Mercury Planner: Pushing the Limits of Partial Delete Relaxation

Michael Katz
IBM Haifa Research Labs
Haifa, Israel
katzm@il.ibm.com

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

Abstract

Mercury is a sequential satisficing planner that 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. Prior work on red-black plan heuristics has identified a powerful tractable fragment requiring the *black causal graph* – the projection of the causal graph onto the black variables – to be a DAG; but all implementations so far use a much simpler fragment where the black causal graph is required to not contain any arcs at all. We close that gap here, and we design techniques aimed at making red-black plans executable, short-cutting the search. Mercury planner is entered into sequential satisficing and agile tracks of the competition.

Planner structure

Mercury planner is a sequential satisficing planner that is implemented in the Fast Downward planning system (Helmert 2006). It performs multiple iterations of heuristic search, starting with a fast and inaccurate greedy best-first search. Once a solution is found, next iterations run weighted A^*, 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. Search algorithms are guided by the red-black heuristic (Katz, Hoffmann, and Domshlak 2013b; 2013a; Katz and Hoffmann 2013), breaking ties using the landmark count heuristic (Porteous, Sebastia, and Hoffmann 2001). In addition, preferred operators are obtained from each of those heuristics. For red-black heuristic, which is based on FF (Hoffmann and Nebel 2001), we decided to use here the preferred operators of FF heuristic. As the rest of the components are well known, in what follows, we describe in detail the main novelty of *Mercury*, red-black heuristic.

Introduction

The *delete relaxation*, where state variables accumulate their values rather than switching between them, has played a key role in the success of satisficing planning systems, e. g. (Bonet and Geffner 2001; Hoffmann and Nebel 2001; Richter and Westphal 2010). Still, the delete relaxation has well-known pitfalls, for example the fundamental inability to account for moves back and forth (as done, e. g., by vehicles in transportation). It has thus been an actively researched question from the outset how to take *some* deletes into account, e. g. (Fox and Long 2001; Gerevini, Saetti, and Serina 2003; Helmert 2004; Helmert and Geffner 2008; Baier and Botea 2009; Cai, Hoffmann, and Helmert 2009; Haslum 2012; Keyder, Hoffmann, and Haslum 2012). Herein, we continue the most recent attempt, *red-black planning* (Katz, Hoffmann, and Domshlak 2013b; 2013a; Katz and Hoffmann 2013) where a subset of *red* state variables takes on the relaxed value-accumulating semantics, while the other *black* variables retain the regular semantics.

Katz et al. (2013b) introduced the red-black framework and conducted a theoretical investigation of tractability. Following up on this (2013a), they devised practical *red-black plan heuristics*, non-admissible heuristics generated by repairing fully delete-relaxed plans into red-black plans. Observing that this technique often suffers from dramatic overestimation incurred by following arbitrary decisions taken in delete-relaxed plans, Katz and Hoffmann (2013) refined the approach to rely less on such decisions, yielding a more flexible algorithm delivering better search guidance.

The *black causal graph* is the projection of the causal graph onto the black variables only. Both Katz et al. (2013a) and Katz and Hoffmann (2013) exploit, in theory, a tractable fragment characterized by *DAG black causal graphs*, but confine themselves to *arc-empty black causal graphs* – no arcs at all – in practice. Thus current red-black plan heuristics are based on a simplistic, almost trivial, tractable fragment of red-black planning. We herein close that gap, designing *red-black DAG heuristics* exploiting the full tractable fragment previously identified. To that end, we augment Katz and Hoffmann's implementation with a DAG-planning algorithm (executed several times within every call to the heuristic function). We devise some enhancements targeted at making the resulting red-black plans executable in the real task, stopping the search if they succeed in reaching the goal.

Background

Our approach is placed in the *finite-domain representation (FDR)* framework. We introduce FDR and its delete-relaxation as special cases of red-black planning. A **red-**

black (RB) planning task is a tuple $\Pi = \langle V^{\mathsf{B}}, V^{\mathsf{R}}, A, I, G \rangle$. V^{B} is a set of *black state variables* and V^{R} is a set of *red state variables*, where $V^{\mathsf{B}} \cap V^{\mathsf{R}} = \emptyset$ and each $v \in V := V^{\mathsf{B}} \cup V^{\mathsf{R}}$ is associated with a finite domain $\mathcal{D}(v)$. The *initial state* I is a complete assignment to V, the *goal* G is a partial assignment to V. Each action a is a pair $\langle \mathsf{pre}(a), \mathsf{eff}(a) \rangle$ of partial assignments to V called *precondition* and *effect*. We often refer to (partial) assignments as sets of *facts*, i. e., variable-value pairs $v = d$. For a partial assignment p, $\mathcal{V}(p)$ denotes the subset of V instantiated by p. For $V' \subseteq \mathcal{V}(p)$, $p[V']$ denotes the value of V' in p.

