

Explicit Conjunctions without Compilation: Computing $h^{\text{FF}}(\Pi^C)$ in Polynomial Time

Jörg Hoffmann and Maximilian Fickert

Saarland University
Saarbrücken, Germany

hoffmann@cs.uni-saarland.de, s9mafick@stud.uni-saarland.de

Abstract

A successful partial delete relaxation method is to compute h^{FF} in a compiled planning task Π^C which represents a set C of conjunctions explicitly. While this compilation view of such partial delete relaxation is simple and elegant, its meaning with respect to the original planning task is opaque. We provide a direct characterization of $h^+(\Pi^C)$, without compilation, making explicit how it arises from a “marriage” of the critical-path heuristic h^m with (a somewhat novel view of) h^+ . This explicit view allows us to derive a direct characterization of $h^{\text{FF}}(\Pi^C)$, which in turn allows us to compute a version of that heuristic function in time polynomial in $|C|$.

Introduction

Explicit conjunctions were first introduced (Haslum 2009) to characterize critical-path heuristics (Haslum and Geffner 2000) as $h^m = h^1(\Pi^m)$, where Π^m is a compiled task representing each m -conjunction c via a newly introduced π -fluent π_c . A modified compilation Π^C (Haslum 2012), for arbitrary sets C of conjunctions, was shown to yield a *partial delete relaxation* method, guaranteeing to *converge* to h^* , i. e., $h^+(\Pi^C) = h^*$ for appropriately chosen C . The size of Π^C is worst-case exponential in $|C|$ because it explicitly enumerates every subset $C' \subseteq C$ of conjunctions that any application of an action a from the original planning task may be used to support. This size explosion was tackled by the Π_{ce}^C compilation (Keyder, Hoffmann, and Haslum 2012; 2014), which handles each c by a separate conditional effect. Π_{ce}^C still guarantees convergence, but loses information as it ignores *cross-context* conditions, i. e. precondition π -fluents which arise only from the combination of several $c \in C'$.

We provide a direct formulation, without compilation, of *delete relaxation over explicit conjunctions*. This makes explicit some previously opaque aspects of the approach, in particular explaining the complexity difference between Π^C and Π_{ce}^C in terms of a subgoal-choice problem easy for Π_{ce}^C but hard for Π^C . By solving that problem greedily, we compute relaxed plans for Π^C in time polynomial in $|C|$. This supersedes Π_{ce}^C , in terms of achieving the same complexity reduction without having to ignore cross-context conditions. (Alcazar et al. (2013) pursued a similar direction but, as we will detail, trivialized the subgoaling and lost convergence.)

Copyright © 2015, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Preliminaries

We employ the usual STRIPS syntax and semantics. Planning tasks are tuples $\Pi = (F, A, I, G)$ of facts, actions, initial state, and goal, each $a \in A$ being a triple $(pre(a), add(a), del(a))$ where $add(a) \cap del(a) = \emptyset$. For simplicity, we consider uniform costs (all action costs are 1); our results straightforwardly extend to arbitrary action costs.

We assume the standard notions of h^+ and h^* . We characterize heuristic functions in terms of equations over regressed subgoals. The *regression of fact set g over action a* , $R(g, a)$, is defined if $add(a) \cap g \neq \emptyset$ and $del(a) \cap g = \emptyset$. In that case, $R(g, a) = (g \setminus add(a)) \cup pre(a)$; otherwise, we write $R(g, a) = \perp$. To simplify notation, we will often identify a heuristic h with its value $h(I)$ in the initial state. All statements made generalize to arbitrary states s by setting $I := s$. By $h(\Pi')$, we denote a heuristic for Π whose value is given by applying h in a modified task Π' . To make explicit that h is computed on Π itself, we write $h(\Pi)$.

