

# Fork-Decoupled State Space Search (Technical Report)

Daniel Gnad and Jörg Hoffmann

Saarland University  
Saarbrücken, Germany  
{gnad, hoffmann}@cs.uni-saarland.de

## Abstract

Factored planning decouples planning tasks into subsets (*factors*) of state variables. The traditional focus is on handling complex cross-factor interactions. Departing from this, we introduce a form of *target-profile factoring*, forcing the cross-factor interactions to take the form of a *fork*, with several *leaf factors* and one potentially very large *root factor*. We show that forward state space search gracefully extends to such structure, by augmenting regular search on the root factor with maintenance of the cheapest compliant paths within each leaf factor. We analyze how to guarantee optimality. Connecting to standard heuristics, the performance improvements relative to  $A^*$  are substantial, and sometimes dramatic: In four IPC benchmark domains, fork-decoupled state space search outperforms standard state space search even when using  $h^{\max}$  in the former vs. LM-cut in the latter.

## Introduction

Factored planning decouples planning tasks into subsets, *factors*, of state variables. In *localized* factored planning (Laneky and Getoor 1995; Amir and Engelhardt 2003; Brafman and Domshlak 2006; 2008; 2013; Fabre et al. 2010), two factors interact if they are affected by common actions, and a global plan needs to comply with these *cross-factor interactions*. In *hierarchical* factored planning (e. g. (Sacerdoti 1974; Knoblock 1994; Kelareva et al. 2007)), the factors are used within a hierarchy of increasingly more detailed abstraction levels, accumulating the factors processed so far as search proceeds down the hierarchy.

In localized factored planning, much of the complexity stems from having to resolve complex cross-factor interactions. In hierarchical factored planning, at higher levels we don't know whether the abstract plan will be realizable at lower levels, leading to backtracking across levels. So what both localized and hierarchical factored planning algorithms share is *the ability to deal with very complex forms of cross-factor interactions*, and the difficulties entailed.

But do we actually need that ability? Provided we are willing to handle very large “factors”, we can force the cross-factor interactions to be simple. We baptize this approach *target-profile factoring*: fixing a particular *profile* which we want the factoring to induce. Here, we instantiate this with *fork* (Katz and Domshlak 2008) profiles, where a single *root factor* provides preconditions required for mov-

ing several *leaf factors*. The root factor may be arbitrarily large, as we will tackle it by (heuristic) search. To obtain such a *fork factoring*, we view the causal graph as a DAG of strongly connected components (SCC) in which the root SCCs are at the top and the leaf SCCs at the bottom. We draw a horizontal line anywhere through that DAG, perceive the (entire) top part as the root factor, and perceive each weakly connected component inside the bottom part as a leaf factor. Of course, not every planning task has a useful fork factoring (e. g. if the causal graph is a single SCC). We will see this within negligible runtime at factoring time, and if so, simply *abstain* from solving the task.

The key advantage of fork factorings is a particular form of *conditional independence*: *for any fixed root path, the leaves can be moved independently of each other*. We exploit this by augmenting regular search on the root factor with maintenance of the cheapest *compliant paths* (paths that comply with the moves of the root) independently within each leaf factor. Compared to standard state space search, such *fork-decoupled state space search* can drastically reduce search space size, and in contrast to hierarchical factoring it avoids backtracking across levels.

We show how to naturally extend standard concepts such as  $A^*$  and admissible heuristics, how to connect to standard heuristics, and how to guarantee (global cost-)optimality. Empirically, for satisficing planning the decoupling helps dramatically in transportation with fuel consumption (IPC NoMystery), but only there, because other fork-like domains are already easy for existing heuristic search techniques. For optimal planning, the empirical results are almost consistently good, and, in some domains, quite amazing.

It is worth pointing out at this early stage already that target-profile factoring, beyond fork profiles, suggests an entirely new way of exploiting structure:

*Instead of relaxing the planning task into a (structurally defined) fragment to obtain a heuristic function, try to factorize the task into the fragment to obtain a plan.*

This suggests a new direction for causal graph research, designing fragments suited to specialized combinatorial search algorithms, as opposed to tractability analysis. In the long term, this could lead to an entire portfolio of target profiles.

## Preliminaries

We use the finite-domain state variable setting (e.g. (Bäckström and Nebel 1995; Helmert 2006)). A *finite-domain representation* planning task, short FDR task, is a quadruple  $\Pi = \langle V, A, I, G \rangle$ .  $V$  is a set of *state variables*, where each  $v \in V$  is associated with a finite domain  $\mathcal{D}(v)$ . We identify (partial) variable assignments with sets of variable/value pairs. A complete assignment to  $V$  is a *state*.  $I$  is the *initial state*, and the goal  $G$  is a partial assignment to  $V$ .  $A$  is a finite set of *actions*. Each action  $a \in A$  is a triple  $\langle \text{pre}(a), \text{eff}(a), \text{cost}(a) \rangle$  where the *precondition*  $\text{pre}(a)$  and *effect*  $\text{eff}(a)$  are partial assignments to  $V$ , and  $\text{cost}(a) \in \mathbb{R}^{0+}$  is the action’s non-negative *cost*.

For a partial assignment  $p$ ,  $\mathcal{V}(p) \subseteq V$  denotes the subset of state variables instantiated by  $p$ . For any  $V' \subseteq \mathcal{V}(p)$ , by  $p[V']$  we denote the value of  $V'$  in  $p$ . An action  $a$  is *applicable* in a state  $s$  if  $\text{pre}(a) \subseteq s$ , i.e., if  $s[v] = \text{pre}(a)[v]$  for all  $v \in \mathcal{V}(\text{pre}(a))$ . Applying  $a$  in  $s$  changes the value of each  $v \in \mathcal{V}(\text{eff}(a))$  to  $\text{eff}(a)[v]$ , and leaves  $s$  unchanged elsewhere; the outcome state is denoted  $s[a]$ . The outcome state of applying a sequence of (respectively applicable) actions is denoted  $s[\langle a_1, \dots, a_n \rangle]$ . A *plan* for  $\Pi$  is an action sequence s.t.  $G \subseteq I[\langle a_1, \dots, a_n \rangle]$ . The plan is *optimal* if its summed-up cost is minimal among all plans for  $\Pi$ .

The *causal graph* of a planning task captures state variable dependencies incurred by co-occurrences in actions (e.g. (Knoblock 1994; Jonsson and Bäckström 1995; Brafman and Domshlak 2003; Helmert 2006)). We use the commonly employed definition in the FDR context, where the causal graph  $CG$  is a directed graph over vertices  $V$ , with an arc from  $v$  to  $v'$ , which we denote  $(v \rightarrow v')$ , if  $v \neq v'$  and there exists an action  $a \in A$  such that  $(v, v') \in [\mathcal{V}(\text{eff}(a)) \cup \mathcal{V}(\text{pre}(a))] \times \mathcal{V}(\text{eff}(a))$ . We assume for simplicity that  $CG$  is weakly connected (else, the task can be equivalently split into several independent tasks).