A state s assigns each $v \in V$ a non-empty subset $s[v] \subseteq \mathcal{D}(v)$, where $|s[v]| = 1$ for all $v \in V^{\mathsf{B}}$. An action a is applicable in state s if $\mathsf{pre}(a)[v] \in s[v]$ for all $v \in \mathcal{V}(\mathsf{pre}(a))$. Applying a in s changes the value of $v \in \mathcal{V}(\mathsf{eff}(a)) \cap V^{\mathsf{B}}$ to $\{\mathsf{eff}(a)[v]\}$, and changes the value of $v \in \mathcal{V}(\mathsf{eff}(a)) \cap V^{\mathsf{R}}$ to $s[v] \cup \{\mathsf{eff}(a)[v]\}$. By $s[\![\langle a_1, \ldots, a_k \rangle]\!]$ we denote the state obtained from sequential application of $a_1, \ldots, a_k$. An action sequence $\langle a_1, \ldots, a_k \rangle$ is a *plan* if $G[v] \in I[\![\langle a_1, \ldots, a_k \rangle]\!][v]$ for all $v \in \mathcal{V}(G)$.

Π is a **finite-domain representation (FDR)** planning task if $V^{\mathsf{R}} = \emptyset$, and is a **monotonic finite-domain representation (MFDR)** planning task if $V^{\mathsf{B}} = \emptyset$. Plans for MFDR tasks (i. e., for delete-relaxed tasks) can be generated in polynomial time. A key part of many satisficing planning systems is based on exploiting this property for deriving heuristic estimates, via delete-relaxing the task at hand. Generalizing this to red-black planning, the **red-black relaxation** of an FDR task Π relative to V^{R} is the RB task $\Pi^{*+}_{V^{\mathsf{R}}} = \langle V \setminus V^{\mathsf{R}}, V^{\mathsf{R}}, A, I, G \rangle$. A plan for $\Pi^{*+}_{V^{\mathsf{R}}}$ is a **red-black plan** for Π, and the length of a shortest possible red-black plan is denoted $h^{*+}_{V^{\mathsf{R}}}(\Pi)$. For arbitrary states s, $h^{*+}_{V^{\mathsf{R}}}(s)$ is defined via the RB task $\langle V \setminus V^{\mathsf{R}}, V^{\mathsf{R}}, A, s, G \rangle$. If $V^{\mathsf{R}} = V$, then red-black plans are **relaxed plans**, and $h^{*+}_{V^{\mathsf{R}}}$ coincides with the optimal delete relaxation heuristic h^+.
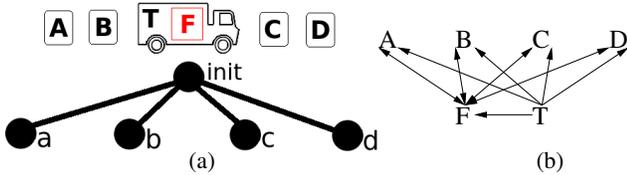
Figure 1: An example (a), and its causal graph (b).

In Figure 1, truck T needs to transport each package $X \in \{A, B, C, D\}$ to its respective goal location $x \in \{a, b, c, d\}$. The truck can only carry one package at a time, encoded by a Boolean variable F ("free"). A real plan has length 15 (8 loads/unloads, 7 drives), a relaxed plan has length 12 (4 drives suffice as there is no need to drive back). If we paint (only) T black, then $h^{*+}_{V^{\mathsf{R}}}(I) = 15$ as desired, but red-black plans may not be applicable in the real task, because F is still red so we can load several packages consecutively. Painting T *and* F black, that possibility disappears.[1]

---

[1] Indeed, all optimal red-black plans (but not some non-optimal ones) then are real plans. We will get back to this below: As we shall see, the ability to increase red-black plan applicability is a