Let C be a set of fact conjunctions. We identify conjunctions with fact sets. We assume throughout that C contains all singleton conjunctions, $\{p\}$ for $p \in F$. The Π^C compilation and its relatives are based on introducing a π -fluent π_c for each $c \in C$. Using the shorthand $X^C := \{\pi_c \mid c \in C \wedge c \subseteq X\}$ for fact sets X , Π^C can be defined as follows:

Definition 1 Π^C is the planning task (F^C, A^C, I^C, G^C) , where A^C contains an action $a^{C'}$, for every pair $a \in A$ and $\emptyset \neq C' \subseteq \{c \in C \mid R(c, a) \neq \perp\}$, with $a^{C'}$ given by $pre(a^{C'}) = (pre(a) \cup \bigcup_{c' \in C'} (c' \setminus add(a)))^C$, and $add(a^{C'}) = \{\pi_{c'} \mid c' \in C'\}$. Π_{nc}^C is identical to Π^C except that $pre(a^{C'}) = pre(a)^C \cup \bigcup_{c' \in C'} (pre(a) \cup (c' \setminus add(a)))^C$.

This changes Haslum’s (2012) definition in minor ways, simplifying our presentation without affecting our results. $\emptyset \neq C'$ and $add(a^{C'}) = \{\pi_{c'} \mid c' \in C'\}$ work as C contains all singleton conjunctions. A *cross-context* condition for $a^{C'}$ is a $c \in C$ where $c \subseteq pre(a) \cup \bigcup_{c' \in C'} (c' \setminus add(a))$ but there exists no single $c' \in C'$ s.t. $c \subseteq pre(a) \cup (c' \setminus add(a))$. Π_{nc}^C ignores cross-context conditions, and is equivalent to the Π_{ce}^C compilation in that $h^+(\Pi_{ce}^C) = h^+(\Pi_{nc}^C)$.

$h^+(\Pi^C)$ w/o Compilation

We commence our investigation by “marrying” critical-path heuristics with the optimal delete-relaxation heuristic h^+ .

This results in a new characterization of $h^+(\Pi^C)$. Critical-path heuristics here will be captured by a straightforward extension h^C of the standard heuristic h^m to consider an arbitrary conjunction set C . We capture h^+ in terms of a novel equation suitable for the subsequent ‘‘marriage’’.

We define $h^C := h(G)$ where $h(\cdot)$ is the function on fact sets g that satisfies $h(g) =$

$$\begin{cases} 0 & g \subseteq I \\ 1 + \min_{a \in A, R(g,a) \neq \perp} h(R(g,a)) & g \in C \\ \max_{g' \subseteq g, g' \in C} h^C(g') & \text{else} \end{cases} \quad (1)$$

It is easy to see that h^C extends h^m , in the sense that $h^C = h^m$ if C consists of all conjunctions of size $\leq m$.

We characterize h^+ as $h^+ := h(G)$ where $h(\cdot)$ is the function on fact sets g that satisfies $h(g) =$

$$\begin{cases} 0 & g \subseteq I \\ 1 + \min_{a \in A, \emptyset \neq g' = \{p \in g \mid R(\{p\}, a) \neq \perp\}} h((g \setminus g') \cup \bigcup_{p \in g'} R(\{p\}, a)) & \text{else} \end{cases} \quad (2)$$

We prove this correct further below. To understand it intuitively, note that, for singleton subgoals, regression ignores the delete list: $R(\{p\}, a) \neq \perp$ iff $p \in \text{add}(a)$. Hence, in the bottom case, we will always have $g' = \text{add}(a) \cap g \neq \emptyset$. As, furthermore, $R(\{p\}, a) = \text{pre}(a)$, the recursive call of h will always be on $(g \setminus \text{add}(a)) \cup \text{pre}(a)$.

While the notation in Equation 2 is cumbersome, it (links to the general case below and) makes visible that the delete relaxation can be understood as *splitting subgoals up into singletons, and considering regression separately with respect to each of these*. Because regression over singleton facts ignores the deletes as outlined, in effect we need to worry only about the part of the subgoal we can support, not about other parts that the same action may contradict. This understanding underlies our ‘‘marriage of h^+ with h^C ’’: instead of the singleton facts in our subgoal, we now have to achieve (apply regression to) its C -conjunctions.