**Example 1** *As an illustrative running example, we consider a transportation task with one package  $p$ , and two trucks  $t_A, t_B$  moving along three locations  $l_1, l_2, l_3$  arranged in a line. Precisely, the FDR planning task  $\Pi = \langle V, A, I, G \rangle$  is as follows.  $V = \{p, t_A, t_B\}$  where  $\mathcal{D}(p) = \{A, B, l_1, l_2, l_3\}$  and  $\mathcal{D}(t_A) = \mathcal{D}(t_B) = \{l_1, l_2, l_3\}$ . The initial state is  $I = \{p = l_1, t_A = l_1, t_B = l_3\}$ , i.e.,  $p$  and  $t_A$  start at  $l_1$ , and  $t_B$  starts at  $l_3$ . The goal is  $G = \{p = l_3\}$ . The actions (uniform costs) are the usual truck moves and load/unload, that is:*

- *move*( $x, y, z$ ) with precondition  $\{t_x = y\}$  and effect  $\{t_x = z\}$ , where  $x \in \{A, B\}$  and  $\{y, z\} \in \{\{1, 2\}, \{2, 3\}\}$ .
- *load*( $x, y$ ) with precondition  $\{t_x = y, p = y\}$  and effect  $\{p = x\}$ , where  $x \in \{A, B\}$  and  $y \in \{l_1, l_2, l_3\}$ .
- *unload*( $x, y$ ) with precondition  $\{t_x = y, p = x\}$  and effect  $\{p = y\}$ , where  $x \in \{A, B\}$  and  $y \in \{l_1, l_2, l_3\}$ .

The causal graph of this task is shown in Figure 1.

## Fork Factoring

Our definitions of factorings and forks follow those by Brafman and Domshlak (2006) and Katz and Domshlak (2008):




Figure 1: The causal graph of our illustrative example.




Figure 2: Examples of “horizontal lines”  $\{T, B\}$  through a DAG of SCCs, and the corresponding fork factorings.

**Definition 1 (Fork Factoring)** *Let  $\Pi$  be an FDR task with variables  $V$ . A factoring  $\mathcal{F}$  is a partition of  $V$  into non-empty subsets  $F$ , called factors. The interaction graph  $IG(\mathcal{F})$  of  $\mathcal{F}$  is the directed graph whose vertices are the factors, and with an arc  $(F \rightarrow F')$  if  $F \neq F'$  and there exist  $v \in F$  and  $v' \in F'$  such that  $(v \rightarrow v')$  is an arc in  $CG$ .*

*$\mathcal{F}$  is a fork factoring if  $|\mathcal{F}| > 1$  and there exists  $F^R \in \mathcal{F}$  s.t. the arcs in  $IG(\mathcal{F})$  are exactly  $\{(F^R \rightarrow F^L) \mid F^L \in \mathcal{F} \setminus \{F^R\}\}$ .  $F^R$  is the root of  $\mathcal{F}$ , and all other factors  $F^L \in \mathcal{F}^L := \mathcal{F} \setminus \{F^R\}$  are leaves. We also consider the trivial fork factoring where  $F^R = V$  and  $\mathcal{F}^L = \emptyset$ , and the pathological fork factoring where  $F^R = \emptyset$  and  $\mathcal{F}^L = \{V\}$ .*

The only global interactions in a fork factoring consist in the root factor establishing preconditions for actions moving individual leaf factors. But when do fork factorings exist, and how to find one? Denote by  $\mathcal{F}^{\text{SCC}}$  the factoring whose factors are the strongly connected components of  $CG$ . Clearly, any fork factoring  $\mathcal{F}$  must be coarser than  $\mathcal{F}^{\text{SCC}}$ , i.e., for every  $F \in \mathcal{F}^{\text{SCC}}$  we must have  $F' \in \mathcal{F}$  with  $F \subseteq F'$ . In other words, we must perform the factoring at the level of SCCs in the causal graph. In particular, if  $|\mathcal{F}^{\text{SCC}}| = 1$  (the causal graph consists of a single SCC), then the only fork factorings are the trivial and pathological ones.

For the case where  $|\mathcal{F}^{\text{SCC}}| > 1$ , consider Figure 2. We view the interaction graph  $IG(\mathcal{F}^{\text{SCC}})$  over SCCs as a DAG where the root SCCs are at the top and the leaf SCCs at the bottom. Let  $\{T, B\}$  (top, bottom) be a “horizontal line” through that DAG, i.e., a partition of  $V$  where every  $F \in \mathcal{F}^{\text{SCC}}$  is fully contained in either of  $T$  or  $B$ , and where the only arc in  $IG(\{T, B\})$  is  $(T \rightarrow B)$ . Let  $\mathcal{W}$  be the set of weakly connected components of  $\mathcal{F}^{\text{SCC}}$  within  $B$ . Then we obtain a fork factoring  $\mathcal{F}$  by setting  $F^R := T$  and  $\mathcal{F}^L := \mathcal{W}$ . Vice versa, any non-trivial fork factoring can be obtained in this manner, except the *redundant* ones where some  $F^L \in \mathcal{F}^L$  contains several weakly connected components from  $\mathcal{W}$ . Thus finding a (non-redundant) fork factoring is equivalent to finding the “horizontal line”  $\{T, B\}$ .

Our concrete strategy for finding that line is quite simple, and will be described in the experiments section. The strategy tries to obtain a factoring with many small leaf factors, which is when our method typically works well: a lot of conditional independence, not much overhead. In particular, if we don’t find a factoring with at least two leaves (which we will know after negligible runtime), we *abstain* from solv-

ing the input task. We can then pass the task on to standard techniques (or, in the long term, to other target profiles).

**Example 2** Consider again our illustrative example. Given its causal graph, cf. Figure 1, there are exactly four possible fork factorings  $\mathcal{F}$  for this planning task:  $\{\{t_A\}, \{t_B, p\}\}$ ,  $\{\{t_B\}, \{t_A, p\}\}$ ,  $\{\{t_A, t_B\}, \{p\}\}$ , and  $\{\{t_A, t_B, p\}\}$ . In the first three cases, there is a single leaf containing the package; the factorings differ regarding whether they include a truck into that leaf, and if so, which one(s). The last case is the trivial factoring, where there is no leaf at all, and all variables are contained in the fork root.

We will henceforth consider the third of these four factorings,  $\{\{t_A, t_B\}, \{p\}\}$ , where the trucks together form the root and the package on its own forms the leaf.




Figure 3: Possible fork factorings in transportation with fuel consumption (left), and job-planning problems (right).

To illustrate the kind of structure we can exploit, consider Figure 3. On the left (as in IPC NoMystery), a truck  $t$  with fuel supply  $f$  transports packages  $p_1, \dots, p_n$ ;  $t$  and  $f$  form an SCC in the causal graph, each  $p_i$  is a leaf SCC on its own. On the right (similar to IPC Woodworking), individual objects  $o_i$  are mutually independent except for sharing the machines, so that each leaf SCC contains the properties pertaining to one  $o_i$ . In both cases, our factoring strategy will come up with the fork factoring shown in the figure.