Tractable fragments of red-black planning have been identified using standard structures. The **causal graph** CG_Π of Π is a digraph with vertices V. An arc (v, v') is in CG_Π if $v \neq v'$ and there exists an action $a \in A$ such that $(v, v') \in [\mathcal{V}(\mathsf{eff}(a)) \cup \mathcal{V}(\mathsf{pre}(a))] \times \mathcal{V}(\mathsf{eff}(a))$. The **domain transition graph** $DTG_\Pi(v)$ of a variable $v \in V$ is a labeled digraph with vertices $\mathcal{D}(v)$. The graph has an arc (d, d') induced by action a if $\mathsf{eff}(a)[v] = d'$, and either $\mathsf{pre}(a)[v] = d$ or $v \notin \mathcal{V}(\mathsf{pre}(a))$. The arc is labeled with its **outside condition** $\mathsf{pre}(a)[V \setminus \{v\}]$ and its **outside effect** $\mathsf{eff}(a)[V \setminus \{v\}]$.

The **black causal graph** CG^{B}_Π of Π is the sub-graph of CG_Π induced by V^{B}. An arc (d, d') is **relaxed side effects invertible**, **RSE-invertible** for short, if there exists an arc (d', d) with outside condition $\phi' \subseteq \phi \cup \psi$ where ϕ and ψ are the outside condition respectively outside effect of (d, d'). A variable v is RSE-invertible if all arcs in $DTG_\Pi(v)$ are RSE-invertible, and an RB task is RSE-invertible if all its black variables are. Prior work on red-black plan heuristics (Katz, Hoffmann, and Domshlak 2013a; Katz and Hoffmann 2013) proved that plan generation for RSE-invertible RB tasks with **DAG** (acyclic) black causal graphs is tractable, but used the much simpler fragment restricted to **arc-empty** black causal graphs in practice. In Figure 1, both T and F are RSE-invertible; if we paint only T black then the black causal graph is arc-empty, and if we paint both T and F black then the black causal graph is (not arc-empty but) a DAG.

Red-Black DAG Heuristics

Katz and Hoffmann (2013) provide an algorithm for RSE-invertible RB tasks with acyclic black causal graphs. To provide the context, Figure 2 shows Katz and Hoffmann's pseudo-code. The algorithm assumes as input the set R^+ of preconditions and goals on red variables in a fully delete-relaxed plan, i. e., $R^+ = G[V^{\mathsf{R}}] \cup \bigcup_{a \in \pi^+} \mathsf{pre}(a)[V^{\mathsf{R}}]$ where π^+ is a relaxed plan for Π. It then successively selects achieving actions for R^+, until all these red facts are true. Throughout the algorithm, *R denotes the set of red facts already achieved by the current red-black plan prefix π; B denotes the set of black variable values that can be achieved using only red outside conditions from R.*

For each action $a \in A'$ selected to achieve new facts from R^+, and for the global goal condition at the end, there may be black variables that do not have the required values. For example, say we paint T and F black in Figure 1. Then R^+ will have the form $\{A = T, A = a, B = T, B = b, C = T, C = c, D = T, D = d\}$. In the initial state, A' will contain only load actions. Say we execute $a = \mathsf{load}(A, init)$, entering $A = T$ into R and thus including $\mathsf{unload}(A, a)$ into A' in the next iteration. Trying to execute that action, we find that its black precondition $T = a$ is not satisfied. The call to ACHIEVE($\{T = a\}$) is responsible for rectifying this.

ACHIEVE(g) creates a task Π^{B} over Π's black variables, asking to achieve g. As Katz and Hoffmann showed, Π^{B} is solvable, has a DAG causal graph, and has strongly connected DTGs (when restricting to actions a where $\mathsf{pre}(a) \subseteq I[\![\pi]\!]$). From this and Theorem 4.4 of Chen and Gimenez

---

main advantage of our red-black DAG heuristics over the simpler red-black plan heuristics devised in earlier work.