Definition 2 The explicit- C delete relaxation heuristic h^{C+} is defined as $h^{C+} := h(\{G\})$, where $h(\cdot)$ is the function on conjunction sets \mathcal{G} that satisfies $h(\mathcal{G}) =$

$$\begin{cases} 0 & \forall g \in \mathcal{G} : g \subseteq I \\ 1 + \min_{a \in A, \emptyset \neq g' \subseteq \{g \in \mathcal{G} \mid R(g,a) \neq \perp\}} h((\mathcal{G} \setminus \mathcal{G}') \cup \{\bigcup_{g \in \mathcal{G}'} R(g,a)\}) & \forall g \in \mathcal{G} : g \in C \\ h(\bigcup_{g \in \mathcal{G}} \{g' \mid g' \subseteq g, g' \in C\}) & \text{else} \end{cases}$$

h_{nc}^{C+} is identical to h^{C+} except that, in the middle case, $\{R(g,a) \mid g \in \mathcal{G}'\}$ is used instead of $\{\bigcup_{g \in \mathcal{G}'} R(g,a)\}$.

Note how Definition 2 merges the ideas underlying h^C and h^+ : We have to support *atomic subgoals* from C individually by regression (h^C); but instead of achieving only the most costly one, we have to achieve all of them, though separately (h^+). Where, previously, atomic subgoals were just facts (and hence we dealt with a set g of facts in the recursion), they now are conjunctions (hence the set \mathcal{G} of conjunctions in the recursion). Akin to Equation 1, the bottom case serves to extract the atomic subgoals (defined by our conjunction set C) from a non-atomic subgoal. Akin to Equation 2, the middle case selects the best action supporting some of our atomic subgoals. Note that using action a to ‘‘support’’ a subgoal g , which is no longer necessarily a singleton, means to comply with the full definition

of regression: a is not allowed to delete any fact in g , and $R(g,a) = (g \setminus \text{add}(a)) \cup \text{pre}(a) \supseteq \text{pre}(a)$.

Importantly, in difference to Equation 2, there is now a choice of \mathcal{G}' . This is the ‘‘subgoal-choice problem’’ mentioned in the introduction. The pairs a, \mathcal{G}' here correspond to the pairs a and C' in the actions $a^{C'}$ of Π^C and Π_{nc}^C . The choice is needed because (a) for Π^C (but not for Π_{nc}^C), larger \mathcal{G}' may give rise to additional cross-context conditions; and because (b) for both Π^C and Π_{nc}^C , it may be advantageous to achieve an atomic subgoal g later on in the recursion, with an action that has a different precondition which is easier to combine with g . To illustrate (a), say that $\text{add}(a) = \{p\}$ and $\mathcal{G} = \{\{p, q_1\}, \{p, q_2\}\}$ where q_1 and q_2 are impossible to achieve together. Then $\mathcal{G}' = \{\{p, q_1\}, \{p, q_2\}\}$ in h^{C+} (corresponding to Π^C) leads to the unsolvable subgoal $\{\{q_1, q_2\}\}$, while $\mathcal{G}' = \{\{p, q_1\}\}$ leads to the subgoal $\{\{q_1\}, \{p, q_2\}\}$ which is solvable because we can achieve each of $\{q_1\}$ and $\{p, q_2\}$ separately. To illustrate (b), say again that $\mathcal{G} = \{\{p, q_1\}, \{p, q_2\}\}$. Say that $\text{add}(a) = \{p\}$ and $\text{pre}(a) = \{r\}$ where the conjunction $\{r, q_1\}$ takes a single step to achieve, but $\{r, q_2\}$ takes $N > 2$ actions to achieve. Say that $\text{add}(a') = \{p\}$ and $\text{pre}(a') = \{r'\}$ where the conjunction $\{r', q_2\}$ takes a single step to achieve, but $\{r', q_1\}$ takes $N > 2$ actions to achieve. Then using $\mathcal{G}' = \{\{p, q_1\}, \{p, q_2\}\}$ for either of a or a' yields a solution of length $N + 2$, while using $\mathcal{G}' = \{\{p, q_1\}\}$ for a and, subsequently, $\mathcal{G}' = \{\{p, q_2\}\}$ for a' , yields a solution of length 4. So the only optimal solutions are such that neither a nor a' make maximal use of the subgoals they could support.

To prove that Definition 2 does indeed capture $h^+(\Pi^C)$, i. e., that $h^+(\Pi^C) = h^{C+}(\Pi)$, we start with the simple case where C contains only the singleton conjunctions:

Lemma 1 For $C = \{\{p\} \mid p \in F\}$, $h^+ = h^{C+}$.