### Fork-Decoupled State Spaces

We assume an FDR task  $\Pi = \langle V, A, I, G \rangle$  and a fork factoring  $\mathcal{F}$  with root  $F^R$  and leaves  $F^L \in \mathcal{F}^L$ . We refer to the actions  $A^R$  affecting the root as *root actions*, notation convention  $a^R$ , and to the actions  $A^L$  affecting  $V \setminus F^R$  (i. e. any of the leaves) as *leaf actions*, notation convention  $a^L$ . For the set of actions affecting one particular  $F^L \in \mathcal{F}^L$ , we write  $A^L|_{F^L}$ . Bear in mind that the root actions  $A^R$  have preconditions and effects only on  $F^R$ , and the leaf actions  $A^L|_{F^L}$  have preconditions only on  $F^R \cup F^L$ , and effects only on  $F^L$ . In particular, the sets  $A^R$  and  $A^L|_{F^L}$  form a partition of the original action set  $A$ . A *root path* is a sequence of root actions applicable to  $I$ ; a *leaf path* is a sequence of leaf actions applicable to  $I$  when ignoring preconditions on the root. Value assignments to  $F^R$  are *root states*, notated  $s^R$ , and value assignments to any  $F^L \in \mathcal{F}^L$  are *leaf states*, notated  $s^L$ . For the leaf states of one particular  $F^L \in \mathcal{F}^L$ , we write  $S^L|_{F^L}$ , and for the set of all leaf states we write  $S^L$ . A root state  $s^R$  is a *goal root state* if  $s^R \supseteq G[F^R]$ , and a leaf state  $s^L \in S^L|_{F^L}$  is a *goal leaf state* if  $s^L \supseteq G[F^L]$ . Finally, for a root state  $s^R$ ,  $A^L|_{s^R}$  denotes the leaf actions enabled by  $s^R$ , i. e., those where  $\text{pre}(a^L)[F^R] \subseteq s^R$ .

The fork-decoupled state space augments root paths with maintenance of cheapest compliant leaf paths. Leaf path  $\pi^L = \langle a_1^L, \dots, a_n^L \rangle$  complies with root path  $\pi^R$  if we can

schedule the  $a_i^L$  at monotonically increasing points alongside  $\pi^R$  so that each  $a_i^L$  is enabled in the respective root state  $s^R$ , i. e.,  $a_i^L \in A^L|_{s^R}$ . Formally:

**Definition 2 (Compliant Paths)** Let  $\Pi$  be an FDR task,  $\mathcal{F}$  a fork factoring, and  $\pi^R$  a root path traversing root states  $\langle s_0^R, \dots, s_n^R \rangle$ . For a leaf path  $\pi^L = \langle a_1^L, \dots, a_m^L \rangle$ , an embedding into  $\pi^R$  is a monotonically increasing function  $e : \{1, \dots, m\} \mapsto \{0, \dots, n\}$  so that, for every  $i \in \{1, \dots, m\}$ ,  $a_i^L \in A^L|_{s_{e(i)}^R}$ . We say that  $\pi^L$  complies with  $\pi^R$ , short  $\pi^L$  is  $\pi^R$ -compliant, if an embedding exists.

Where the root path in question is clear from context, or when discussing compliant paths in general, we will omit “ $\pi^R$ ” and simply talk about *compliant leaf paths*.

**Example 3** In our illustrative example,  $\pi^L = \langle \text{load}(A, l_1) \rangle$  complies with the empty center path, and  $\pi^L = \langle \text{load}(A, l_1), \text{unload}(A, l_2) \rangle$  complies with the center path  $\pi^C = \langle \text{move}(A, l_1, l_2) \rangle$ . But, e. g.,  $\pi^L = \langle \text{load}(A, l_1), \text{unload}(A, l_3) \rangle$  does not comply with  $\pi^C$  because the required precondition  $t_A = l_3$  is not established on  $\pi^C$ . And  $\pi^L = \langle \text{load}(A, l_1), \text{unload}(A, l_2), \text{load}(A, l_2), \text{load}(A, l_1) \rangle$  does not comply with  $\pi^C$  because the last precondition  $t_A = l_1$  does not appear behind  $t_A = l_2$  on  $\pi^C$ .

**Definition 3 (Fork-Decoupled State Space)** Let  $\Pi$  be an FDR task, and  $\mathcal{F}$  a fork factoring. A fork-decoupled state  $s$  is a pair  $\langle \text{rt}(s), \text{prices}(s) \rangle$  with  $\text{rt}(s)$  being a root state, and  $\text{prices}(s) : S^L \mapsto \mathbb{R}^{0+} \cup \{\infty\}$  being a function which assigns every leaf state a non-negative price. The fork-decoupled state space is a labeled transition system  $\Theta^{\mathcal{F}} = \langle S^{\mathcal{F}}, A^R, T^{\mathcal{F}}, I^{\mathcal{F}}, S_G^{\mathcal{F}} \rangle$  as follows:

- (i)  $S^{\mathcal{F}}$  is the set of all fork-decoupled states  $s$ .
- (ii)  $A^R$ , the set of root actions, gives the transition labels.
- (iii)  $T^{\mathcal{F}}$  is the set of transitions, with  $(s \xrightarrow{a} t) \in T^{\mathcal{F}}$  if  $a \in A^R$ ,  $\text{rt}(s)[a] = \text{rt}(t)$ , and, for every leaf state  $t^L$ ,  $\text{prices}(t)[t^L] = \min_{s^L} (\text{prices}(s)[s^L] + c(s^L))$  where  $c(s^L)$  is the cost of a cheapest path of  $A^L|_{\text{rt}(t)}$  actions from  $s^L$  to  $t^L$  if such a path exists, else  $c(s^L) := \infty$ .
- (iv)  $I^{\mathcal{F}}$  is the fork-decoupled initial state, where  $\text{rt}(I^{\mathcal{F}}) := I[F^R]$ , and  $\text{prices}(I^{\mathcal{F}})[s^L]$ , for every leaf state  $s^L \in S^L|_{F^L}$  for some  $F^L \in \mathcal{F}^L$ , is the cost of a cheapest path of  $A^L|_{\text{rt}(I^{\mathcal{F}})}$  actions from  $I[F^L]$  to  $s^L$  if such a path exists, else  $\text{prices}(I^{\mathcal{F}})[s^L] := \infty$ .
- (v)  $S_G^{\mathcal{F}}$  are the fork-decoupled goal states  $s_G$ , where  $\text{rt}(s_G)$  is a goal root state and, for every  $F^L \in \mathcal{F}^L$ , there exists a goal leaf state  $s^L \in S^L|_{F^L}$  s.t.  $\text{prices}(s_G)[s^L] < \infty$ .

The cost of a path  $\pi^{\mathcal{F}} = s_0 \xrightarrow{a_1} s_1 \dots s_{n-1} \xrightarrow{a_n} s_n$  in  $\Theta^{\mathcal{F}}$  is  $\text{cost}(\pi^{\mathcal{F}}) := \sum_{i=1}^n \text{cost}(a_i)$ . A fork-decoupled plan for  $\Pi$  is a path  $\pi^{\mathcal{F}}$  from  $I^{\mathcal{F}}$  to an element of  $S_G^{\mathcal{F}}$ .

