.

Katz, M.; Hoffmann, J.; and Domshlak, C. 2013b. Who said we need to relax *all* variables? In Borrajo, D.; Kambhampati, S.; Oddi, A.; and Fratini, S., eds., *Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling (ICAPS 2013)*, 126–134. AAAI Press.

Porteous, J.; Sebastia, L.; and Hoffmann, J. 2001. On the extraction, ordering, and usage of landmarks in planning. In Cesta, A., and Borrajo, D., eds., *Proceedings of the Sixth European Conference on Planning (ECP 2001)*, 174–182. AAAI Press.

Richter, S., and Westphal, M. 2010. The LAMA planner: Guiding cost-based anytime planning with landmarks. *Journal of Artificial Intelligence Research* 39:127–177.