Proof Sketch: For $C = \{\{p\} \mid p \in F\}$, \mathcal{G} will always be a set of singleton conjunctions, which we can identify with a set g of facts. Re-writing the h^{C+} equation yields:

$$\begin{cases} 0 & g \subseteq I \\ 1 + \min_{a \in A, \emptyset \neq g' \subseteq \{p \in g \mid R(\{p\}, a) \neq \perp\}} h((g \setminus g') \cup \bigcup_{p \in g'} R(\{p\}, a)) & \text{else} \end{cases}$$

Choosing $g' \subseteq \{p \in g \mid R(\{p\}, a) \neq \perp\}$ can only lead to a larger subgoal $(g \setminus g') \cup \text{pre}(a)$, hence we can exclude these choices and the equation simplifies to Equation 2. Simplifying the notations, that becomes:

$$\begin{cases} 0 & g \subseteq I \\ 1 + \min_{a \in A, \emptyset \neq g' = g \cap \text{add}(a)} h((g \setminus \text{add}(a)) \cup \text{pre}(a)) & \text{else} \end{cases}$$

This last equation corresponds to h^+ , proving the claim. ■

A detailed proof of Lemma 1, along with all other proofs omitted below, is available in (Hoffmann and Fickert 2015).

Theorem 1 $h^+(\Pi^C) = h^{C+}(\Pi)$.

Proof Sketch: Denoting h^{C+} for singleton conjunctions only by h^{1+} , with Lemma 1 the claim is equivalent to $h^{1+}(\Pi^C) = h^{C+}(\Pi)$. We prove this by comparing two equations, capturing $h^{1+}(\Pi^C)$ respectively $h^{C+}(\Pi)$.

For $h^{1+}(\Pi^C)$, after some simplifications our equation (called *Equation I*) reads as follows:

$$\begin{cases} 0 & \forall g \in \mathcal{G} : g \subseteq I^C \\ 1 + \min_{a \in A, \emptyset \neq \mathcal{G}' \subseteq \{\{\pi_g\} \in \mathcal{G} \mid R(g, a) \neq \perp\}} & \forall g \in \mathcal{G} : |g| = 1 \\ h((\mathcal{G} \setminus \mathcal{G}') \cup \text{pre}(a^{\mathcal{G}'})) & \text{else} \\ h(\bigcup_{g \in \mathcal{G}} \{g' \mid g' \subseteq g, |g'| = 1\}) & \end{cases}$$

For $h^{C^+}(\Pi)$, we equivalently modify Definition 2 into an equation (called *Equation II*) that works on *completed* subgoals \mathcal{G} only. \mathcal{G} is *completed* if, for all $g \in \mathcal{G}$, every $g' \subseteq g$ with $g' \in C$ is contained in \mathcal{G} as well.

$$\begin{cases} 0 & \forall g \in \mathcal{G} : g \subseteq I \\ 1 + \min_{a \in A, \emptyset \neq \mathcal{G}' \subseteq \{g \in \mathcal{G} \mid R(g, a) \neq \perp\}} & \mathcal{G} \text{ is completed and } \forall g \in \mathcal{G} : g \in C \\ h((\mathcal{G} \setminus \mathcal{G}') \cup \{\bigcup_{g \in \mathcal{G}'} R(g, a)\}) & \text{else} \\ h(\bigcup_{g \in \mathcal{G}} \{g' \mid g' \subseteq g, g' \in C\}) & \end{cases}$$

Equations I and II are isomorphic because Equation II works on C -subgoals and Equation I works on singleton π -fluents representing these same C -subgoals. Spelling this out is (notationally cumbersome but) straightforward. ■

A similar proof shows that $h^+(\Pi_{nc}^C) = h_{nc}^{C^+}(\Pi)$, and thus $h^+(\Pi_{ce}^C) = h_{nc}^{C^+}(\Pi)$. So, as desired, we obtain direct characterizations of both $h^+(\Pi^C)$ and $h^+(\Pi_{ce}^C)$.

With Theorem 1 and known results about $h^+(\Pi^C)$ (Keyder, Hoffmann, and Haslum 2014), delete relaxation over explicit conjunctions behaves exactly as expected: $h^{C^+} \geq h^C$, $h^{C^+} \geq h^+$, $h^{C^+} = \infty$ iff $h^C = \infty$, and $h^{C^+} = h^*$ for appropriately chosen C .