Note that, as stated,  $S^{\mathcal{F}}$  is infinite. The *reachable* part of  $S^{\mathcal{F}}$ , however, is finite because it corresponds to root paths leading to different end states and/or pricing functions, where *prices decrease monotonically along any  $\Theta^{\mathcal{F}}$  path*.

The latter is because  $\text{prices}(s)[t^L]$  forms part of the minimization in (iii), namely when setting  $s^L := t^L$  and using the empty path of  $A^L|_{\text{rt}(t)}$  actions.<sup>1</sup>

Intuitively,  $\Theta^{\mathcal{F}}$  encapsulates the idea of searching only in the root part, keeping track of what the leaves *could* do via the pricing functions, and selecting the *actually used* leaf paths as needed once we can achieve the goal. Observe in particular that the cost  $\text{cost}(\pi^{\mathcal{F}})$  of a fork-decoupled plan  $\pi^{\mathcal{F}}$  accounts only for the root moves, not for the leaves. In that sense, fork-decoupled goal states  $s_G$  are *goal states with price tags*: We still have to “buy” the leaf goals, i. e., for each leaf factor  $F^L$ , we must pay  $\text{prices}(s_G)[s^L]$  for a cheapest goal leaf state  $s^L \in S^L|_{F^L}$ . We will get back to this below; for now let us illustrate Definition 3 with an example.

**Example 4** Consider again our illustrative example, with the fork factoring  $\{\{t_A, t_B\}, \{p\}\}$ , i. e., the fork root is  $F^R = \{t_A, t_B\}$  and the only leaf is  $F^L = \{p\}$

Consider Definition 3 (i). The fork-decoupled states  $S^{\mathcal{F}}$  here are assignments to  $\{t_A, t_B\}$  along with a price for every leaf state, i. e., for every value of the leaf variable  $p$ . For Definition 3 (ii), the transition labels are the truck-moving actions,  $\text{move}(x, y)$ .

For Definition 3 (iv), the fork-decoupled initial state  $I^{\mathcal{F}}$ , we clearly have  $\text{rt}(I^{\mathcal{F}}) = \{t_A = l_1, t_B = l_3\}$ . The pricing function  $\text{prices}(I^{\mathcal{F}})$  assigns 0 to  $p = l_1$ , assigns 1 to  $p = A$ , and assigns  $\infty$  to all other leaf states, because  $A^L|_{\text{rt}(I^{\mathcal{F}})}$  contains exactly the load/unload actions for  $t_A$  at  $l_1$ , and for  $t_B$  at  $l_3$ . Note here that the “price” of  $p = A$  is not a cost we have already spent. It is the cost we would have to spend if we decided to “buy”  $p = A$  in  $I^{\mathcal{F}}$ , i. e., if we decided to augment the empty root path with a compliant leaf path achieving  $p = A$ .

To illustrate Definition 3 (iii) and (v), let us apply some root actions and consider the corresponding transitions in  $\Theta^{\mathcal{F}}$ . Say we obtain  $s_1$  from  $I^{\mathcal{F}}$  by moving  $t_A$  to  $l_2$ . Then  $\text{rt}(s_1) = \{t_A = l_2, t_B = l_3\}$ . The price tags  $\text{prices}(s_1)$  are still 0 for  $p = l_1$  and 1 for  $p = A$ , as these were the prices in  $I^{\mathcal{F}}$ , no cheaper compliant paths become available in  $s_1$ , and  $\text{prices}(s)[t^L]$  forms part of the minimization in (iii) as pointed out above. In other words, the best option for these leaf states is to not move at  $s_1$ , instead using the compliant paths that were already available at  $I^{\mathcal{F}}$ . The only change in  $\text{prices}(s_1)$  is that  $p = l_2$  gets price 2, accounting for “unload( $p, t_A, l_2$ )” which is compliant at  $s_1$  (but not at  $I^{\mathcal{F}}$ ).

Now say we obtain  $s_2$  from  $s_1$  by moving  $t_B$  to  $l_2$ , and  $s_3$  from  $s_2$  by moving  $t_B$  back to  $l_3$ . Then  $s_3$  is a fork-decoupled goal state, because  $\text{prices}(s_2)[p = B] = 3$  and  $\text{prices}(s_3)[p = l_3] = 4$ . Tracing back the compliant leaf path supporting  $p = l_3$ , we get the following 7-step global plan: load  $p$  onto  $t_A$ , move  $t_A$  to  $l_2$ , unload  $p$ , move  $t_B$  to  $l_2$ , load  $p$  onto  $t_2$ , move  $t_B$  to  $l_3$ , unload  $p$ .

<sup>1</sup>The reader who is reminded, at the word “monotonic”, of delete-relaxed planning (McDermott 1999; Bonet and Geffner 2001), or of red-black planning (Domshlak, Hoffmann, and Katz 2015), is right on target. Indeed, thinking about delete-relaxed causal graph leaf variables in red-black planning is what lead us to the idea of fork-decoupled search in the first place.

Finally, say we obtain  $s_4$  by moving  $t_A$  to  $l_3$  in  $s_3$ . The compliant leaf path supporting  $p = l_3$  in  $s_3$  to yield the price tag 4 is then superseded by the new path “load( $p, t_A, l_1$ ), unload( $p, t_A, l_3$ )” yielding  $\text{prices}(s_4)[p = l_3] = 2$ . Observe that we now get a 6-step global plan, i. e., a plan better than that of the fork-decoupled goal state  $s_3$  we passed through on the way to  $s_4$ . This is because, as indicated, fork-decoupled goal states have price tags. The price tag of  $s_4$  is cheaper than that of  $s_3$ , hence continuing search behind the goal state  $s_3$  leads to a strictly better solution. We will see later how to obtain a standard structure for search.

We will see later how to get rid of goal state price tags to obtain a standard structure for search. For now, we establish correctness of the fork-decoupled state space relative to the original planning task. The central observation is:

**Lemma 1** Let  $\Pi$  be an FDR task, and  $\mathcal{F}$  a fork factoring. For every root path  $\pi^R$ , denoting by  $s$  the fork-decoupled state reached by  $\pi^R$  in  $\Theta^{\mathcal{F}}$ , and for every leaf state  $s^L$ ,  $\text{prices}(s)[s^L]$  is exactly the cost of a cheapest leaf path  $\pi^L$  that complies with  $\pi^R$  and achieves  $s^L$ .