Algorithm : REDBLACKPLANNING(Π, R^+)

main
  $// \ \Pi = \langle V^{\mathsf{B}}, V^{\mathsf{R}}, A, I, G \rangle$
  global R, $B \leftarrow \emptyset$, $\pi \leftarrow \langle \rangle$
  UPDATE()
  while $R \not\supseteq R^+$
        $\begin{cases} A' = \{a \in A \mid \mathsf{pre}(a) \subseteq B \cup R, \mathsf{eff}(a) \cap (R^+ \setminus R) \neq \emptyset\} \\ \text{Select } a \in A' \\ \textbf{if } \mathsf{pre}(a)[V^{\mathsf{B}}] \not\subseteq I[\![\pi]\!] \\ \quad \textbf{then } \pi \leftarrow \pi \circ \text{ACHIEVE}(\mathsf{pre}(a)[V^{\mathsf{B}}]) \\ \pi \leftarrow \pi \circ \langle a \rangle \\ \text{UPDATE}() \end{cases}$
  do
  if $G[V^{\mathsf{B}}] \not\subseteq I[\![\pi]\!]$
    then $\pi \leftarrow \pi \circ \text{ACHIEVE}(G[V^{\mathsf{B}}])$
  return π

procedure UPDATE()
  $R \leftarrow I[\![\pi]\!][V^{\mathsf{R}}]$
  $B \leftarrow B \cup I[\![\pi]\!][V^{\mathsf{B}}]$
  for $v \in V^{\mathsf{B}}$, ordered topologically by the black causal graph
    do $B \leftarrow B \cup DTG_{\Pi}(v)|_{R \cup B}$

procedure ACHIEVE(g)
  $I^{\mathsf{B}} \leftarrow I[\![\pi]\!][V^{\mathsf{B}}]$
  $G^{\mathsf{B}} \leftarrow g$
  $A^{\mathsf{B}} \leftarrow \{a^{\mathsf{B}} \mid a \in A, a^{\mathsf{B}} = \langle \mathsf{pre}(a)[V^{\mathsf{B}}], \mathsf{eff}(a)[V^{\mathsf{B}}] \rangle,$
        $\mathsf{pre}(a) \subseteq R \cup B, \mathsf{eff}(a)[V^{\mathsf{B}}] \subseteq B\}$
  $\langle a_1'^{\mathsf{B}}, \ldots, a_k'^{\mathsf{B}} \rangle \leftarrow$ an FDR plan for $\Pi^{\mathsf{B}} = \langle V^{\mathsf{B}}, A^{\mathsf{B}}, I^{\mathsf{B}}, G^{\mathsf{B}} \rangle$
  return $\langle a_1', \ldots, a_k' \rangle$

Figure 2: Red-black planning algorithm. $R^+ = G[V^{\mathsf{R}}] \cup \bigcup_{a \in \pi^+} \mathsf{pre}(a)[V^{\mathsf{R}}]$ where π^+ is a relaxed plan for Π.

(2010), it directly follows that a plan for Π^{B}, *in a succinct plan representation*, can be generated in polynomial time.

The "succinct plan representation" just mentioned consists of recursive macro actions for pairs of initial-value/other-value within each variable's DTG; it is required as plans for Π^{B} may be exponentially long. Chen and Gimenez' algorithm handling these macros involves the exhaustive enumeration of shortest paths for the mentioned value pairs in all DTGs, and it returns highly redundant plans moving precondition variables back to their initial value in between every two requests. For example, if a truck unloads two packages at the same location, then it is moved back to its start location in between the two unload actions.

Katz and Hoffmann (2013) shunned the complexity of DAG planning, and considered Π^{B} with arc-empty causal graphs, solving which is trivial. In our work, after exploring a few options, we decided to use the simple algorithm in Figure 3: Starting at the leaf variables and working up to the roots, the partial plan π^{B} is augmented with plan fragments bringing the supporting variables into place (a similar algorithm was mentioned, but not used, by Helmert (2006)).

Proposition 1 *The algorithm* DAGPLANNING(Π^{B}) *is sound and complete, and its runtime is polynomial in the size of* Π^{B} *and the length of the plan* π^{B} *returned.*

Note here that the length of π^{B} is worst-case exponential in the size of Π^{B}, and so is the runtime of

Algorithm : DAGPLANNING(Π^{B})

main
  $\pi^{\mathsf{B}} \leftarrow \langle \rangle$
  for $i = n$ **downto** 1
        $\begin{cases} // \ \textit{Denote } \pi^{\mathsf{B}} = \langle a_1, \ldots, a_k \rangle \\ d \leftarrow I[v_i] \\ \textbf{for } j = 1 \textbf{ to } k \\ \quad \begin{cases} \pi_j \leftarrow \langle \rangle \\ \textbf{if } \mathsf{pre}(a_j)[v_i] \text{ is defined} \\ \quad \textbf{then } \begin{cases} \pi_j \leftarrow \pi_{v_i}(d, \mathsf{pre}(a_j)[v_i]) \\ d \leftarrow \mathsf{pre}(a_j)[v_i] \end{cases} \end{cases} \\ \pi_{k+1} \leftarrow \langle \rangle \\ \textbf{if } G[v_i] \text{ is defined} \\ \quad \textbf{then } \pi_{k+1} \leftarrow \pi_{v_i}(d, G[v_i]) \\ \pi^{\mathsf{B}} \leftarrow \pi_1 \cdot \langle a_1 \rangle \cdot \ldots \cdot \pi_k \cdot \langle a_k \rangle \cdot \pi_{k+1} \end{cases}$
  do
  return π^{B}