$h^{\text{FF}}(\Pi^C)$ w/o Compilation

We introduce a direct characterization, denoted $h^{C^{\text{FF}}}$, of $h^{\text{FF}}(\Pi^C)$. We will build on the following characterization of standard relaxed plans: $\pi^{\text{FF}} := \pi(G)$ where $\pi(\cdot)$ is a partial function on fact sets g that satisfies $\pi(g) =$

$$\begin{cases} \emptyset & g \subseteq I \\ \pi(g_r) \cup \{a\} \text{ where } a \in A, \text{add}(a) \cap g \neq \emptyset, & (3) \\ \text{and } h^1(g_r) < h^1(g \cap \text{add}(a)) & \text{else} \end{cases}$$

with $g_r := (g \setminus \text{add}(a)) \cup \text{pre}(a)$. As $g \cap \text{add}(a)$ always contains a fact p with $h^1(\{p\}) = h^1(g \cap \text{add}(a))$, this corresponds to relaxed plan extraction from the h^1 best-supporter function (Keyder and Geffner 2008). Note that $\pi(g)$ is undefined for *unsolvable subgoals* g where a *feasible* action a as requested does not exist. We say that $\pi^{\text{FF}} = \pi(G)$ is *supported* if the equation has a solution π whose domain contains G ; similar for the equations below.

From Equation 3, we obtain $h^{C^{\text{FF}}}$ by similar generalizations as we made to get from Equation 2 to h^{C^+} :

Definition 3 The explicit- C FF heuristic $h^{C^{\text{FF}}}$ is defined as $h^{C^{\text{FF}}} := \infty$ if $h^C = \infty$, else $h^{C^{\text{FF}}} := |\pi^{C^{\text{FF}}}|$ with $\pi^{C^{\text{FF}}} := \pi(\{G\})$ where $\pi(\cdot)$ is a partial function on conjunction sets \mathcal{G} that satisfies $\pi(\mathcal{G}) =$

$$\begin{cases} \emptyset & \forall g \in \mathcal{G} : g \subseteq I \\ \pi(\mathcal{G}_r) \cup \{(a, \mathcal{G}')\} \text{ where } a \in A, & \\ \emptyset \neq \mathcal{G}' \subseteq \{g \in \mathcal{G} \mid R(g, a) \neq \perp\}, & \\ \text{and } h^C(\mathcal{G}_r) < h^C(\mathcal{G}') & \forall g \in \mathcal{G} : g \in C \\ \pi(\bigcup_{g \in \mathcal{G}} \{g' \subseteq g \mid g' \in C\}) & \text{else} \end{cases}$$

Here, $\mathcal{G}_r := (\mathcal{G} \setminus \mathcal{G}') \cup \{\bigcup_{g \in \mathcal{G}'} R(g, a)\}$ and h^C is extended to sets \mathcal{G} of conjunctions by $h^C(\mathcal{G}) := \max_{g \in \mathcal{G}} h^C$.

$h_{nc}^{C^{\text{FF}}}$ and $\pi_{nc}^{C^{\text{FF}}}$ are defined in the same way except that $\mathcal{G}_r := (\mathcal{G} \setminus \mathcal{G}') \cup \{R(g, a) \mid g \in \mathcal{G}'\}$.

The step from Equation 3 to Definition 3 should be clear given the discussion in the previous section. As before, atomic subgoals are conjunctions instead of single facts, each atomic subgoal must be supported correctly, and we minimize over choices of \mathcal{G}' . The equivalent of h^1 given conjunctions C is h^C . Maintaining a set of pairs (a, \mathcal{G}') instead of a set of actions is required because the same action may be used several times, for different purposes (exactly as for the “action representatives” $a^{C'}$ in Π^C (Haslum 2012)).

Theorem 2 $\pi^{C^{\text{FF}}}$ is supported iff $h^C < \infty$.

Proof Sketch: If $h^C = \infty$, then the top level subgoal already is unsolvable. Vice versa, if $h^C < \infty$, then the $\pi^{C^{\text{FF}}}$ equation has a solution for $\pi(\{G\})$ even when restricting the choice of \mathcal{G}' , in the middle case, to singletons (i. e., to single conjunctions $\mathcal{G}' = \{g\}$). Intuitively, this is because h^C corresponds to reasoning over singleton conjunctions. ■

Theorem 3 $\pi^{C^{\text{FF}}}$, if supported, is a relaxed plan for Π^C .