**Proof:** Showing this by induction, the claim holds directly by construction, Definition 3 (iv), in  $I^{\mathcal{F}}$ . For the inductive step, we show that if the claim holds for a fork-decoupled state  $s$ , and  $(s \xrightarrow{a} t) \in T^{\mathcal{F}}$ , then the claim holds for  $t$  as well. Let  $\pi^R = \langle a_1^R, \dots, a_n^R \rangle$  be any root path. Denote  $\pi^R$  itself as  $\pi^R[t]$ , and denote the prefix of  $\pi^R$  up to  $a_{n-1}^R$  as  $\pi^R[s]$ . Let  $s$  and  $t$  be the fork-decoupled states reached by  $\pi^R[s]$  respectively  $\pi^R[t]$ . Say  $\pi^L[t]$  is any  $\pi^R[t]$ -compliant leaf path. Let  $\pi^L[s]$  be the longest  $\pi^R[s]$ -compliant prefix of  $\pi^L[t]$ . Then the postfix of  $\pi^L[t]$  after  $\pi^L[s]$ , if non-empty, must consist of actions in  $A^L|_{\text{rt}(t)}$ . In other words, the compliant leaf paths for  $t$  are (possibly empty) extensions of those for  $s$ . If  $\pi^L[t]$  is a cheapest  $\pi^R[t]$ -compliant leaf path ending in  $t^L$ , then clearly  $\pi^L[s]$  is a cheapest  $\pi^R[s]$ -compliant leaf path to its end point  $s^L$ , and the postfix of  $\pi^L[t]$  behind  $\pi^L[s]$  is a cheapest path of  $A^L|_{\text{rt}(t)}$  actions from  $s^L$  to  $t^L$ . Hence the claim for  $t$  follows directly from the induction assumption on  $s$  –  $\text{prices}(s)[s^L]$  equals the cost of  $\pi^L[s]$  – and construction, Definition 3 (iii).  $\square$

To formalize the “goal state price tag”, recall that each leaf factor may be the cross-product of several state variables, and hence may have more than one goal leaf state. So we need to minimize over these leaf goal states:

**Definition 4 (Goal Price)** Let  $\Pi$  be an FDR task,  $\mathcal{F}$  a fork factoring, and  $\Theta^{\mathcal{F}}$  the fork-decoupled state space.

The goal price of  $s_G \in S_G^{\mathcal{F}}$  is the sum over minimal leaf goal state prices, i. e.,  $\text{Gprice}(s_G) := \sum_{F^L \in \mathcal{F}^L} \min\{\text{prices}(s_G)[s^L] \mid s^L \in S^L|_{F^L}, s^L \supseteq G[F^L]\}$ .

The global cost of a  $\Theta^{\mathcal{F}}$  path  $\pi^{\mathcal{F}}$  ending in  $s_G \in S_G^{\mathcal{F}}$  is  $\text{GlobalCost}(\pi^{\mathcal{F}}) := \text{cost}(\pi^{\mathcal{F}}) + \text{Gprice}(s_G)$ . A fork-decoupled plan is fork-optimal if its global cost is minimal among all fork-decoupled plans.

For a fork-decoupled plan  $\pi^{\mathcal{F}}$  ending in  $s_G \in S_G^{\mathcal{F}}$ , define  $\text{GlobalPlan}(\pi^{\mathcal{F}})$  as  $\pi^{\mathcal{F}}$  augmented, for every  $F^L \in \mathcal{F}^L$ , with a cheapest compliant leaf path  $\pi^L$  achieving  $\arg \min\{\text{prices}(s_G)[s^L] \mid s^L \in S^L|_{F^L}, s^L \supseteq G[F^L]\}$ . Vice

versa, given a plan  $\pi$  for  $\Pi$ , let  $\text{ForkPlan}(\pi)$  be the  $\Theta^{\mathcal{F}}$  path  $\pi^{\mathcal{F}}$  induced by the root actions in  $\pi$ . We have:

**Theorem 1** *Let  $\Pi$  be an FDR task, and  $\mathcal{F}$  a fork factoring. For any (fork-optimal) fork-decoupled plan  $\pi^{\mathcal{F}}$  for  $\Pi$ ,  $\text{GlobalPlan}(\pi^{\mathcal{F}})$  is an (optimal) plan for  $\Pi$ . Vice versa, for any (optimal) plan  $\pi$  for  $\Pi$ ,  $\text{ForkPlan}(\pi)$  is a (fork-optimal) fork-decoupled plan for  $\Pi$ .*

**Proof:** Because  $\mathcal{F}$  is a fork factoring, (\*) any plan for  $\Pi$  can be perceived as a root path  $\pi^R$  achieving the goal for  $F^R$ , augmented for every leaf  $F^L$  with a compliant leaf path  $\pi^L$  achieving  $G[F^L]$ .

Let  $\pi^{\mathcal{F}}$  be a fork-decoupled plan. By construction and by (\*),  $\text{GlobalPlan}(\pi^{\mathcal{F}})$  is a plan for  $\Pi$ . By Lemma 1, the global cost of  $\pi^{\mathcal{F}}$  is equal to the cost of  $\text{GlobalPlan}(\pi^{\mathcal{F}})$ .

Vice versa, let  $\pi$  be a plan. By construction and by (\*),  $\text{ForkPlan}(\pi)$  is a fork-decoupled plan for  $\Pi$ . Assuming now that  $\pi$  is optimal, clearly its compliant leaf paths as per  $\pi^L$  are cheapest compliant leaf paths. Hence, by Lemma 1, the global cost of  $\text{ForkPlan}(\pi)$  is equal to the cost of  $\pi$ .

We now know, in particular, that fork-optimal fork-decoupled plans correspond to plans for  $\Pi$  of the same cost, and that optimal plans for  $\Pi$  correspond to fork-decoupled plans of the same cost. Hence optimal cost is the same on both sides, which concludes our argument.  $\square$

In other words, (fork-optimal) search in  $\Theta^{\mathcal{F}}$  is a form of (optimal) planning for  $\Pi$ .

What can we say about the *size* of  $\Theta^{\mathcal{F}}$ ? It captures a search over root *paths*, as opposed to root *states*, because even if two paths lead to the same root state, they may lead to different leaf prices. So what is the relation between the reachable states in  $\Theta^{\mathcal{F}}$  (denoted  $S_R^{\mathcal{F}}$ ) and the reachable states in the standard state space (denoted  $S_R$ )? A fork-decoupled state  $s \in S_R^{\mathcal{F}}$  can be viewed as a “condensed” set of (priced) world states  $w \in S_R$ , namely all those  $w$  consisting of the same root state and, for each leaf factor, any leaf state  $s^L$  reachable in  $s$  i. e. where  $\text{prices}(s)[s^L] < \infty$ . Condensing world states in this way can dramatically reduce state space size.

A pathological case arises when using the pathological fork factoring (hence the name), where the entire task is a single “leaf”. Then, obviously,  $|S_R^{\mathcal{F}}| = 1$ , but that does not buy us anything because the pricing function in the fork-decoupled initial state enumerates the cheapest distances to all states in  $S_R$ .