Figure 3: Planning algorithm for FDR tasks Π^{B} with DAG causal graph $CG_{\Pi^{\mathsf{B}}}$ and strongly connected DTGs. $v_1, \ldots, v_n$ is an ordering of variables V consistent with the topology of $CG_{\Pi^{\mathsf{B}}}$. $\pi_v(d, d')$ denotes an action sequence constituting a shortest path in $DTG_v(\Pi)$ from d to d'.

DAGPLANNING(Π^{B}). We trade the theoretical worst-case efficiency of Chen and Gimenez' algorithm against the practical advantage of not having to rely on exhaustive computation of shortest paths – anew for every call of DAGPLANNING, with "initial values" and DTGs from Π^{B} – for input tasks Π^{B} that typically have small plans (achieving the next action's black preconditions) anyhow.[2]

Unlike the macro-based algorithm of Chen and Gimenez, our DAGPLANNING algorithm does not superfluously keep switching supporting variables back to their initial values. But it is not especially clever, either: If variable v_0 supports two otherwise independent leaf variables v_1 and v_2, then the sub-plans for v_1 and v_2 will be inserted sequentially into π^{B}, losing any potential for synergies in the values of v_0 required.

Painting Strategy

Katz and Hoffmann explored a variety of *painting strategies*, i.e., strategies for selecting the black variables. We kept this simple here because, as we noticed, there actually is little choice, at least when accepting the rationale that we should paint black as many variables as possible: In most IPC domains, there are at most 2 possible paintings per task. To illustrate, consider Figure 1: We can paint T and F black, or paint T and the packages black. All other paintings either do not yield a DAG black causal graph, or are not set-inclusion maximal among such paintings. We thus adopted one of Katz and Hoffmann's basic strategies, namely ordering the variables by causal graph level, and iteratively painting variables red until the black causal graph is a DAG (Katz and

---

[2]One could estimate DAG plan length (e.g., using Helmert's (2006) causal graph heuristic), computing a red-black plan *length estimate* only. But that would forgo the possibility to actually execute DAG red-black plans, which is a key advantage in practice.

Hoffmann's original strategies continue until that graph is arc-empty).

Enhancing Red-Black Plan Applicability

One crucial advantage of red-black plans, over fully-delete relaxed plans, is that they have a much higher chance of actually working for the original planning task. This is especially so for the more powerful DAG red-black plans we generate here. In Figure 1, as already mentioned, if we paint just T black then the red-black plan *might* work; but if we paint both T and F black – moving to a non-trivial DAG black causal graph – then *every optimal red-black plan definitely works*. A simple possibility for exploiting this, already implemented in Katz and Hoffmann's (2013) earlier work, is to *stop search* if the red-black plan generated for a search state s is a plan for s in the original task.

There is a catch here, though – the red-black plans we generate are not optimal and thus are not guaranteed to execute in Figure 1. In our experiments, we observed that the red-black plans often were not executable due to simple reasons. We fixed this by augmenting the algorithms with the two following applicability enhancements.

(1) Say that, as above, $R^+ = \{A = T, A = a, B = T, B = b, C = T, C = c, D = T, D = d\}$ and REDBLACKPLANNING started by selecting load$(A, init)$. Unload(A, a) *might* be next, but the algorithm might just as well select load$(B, init)$. With T and F black, load$(B, init)$ has the black precondition $F = true$. Calling ACHIEVE$(\{F = true\})$ will obtain that precondition using unload$(A, init)$. Note here that variable A is red so the detrimental side effect is ignored. The same phenomenon may occur in any domain with renewable resources (like transportation capacity). We tackle it by giving a preference to actions $a \in A'$ getting whose black preconditions does not involve deleting R^+ facts already achieved beforehand. To avoid excessive overhead, we approximate this by recording, in a pre-process, which red facts may be deleted by moving each black variable, and prefer an action if none of its black preconditions may incur any such side effects.