Proof Sketch: Consider the sequence of pairs (a_i, \mathcal{G}'_i) selected in a solution for $\pi^{C^{\text{FF}}}$. It is easy to prove by induction over i that the sequence $a_i^{\mathcal{G}'_i}$ is a relaxed plan in Π^C . ■

Similar proofs show that the same properties hold for $\pi_{nc}^{C^{\text{FF}}}$.

There is an exponential number of choices for \mathcal{G}' in the $\pi^{C^{\text{FF}}}$ and $\pi_{nc}^{C^{\text{FF}}}$ equations. The Π^C compilation can be understood as enumerating these choices explicitly in memory, via the action representatives $a^{C'}$. So how do we solve these equations in polynomial time? The answer is, for $\pi_{nc}^{C^{\text{FF}}}$ the equation simplifies to a unique choice of \mathcal{G}' , and for $\pi^{C^{\text{FF}}}$ we can approximate that choice greedily.

For $\pi_{nc}^{C^{\text{FF}}}$, \mathcal{G}' is *decomposable* in the sense that, for $\mathcal{G}' = \mathcal{G}'^1 \cup \mathcal{G}'^2$ and corresponding new generated subgoals $\mathcal{G}_r, \mathcal{G}_r^1, \mathcal{G}_r^2$, we have $\mathcal{G}_r = \mathcal{G}_r^1 \cup \mathcal{G}_r^2$. So \mathcal{G}' is *feasible*, i. e. $h^C(\mathcal{G}_r) < h^C(\mathcal{G}')$, iff each of its elements is, and we can restrict the choice of \mathcal{G}' to be maximal, $\mathcal{G}' := \{g \in \mathcal{G} \mid R(g, a) \neq \perp, h^C(R(g, a)) < h^C(g)\}$. This essentially corresponds to what Π_{ce}^C achieves via conditional effects.

For Π^C , due to cross-context conditions, \mathcal{G}' is not decomposable. However, to get a practical heuristic function, all we need is for the choice of \mathcal{G}' to be complete (we do find feasible \mathcal{G}' if there exists one) and sound (any \mathcal{G}' we choose is feasible). Completeness is preserved even for singleton \mathcal{G}' , cf. Theorem 2. Soundness can be ensured easily during relaxed plan extraction. Our implementation keeps greedily adding new $g \in \mathcal{G}$ into \mathcal{G}' , until adding any single more g would result in $h^C(\{\bigcup_{g \in \mathcal{G}'} R(g, a)\}) \not< h^C(\mathcal{G}')$. It is interesting to note that an optimal selection of \mathcal{G}' would be hard:

Theorem 4 Given an integer K , in $\pi^{C^{\text{FF}}}$ it is **NP-complete** to decide whether there exists a feasible \mathcal{G}' with $|\mathcal{G}'| \geq K$.

Proof Sketch: Membership by guess and check. Hardness via a reduction of Hitting Set: Given a collection of sets $b \subseteq E$, the construction is such that choosing \mathcal{G}' amounts to choosing $E' \subseteq E$, where E' is feasible iff there is no b with $b \subseteq E'$. $E' \setminus E$ then is a hitting set, and maximizing $|E'|$ is equivalent to finding a minimum-size such set. ■

This result somewhat “explains” the complexity difference between Π^C and Π_{ce}^C : Π_{ce}^C exploits decomposability

for easy choice of \mathcal{G}' , whereas that choice is hard in Π^C . The Π^C compilation enumerates all possible choices, which our π^{CFF} algorithm avoids using a simple greedy algorithm.