The genuinely good cases are those where the decoupling exponentially reduces the state space, but each fork-decoupled state has polynomial size. Here is one such case:

**Example 5** *Consider the variant of our illustrative example with only the single truck  $t_A$ , but  $n$  packages  $p_1, \dots, p_n$  and a line  $l_1, \dots, l_m$  of locations. Initially, both truck and packages are at  $l_1$ . Say the fork factoring is that where the root is  $t_A$  and each  $p_i$  is a leaf. Obviously,  $|S_R| = m * (m + 1)^n$ . However, each fork-decoupled state has polynomial size  $O(n * m)$ , and  $|S_R^{\mathcal{F}}| = \frac{m(m+1)}{2}$ . The number of (reachable) fork-decoupled states does not even depend on the number of packages!*

*To see this, consider how fork-decoupled states are reached from the initial state. In the initial fork-decoupled state, all packages may be loaded so each has the reached value set  $\{l_1, A\}$ . That is, the pricing function assigns finite prices to the leaf states (values of each  $p_i$ )  $l_1$  and  $A$ , and assigns price  $\infty$  to all other leaf states. We will ignore the actual prices here, distinguishing only what is reached, as it is easy to see that the price for the same leaf will always be the same anyhow.*

*After moving to  $l_2$ , the reached value sets are  $\{l_1, A, l_2\}$ . Moving back to  $l_1$ , we obtain the new fork-decoupled state with reached value sets  $\{l_1, A, l_2\}$  and where the truck is at  $l_1$ . Similarly, for each location  $l_i$  where  $i > 2$ , the new fork-decoupled states are those where the truck reaches  $l_i$  itself for the first time, yielding the new reached value sets  $\{l_1, A, l_2, \dots, l_i\}$ , plus the  $i - 1$  fork-decoupled states reached by going back to  $l_{i-1}, \dots, l_1$  from there. Hence, overall we get  $\sum_{i=1}^m i = \frac{m(m+1)}{2}$  reachable fork-decoupled states as advertised.*

That said, the condensed state sets may overlap for distinct fork-decoupled states. One can exploit this to construct bad cases, in which  $S_R^{\mathcal{F}}$  enumerates combinations of leaf states for the same leaf factor, and  $S_R^{\mathcal{F}}$  is exponentially larger than  $S_R$ . Here is one such case:

**Example 6** *Consider an FDR task with two variables,  $r$  and  $l$ . Our fork factoring will have  $r$  as the root and  $l$  as the leaf. The domain of  $r$  is  $\{d_1, d_1^a, d_1^b, d_2, d_2^a, d_2^b, \dots, d_n, d_n^a, d_n^b\}$ , arranged such that we can move from each  $d_i$  to each  $d_i^j$ , and from each  $d_i^j$  where  $i < n$  to  $d_{i+1}$ . The initial value is  $d_1$ , and there is no goal on  $r$ . The domain of  $l$  is  $e_1, e_1^a, e_1^b, e_2, e_2^a, e_2^b, \dots, e_n, e_n^a, e_n^b, e_{n+1}$ , arranged such that, with precondition  $d_i^j$ , we can move from  $e_i$  to  $e_i^j$  and from  $e_i^j$  to  $e_{i+1}$ . The initial value is  $e_1$ , and the goal value is  $e_{n+1}$ .*

*In words, the domain of each variable forms a chain of alternative branches  $a$  and  $b$ , where each such branch is followed by a common culmination point. At each step along its chain, the leaf variable has two alternative options to move forward. The fork-decoupled state space will “remember” which option was taken, in the sense that exactly one of the respective options –  $e_i^a$  or  $e_i^b$  – will have a finite price in all subsequent fork-decoupled states. Hence the fork-decoupled state space grows exponentially in  $n$ . In the standard state space, on the other hand, once the leaf reached the common culmination point  $e_{i+1}$ , it will not matter whether we got there via  $e_i^a$  or  $e_i^b$ . Precisely, the standard state space has polynomial size simply because it is bounded by the product of variable domain sizes,  $O(n^2)$ .*

Note that this example is contrived. On IPC benchmarks, with our current factoring strategy,  $S_R^{\mathcal{F}}$  is never larger than  $S_R$ , and is typically *much* smaller. The curious reader is invited to skip forward to Table 1 for a moment.

## Heuristic Functions

In fork-decoupled search, two *different* kinds of heuristic functions are of interest:

**Definition 5 (Heuristics)** Let  $\Pi$  be an FDR task, and  $\mathcal{F}$  a fork factoring. A heuristic is a function  $h : S^{\mathcal{F}} \mapsto \mathbb{R}^{0+} \cup \{\infty\}$ .

The root-perfect heuristic,  $h^{R*}$ , is that where  $h^{R*}(s)$  for any  $s$  is the minimum over  $\text{cost}(\pi^{\mathcal{F}})$  for all  $\Theta^{\mathcal{F}}$  paths  $\pi^{\mathcal{F}}$  from  $s$  to some  $s_G \in S_G^{\mathcal{F}}$ . The fork-perfect heuristic,  $h^{\mathcal{F}*}$ , is that where  $h^{\mathcal{F}*}(s)$  for any  $s$  is the minimum over  $\text{GlobalCost}(\pi^{\mathcal{F}})$  for all such paths.

We say that  $h$  is root-admissible if  $h \leq h^{R*}$ , and fork-admissible if  $h \leq h^{\mathcal{F}*}$ .

The conceptual distinction between  $h^{R*}$  and  $h^{\mathcal{F}*}$  lies in that  $h^{R*}$  cares only about how much work is left for the root factor, i. e., the cost of a root path sufficient to enable every leaf to reach its goal *somehow*. In contrast,  $h^{\mathcal{F}*}$  accounts for the combined cost of root and leaves, i. e. for the best extension of our current root path into an overall fork-decoupled plan (and thus a plan for the original input task  $\Pi$ , cf. Theorem 1). We will refer to heuristics attempting to estimate  $h^{R*}$  as *root heuristics*, and to heuristics attempting to estimate  $h^{\mathcal{F}*}$  as *fork heuristics*, and distinguish them notationally by superscripts “ $R$ ” respectively “ $\mathcal{F}$ ”.

Observe that  $h^{\mathcal{F}*}$  is notably different from typical remaining cost estimators (including, e. g.,  $h^{R*}$ ) due to the goal state price tags, captured in Definition 5 by the goal price component of  $\text{GlobalCost}(\pi^{\mathcal{F}})$ . These price tags account for the entire cost of moving the leaves, including actions that will be scheduled *before* (and up to) the state  $s$  in question. In particular,  $h^{\mathcal{F}*}$  is *not* 0 on fork-decoupled goal states: we still have to pay the goal prize, moving the leaves into place.

Observe furthermore that  $h^{R*}$  is a special case of  $h^{\mathcal{F}*}$ : We can compute  $h^{R*}$  as  $h^{\mathcal{F}*}$  in a modified planning task where the cost of all leaf actions is set to 0.  $h^{R*}$  keeps track only of which leaf states are reachable, not of the associated cost. This may lead to qualitatively different decisions, i. e.,  $h^{R*}$  and  $h^{\mathcal{F}*}$  may *disagree*. Using a transportation example again, say that there are two alternative kinds of plans, (a) ones that pass the packages through several trucks, loading/unloading every time, vs. (b) ones that make more truck moves but have to load/unload each package only once and thus are better globally. Then  $h^{R*}$  will draw search towards plans (a), whereas  $h^{\mathcal{F}*}$  will draw search towards plans (b).