(2) Our second enhancement pertains to the DTG paths chosen for the black precondition variables in DAGPLANNING (after REDBLACKPLANNING has already selected the next action). The red outside conditions are by design all reached (contained in R), but we can prefer paths whose red outside conditions are "active", i. e., true when executing the current red-black plan prefix in the real task. (E.g., if a capacity variable is red, then this will prefer loads/unloads that use the actual capacity instead of an arbitrary one.) In some special cases, non-active red outside conditions can be easily fixed by inserting additional supporting actions.

Supported Features

In contrast to previous years, a support for conditional effects is currently mandated. Since there is no straightforward adaptation of the red-black heuristics to the formalism that supports conditional effects, we have chosen here to compile 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.

References

Baier, J. A., and Botea, A. 2009. Improving planning performance using low-conflict relaxed plans. In Gerevini, A.; Howe, A.; Cesta, A.; and Refanidis, I., eds., *Proceedings of the 19th International Conference on Automated Planning and Scheduling (ICAPS'09)*. AAAI Press.

Bonet, B., and Geffner, H. 2001. Planning as heuristic search. *Artificial Intelligence* 129(1–2):5–33.

Cai, D.; Hoffmann, J.; and Helmert, M. 2009. Enhancing the context-enhanced additive heuristic with precedence constraints. In Gerevini, A.; Howe, A.; Cesta, A.; and Refanidis, I., eds., *Proceedings of the 19th International Conference on Automated Planning and Scheduling (ICAPS'09)*, 50–57. AAAI Press.

Chen, H., and Giménez, O. 2010. Causal graphs and structurally restricted planning. *Journal of Computer and System Sciences* 76(7):579–592.

Fox, M., and Long, D. 2001. Stan4: A hybrid planning strategy based on subproblem abstraction. *The AI Magazine* 22(3):81–84.

Gerevini, A.; Saetti, A.; and Serina, I. 2003. Planning through stochastic local search and temporal action graphs. *Journal of Artificial Intelligence Research* 20:239–290.

Haslum, P. 2012. Incremental lower bounds for additive cost planning problems. In Bonet, B.; McCluskey, L.; Silva, J. R.; and Williams, B., eds., *Proceedings of the 22nd International Conference on Automated Planning and Scheduling (ICAPS'12)*, 74–82. AAAI Press.

Helmert, M., and Geffner, H. 2008. Unifying the causal graph and additive heuristics. In Rintanen, J.; Nebel, B.; Beck, J. C.; and Hansen, E., eds., *Proceedings of the 18th International Conference on Automated Planning and Scheduling (ICAPS'08)*, 140–147. AAAI Press.

Helmert, M. 2004. A planning heuristic based on causal graph analysis. In Koenig, S.; Zilberstein, S.; and Koehler, J., eds., *Proceedings of the 14th International Conference on Automated Planning and Scheduling (ICAPS'04)*, 161–170. Whistler, Canada: Morgan Kaufmann.

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 6th Annual Symposium on Combinatorial Search (SOCS'13)*, 105–113. AAAI Press.

Katz, M.; Hoffmann, J.; and Domshlak, C. 2013a. Red-black relaxed plan heuristics. In desJardins, M., and Littman, M., eds., *Proceedings of the 27th National Conference of the American Association for Artificial Intelligence (AAAI'13)*, 489–495. Bellevue, WA, USA: AAAI Press.

Katz, M.; Hoffmann, J.; and Domshlak, C. 2013b. Who said we need to relax *all* variables? In Borrajo, D.; Fratini, S.; Kambhampati, S.; and Oddi, A., eds., *Proceedings of the 23rd International Conference on Automated Planning and Scheduling (ICAPS'13)*, 126–134. Rome, Italy: AAAI Press.

Keyder, E.; Hoffmann, J.; and Haslum, P. 2012. Semi-relaxed plan heuristics. In Bonet, B.; McCluskey, L.; Silva, J. R.; and Williams, B., eds., *Proceedings of the 22nd International Conference on Automated Planning and Scheduling (ICAPS'12)*. 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., *Recent Advances in AI Planning. 6th European Conference on Planning (ECP-01)*, Lecture Notes in Artificial Intelligence, 37–48. Toledo, Spain: Springer-Verlag.

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