Alcazar et al. (2013) earlier on introduced a heuristic “FF^m” which, like ours, handles explicit conjunctions without compilation. FF^m is like h^m in that it deals with all size- m conjunctions. Extending FF^m to arbitrary conjunction sets C , it corresponds in our notation to this equation:

$$\begin{cases} \emptyset & g \subseteq I \\ \pi(R(g, a)) \cup \{a\} & \text{where } a \in A, R(g, a) \neq \perp, \text{ and} \\ h^C(R(g, a)) < h^C(g \cap \text{add}(a)) & g \in C \\ \bigcup_{g' \subseteq g, g' \in C} \pi(g') & \text{else} \end{cases} \quad (4)$$

This suffers from two major weaknesses, with respect to our definition of π^{CFF} above. (1) It uses “ $\cup\{a\}$ ” instead of “ $\cup\{(a, g)\}$ ”, which bereaves the heuristic of almost all its power. It is now bounded from above by $|A|$, in contrast to h^{C+} which converges to h^* . (2) It restricts the choice of \mathcal{G}' to singletons, and thus over-simplifies the subgoal-choice problem, running the risk of excessively long relaxed plans (or rather, the heuristic would be running that risk were it not trivialized as per (1)). Alcazar et al. effectively tackle Π_{nc}^C rather than Π^C because, with singleton \mathcal{G}' , cross-context conditions never occur. Indeed, using “ $\cup\{(a, g)\}$ ” rather than “ $\cup\{a\}$ ”, Equation 4 is exactly what both π^{CFF} and π_{nc}^{CFF} simplify to when restricting the choice of \mathcal{G}' to singletons.

Experiments

We implemented h^C , h^{CFF} , and h_{nc}^{CFF} in FD (Helmert 2006). For h^C , we extended FF’s counter-based implementation of relaxed planning graphs (Hoffmann and Nebel 2001). Instead of a counter for each action precondition, we maintain a counter for each pair (a, c) of action a and conjunction $c \in C$ where $R(c, a) \neq \perp$ and $R(c, a)$ contains no mutex fact pair. The last condition prunes useless counters, and is similar to *mutex pruning* of useless actions/conditional effects in Π^C/Π_{ce}^C as discussed by Keyder et al. (2014).

For each heuristic in our experiment, we implemented and ran three tie-breaking methods for relaxed plan extraction, FF’s “difficulty” measure vs. arbitrary vs. random.¹ We ran all heuristics in FD’s lazy-greedy search with a dual open queue for preferred operators, with time/memory bounds of 30 minutes/2 GB on Intel E5-2660 machines running at 2.20 GHz, on the IPC’11 and IPC’14 satisficing benchmarks.

The primary practical advantage of our work lies in the different computation of relaxed plans for Π^C , via greedy selection of \mathcal{G}' in h^{CFF} vs. enumerative such selection in $h^{FF}(\Pi^C)$. This corresponds to a different trade-off of heuristic function speed vs. accuracy. Consider Figure 1.

Keyder et al. generate C by iterative refinement of a relaxed plan for the initial state, stopping when either a bound x on the growth of the action set in Π^C relative to the original action set A , or a time-out of 15 minutes, is reached. To examine the effect of large C on h^{CFF} vs. $h^{FF}(\Pi^C)$, we ran this process with $x = \infty$ (but keeping the time bound). As

¹The performance variance over different random seeds is consistently small. Depending on the heuristic and domain, the difference to non-random tie-breaking, however, can be large.




Figure 1: States per second (left) and initial state heuristic value (right) for h^{CFF} (x -axis) vs. $h^{FF}(\Pi^C)$ (y -axis) on the IPC’11 and IPC’14 satisficing benchmarks, with large C (size bound $x = \infty$). Both heuristics are run with FF’s “difficulty” tie breaking.

reported by Keyder et al., on IPC benchmarks, mutex pruning avoids the exponential blow-up of Π^C in $|C|$ completely. Still, an overhead remains, leading to the clear speed-up in Figure 1 (left). In Figure 1 (right), we see that the accuracy price paid, measured in terms of relaxed plan length on the initial state, is benign in comparison.

For $x = \infty$, this results in a substantial performance advantage, overall coverage being 256 with h^{CFF} vs. 217 with $h^{FF}(\Pi^C)$. However, in most domains the best performance is obtained with small size bounds x . For $x = 2$, the overall best setting in Keyder et al.’s experiments, the comparison between h^{CFF} and $h^{FF}(\Pi^C)$ is dominated by the variance over tie-breaking. The best overall coverage is 314 for h^{CFF} (random tie-breaking) vs. 301 for $h^{FF}(\Pi^C)$ (difficulty tie-breaking). There are some cases in which increasing x helps, namely ChildSnack, CityCar, Maintenance, ParcPrinter, and Tetris; there are rare cases where this leads to improved overall best-possible performance with h^{CFF} compared to $h^{FF}(\Pi^C)$. In Maintenance, for example, h^{CFF} has coverage 12 for $x = 2$, 14 for $x = 4$, and 15 for $x = 8$, while the peak coverage obtained with $h^{FF}(\Pi^C)$ is 13 for $x = 32$.