Turning our attention to the computation of practical heuristics, we consider  $h^{\mathcal{F}*}$  first. Note that, to estimate  $h^{\mathcal{F}*}$ , we have to account for the option to schedule leaf actions *before* (and up to) the state  $s$  in question. This is both, a threat to informativity (we cannot simply discount the cost of such actions) and to admissibility (we cannot simply pretend that this is not possible). It turns out that both can be tackled using the pricing information maintained in  $\Theta^{\mathcal{F}}$ :

**Definition 6 (From Classical  $h$  to  $h^{\mathcal{F}*}$  Estimation)** Let  $h$  be any FDR planning heuristic,  $\Pi = \langle V, A, I, G \rangle$  an FDR task,  $\mathcal{F}$  a fork factoring, and  $s$  a fork-decoupled state. The leaves-for-pay fork heuristic estimate is  $h_{L\mathcal{S}}^{\mathcal{F}*}(s) := h(s_{L\mathcal{S}})$  using the FDR task  $\Pi_{L\mathcal{S}} = \langle V_{L\mathcal{S}}, A_{L\mathcal{S}}, s_{L\mathcal{S}}, G \rangle$  as follows. We set  $V_{L\mathcal{S}} := V \cup \{ch[F^L] \mid F^L \in \mathcal{F}^L\}$  where each  $ch[F^L]$  is Boolean, and for each  $v \notin F^C$ , we add the new value  $u$  into  $\mathcal{D}(v)$ . The initial state is  $s_{L\mathcal{S}} := \text{rt}(s) \cup \{v = u \mid v \notin F^C\} \cup \{ch[F^L] = 0 \mid F^L \in \mathcal{F}^L\}$ . The actions  $A_{L\mathcal{S}}$  are the

previous ones, plus, for every  $F^L \in \mathcal{F}^L$  and  $s^L \in S^L|_{F^L}$  where  $\text{prices}(s)[s^L] < \infty$ , a new action  $a[s^L]$  with precondition  $\{ch[F^L] = 0\}$ , effect  $\text{eff}(a[s^L]) := s^L \cup \{ch[F^L] = 1\}$ , and cost  $\text{cost}(a[s^L]) := \text{prices}(s)[s^L]$ .

This construction caters for “leaf actions to be scheduled before and up to  $s$ ”, but not by allowing to explicitly insert such actions. We instead allow the planner on  $\Pi_{L\mathcal{S}}$  to pay the prize that *would* be incurred by doing so, using the new actions  $a[s^L]$ .

**Example 7** Consider the state  $s_1$  from Example 2, where we moved  $t_A$  to  $l_2$  and have  $\text{rt}(s_1) = \{t_A = l_2, t_B = l_3\}$  as well as finite-value price tags  $\text{prices}(s_1)[p = l_1] = 0$ ,  $\text{prices}(s_1)[p = A] = 1$ , and  $\text{prices}(s_1)[p = l_2] = 2$ . The task  $\Pi_{L\mathcal{S}}$  from Definition 6 then takes the following form. The variables and goal are  $\{p, t_A, t_B, ch[p]\}$  respectively  $\{p = l_3\}$ . The new initial state  $s_{L\mathcal{S}}$  uses the root state of  $s_1$ ,  $\{t_A = l_2, t_B = l_3\}$ , uses the new “ $u$ ” value of the leaf variable,  $p = u$ , and it sets  $ch[p] = 0$ . We have the same move/load/unload actions as before, plus the new leaf actions  $a[s^L]$  with precondition  $ch[p] = 0$ , effects  $s^L$  as well as  $ch[p] = 1$ , and cost  $\text{cost}(a[s^L]) := \text{prices}(s)[s^L]$ . The latter actions are present for each  $s^L$  with a finite-value price tag in  $s_1$ , i. e., for  $p = l_1$  (price 0), for  $p = A$  (price 1, and for  $p = l_2$  (price 2).

Clearly, the only optimal plan for  $\Pi_{L\mathcal{S}}$  is to use  $a[p = A]$  to get the package into  $t_A$ , then move  $t_A$  over to  $l_3$  and unload the package. So  $h^*(s_{L\mathcal{S}})$  in  $\Pi_{L\mathcal{S}}$  equals 3, which is the same as  $h^{\mathcal{F}*}(s_1)$ . Note that existing heuristics like  $h^{\text{FF}}$  (Hoffmann and Nebel 2001) and LM-cut (Helmert and Domshlak 2009) would return this same heuristic value.

It was no coincidence that  $h^*$  in this example resulted in the fork-perfect heuristic. This is true in general, and furthermore the leaves-for-pay fork heuristic preserves admissibility:

**Theorem 2** If  $h$  is admissible, then  $h_{L\mathcal{S}}^{\mathcal{F}}$  is fork-admissible. If  $h = h^*$ , then  $h_{L\mathcal{S}}^{\mathcal{F}} = h^{\mathcal{F}*}$ .

**Proof:** Let  $s \in S^{\mathcal{F}}$ . Say  $\pi^{\mathcal{F}}$  is any  $\Theta^{\mathcal{F}}$  path from  $s$  to some  $s_G \in S_G^{\mathcal{F}}$ . We construct a plan  $\pi$  for  $\Pi_{L\mathcal{S}}$  by augmenting  $\pi^{\mathcal{F}}$  with cheapest compliant leaf paths achieving the leaf goals when starting from the leaf states reached in  $s$ . By construction, the cost of  $\pi$  in  $\Pi_{L\mathcal{S}}$  is equal to  $\text{cost}(\pi^{\mathcal{F}}) + \text{Gprice}(s_G) = \text{GlobalCost}(\pi^{\mathcal{F}})$ . Hence  $h_{L\mathcal{S}}^{\mathcal{F}}(s) \geq h^*(s_{L\mathcal{S}})$ .

Vice versa, let  $\pi$  be any plan for  $\Pi_{L\mathcal{S}}$ . Then the root action sub-sequence  $\pi^R$  of  $\pi$  must be augmentable with compliant leaf paths achieving the leaf goals when starting from a leaf state  $s^L$  reached in  $s$ . So  $\pi^R$  induces a  $\Theta^{\mathcal{F}}$  path  $\pi^{\mathcal{F}}$  from  $s$  to some  $s_G \in S_G^{\mathcal{F}}$ . Further, as the construction of  $\Pi_{L\mathcal{S}}$  forces the plan to buy, for each leaf factor with a goal, exactly one such  $s^L$ ,  $\pi$  must for each  $F^L$  contain a  $\pi^R$ -compliant leaf path  $\pi^L$  from such an  $s^L$  to its goal. As  $\pi$  is optimal, the  $\pi^L$  are cheapest such paths. So the cost of  $\pi$  is  $\text{cost}(\pi^{\mathcal{F}}) + \text{Gprice}(s_G) = \text{GlobalCost}(\pi^{\mathcal{F}})$ , and hence  $h^*(s_{L\mathcal{S}}) \geq h^{\mathcal{F}*}(s)$ . We get  $h^*(s_{L\mathcal{S}}) = h^{\mathcal{F}*}(s)$ , from which the claims follow.  $\square$