Comparing h^{CFF} to $h^{FF}(\Pi_{ce}^C)$, which like h^{CFF} avoids the worst-case exponential blow-up of Π^C , our data shows that the more complicated relaxed plan extraction in h^{CFF} does result in a runtime overhead. For $x = \infty$, states per second are typically reduced by factors between 1 and 5. This pays off only if accounting for cross-context conditions gives an advantage in informativity, which in the IPC benchmarks happens rarely. In Maintenance, the peak coverage obtained with $h^{FF}(\Pi_{ce}^C)$ is 11 (compared to 15 with h^{CFF}). For $x = 2$, as in the comparison h^{CFF} vs. $h^{FF}(\Pi^C)$, the comparison h^{CFF} vs. $h^{FF}(\Pi_{ce}^C)$ is dominated by the tie-breaking differences. The same is true of h_{nc}^{CFF} vs. $h^{FF}(\Pi_{ce}^C)$.

Conclusion

Our direct formulation of delete relaxation over explicit conjunctions is nice in being less opaque than compilations, capturing such partial delete relaxation in terms of separate regression steps over conjunctive subgoals; and in enabling polynomial-time relaxed plans for Π^C . On IPC benchmarks,

the benefits are visible but minor. It remains to be seen whether our formulation is fruitful for advanced research, such as deeper theoretical analyses of the approach.

Acknowledgments

This work was partially supported by the German Research Foundation (DFG), under grant HO 2169/5-1, and by the EU FP7 Programme under grant agreement no. 295261 (MEALS).

References

- Alcázar, V.; Borrajo, D.; Fernández, S.; and Fuentetaja, R. 2013. Revisiting regression in planning. In Rossi, F., ed., *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI'13)*, 2254–2260. AAAI Press/IJCAI.
- Bonet, B.; McCluskey, L.; Silva, J. R.; and Williams, B., eds. 2012. *Proceedings of the 22nd International Conference on Automated Planning and Scheduling (ICAPS'12)*. AAAI Press.
- Haslum, P., and Geffner, H. 2000. Admissible heuristics for optimal planning. In Chien, S.; Kambhampati, R.; and Knoblock, C., eds., *Proceedings of the 5th International Conference on Artificial Intelligence Planning Systems (AIPS-00)*, 140–149. Breckenridge, CO: AAAI Press, Menlo Park.
- Haslum, P. 2009. $h^m(P) = h^1(P^m)$: Alternative characterisations of the generalisation from h^{\max} to h^m . In Gerevini, A.; Howe, A.; Cesta, A.; and Refanidis, I., eds., *Proceedings of the 19th International Conference on Automated Planning and Scheduling (ICAPS'09)*, 354–357. AAAI Press.
- Haslum, P. 2012. Incremental lower bounds for additive cost planning problems. In Bonet et al. (2012), 74–82.
- Helmert, M. 2006. The Fast Downward planning system. *Journal of Artificial Intelligence Research* 26:191–246.
- Hoffmann, J., and Fickert, M. 2015. Explicit conjunctions w/o compilation: Computing $h^{FF}(\pi^c)$ in polynomial time (technical report). Technical report, Saarland University. Available at <http://fai.cs.uni-saarland.de/hoffmann/papers/icaps15b-tr.pdf>.
- Hoffmann, J., and Nebel, B. 2001. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research* 14:253–302.
- Keyder, E., and Geffner, H. 2008. Heuristics for planning with action costs revisited. In Ghallab, M., ed., *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI-08)*, 588–592. Patras, Greece: Wiley.
- Keyder, E.; Hoffmann, J.; and Haslum, P. 2012. Semi-relaxed plan heuristics. In Bonet et al. (2012), 128–136.
- Keyder, E.; Hoffmann, J.; and Haslum, P. 2014. Improving delete relaxation heuristics through explicitly represented conjunctions. *Journal of Artificial Intelligence Research* 50:487–533.