We now get an  $h^{R*}$  estimate as a corollary via the aforementioned relation to  $h^{\mathcal{F}*}$ :

**Definition 7 (From Classical  $h$  to  $h^{R^*}$  Estimation)** Let  $h$  be any FDR planning heuristic,  $\Pi = \langle V, A, I, G \rangle$  an FDR task,  $\mathcal{F}$  a fork factoring, and  $s$  a fork-decoupled state. The leaves-for-free root heuristic estimate is  $h_{L0\$}^R(s) := h(s_{L0\$})$  in the FDR task  $\Pi_{L0\$} = \langle V_{L0\$}, A_{L0\$}, s_{L0\$}, G \rangle$  identical to  $\Pi_{L\$}$  as in Definition 6, except that we set  $\text{cost}(a^L) := 0$  for all leaf actions  $a^L \in A^L$ , and define  $\text{cost}(a[s^L]) := 0$ .

**Example 8** Reconsider Example 7. In difference to  $\Pi_{L\$}$  as constructed there, in  $\Pi_{L0\$}$  all leaf actions, including the new ones  $a[p = l_1]$ ,  $a[p = A]$ , and  $a[p = l_1]$ , have 0 cost. Other than that, the two tasks are identical. Clearly, the only optimal plan still is to use  $a[p = A]$  to get the package into  $t_A$ , then move  $t_A$  over to  $l_2$  and unload the package. Only now the cost of that plan is 1, which is exactly  $h^{R^*}(s_1)$ .

**Theorem 3** If  $h$  is admissible, then  $h_{L0\$}^R$  is root-admissible. If  $h = h^*$ , then  $h_{L0\$}^R = h^{R^*}$ .

**Proof:** Let  $s \in S^{\mathcal{F}}$ , reached from  $I^{\mathcal{F}}$  in  $\Theta^{\mathcal{F}}$  using the root action sequence  $\pi^R$ . Consider the state  $s_0$  that would be reached by  $\pi^R$  in the modification  $\Pi_0$  of  $\Pi$  where all leaf action costs are set to 0. Then  $h^{R^*}[\Pi](s) = h^{\mathcal{F}^*}[\Pi_0](s_0)$ . Applying Definition 6 to  $\Pi_0$  and  $s_0$ , we get  $\Pi_{L\$} = \Pi_{L0\$}$  and  $s_{L\$} = s_{L0\$}$ , so that as argued in the proof of Theorem 2 we have  $h^{\mathcal{F}^*}[\Pi_0](s_0) = h^*[\Pi_{L0\$}](s_{L0\$})$ . Overall,  $h^*(s_{L0\$}) = h^{R^*}(s)$ , from which the claims follow.  $\square$

## Search Algorithms

As usual, we say that a search algorithm is *complete* if it returns a solution whenever one exists. We say that a search algorithm is *optimal* if, whenever it returns a solution, its cost is minimal among all solutions. Disregarding optimality, we can run any search algorithm on the fork-decoupled state space, stopping at the first fork-decoupled goal state. For optimal planning, matters are more subtle. One of our methods is formulated on a modified state space where the goal price tags are explicit:

**Definition 8 (Explicit Goal Prices)** Let  $\Pi$  be an FDR task,  $\mathcal{F}$  a fork factoring, and  $\Theta^{\mathcal{F}} = \langle S^{\mathcal{F}}, A^R, T^{\mathcal{F}}, I^{\mathcal{F}}, S_G^{\mathcal{F}} \rangle$  the fork-decoupled state space. The explicit goal-price state space  $\Theta_{G\$}^{\mathcal{F}}$  is like  $\Theta^{\mathcal{F}}$  but with a new state  $G'$  and, for every  $s_G \in S_G^{\mathcal{F}}$ , a new transition  $s_G \rightarrow G'$  with cost  $\text{Gprice}(s_G)$ ; the set of goal states is set to be  $\{G'\}$ .

The explicit goal-price state space is useful because it allows us to formulate search on  $\Theta^{\mathcal{F}}$ , which is non-standard due to the goal state price tags, in terms of standard search on  $\Theta_{G\$}^{\mathcal{F}}$ , via this correspondence:

**Lemma 2** The solution paths in  $\Theta_{G\$}^{\mathcal{F}}$  are in one-to-one correspondence with the fork-decoupled plans for  $\Pi$ .

**Proof:** This is direct from construction. Every solution path  $\pi$  for  $\Theta_{G\$}^{\mathcal{F}}$  must have the form  $\pi^{\mathcal{F}} \circ \langle s_G \rightarrow G' \rangle$  where  $s_G \in S_G^{\mathcal{F}}$ , the weight of  $s_G \rightarrow G'$  is  $\text{Gprice}(s_G)$ , and  $\pi^{\mathcal{F}}$  is a fork-decoupled plan. Clearly,  $\text{cost}(\pi^{\mathcal{F}}) + \text{Gprice}(s_G) = \text{GlobalCost}(\pi^{\mathcal{F}})$  equals the cost of  $\pi$ . Similar vice versa.  $\square$

Consider Figure 4.  $\text{FDX}$  just runs any search algorithm  $X$  on  $\Theta_{G\$}^{\mathcal{F}}$ . We have:

**Algorithm Fork-Decoupled  $X$  (FDX):**  
**Input:** FDR planning task  $\Pi$ , fork factoring  $\mathcal{F}$   
 Heuristic search algorithm  $X$   
 Fork heuristic  $h^{\mathcal{F}}$   
**Output:** A plan for  $\Pi$ , or “failed”  
 Let  $h_{G\$}^{\mathcal{F}} := \begin{cases} h^{\mathcal{F}}(s) & s \in S^{\mathcal{F}} \\ 0 & s = G' \end{cases}$   
 Run  $X$  with  $h_{G\$}^{\mathcal{F}}$  on  $\Theta_{G\$}^{\mathcal{F}}$   
 If  $X$  found a solution path  $\pi = \pi^{\mathcal{F}} \circ \langle s_G \rightarrow G' \rangle$   
   return  $\text{GlobalPlan}(\pi^{\mathcal{F}})$   
 else return “failed”

Figure 4: Exploiting any known search algorithm  $X$ .

**Theorem 4** If  $X$  is complete, then  $\text{FDX}$  is complete. If  $X$  is optimal for admissible heuristics, then  $\text{FDX}$  is optimal for fork-admissible heuristics.

**Proof:** Completeness follows directly from Lemma 2 and Theorem 1. Optimality follows in the same manner, as, for fork-admissible  $h^{\mathcal{F}}$ ,  $h_{G\$}^{\mathcal{F}}$  is admissible in  $\Theta_{G\$}^{\mathcal{F}}$ .  $\square$

Consider now Figure 5.  $\text{AFRA}^*$  guides  $A^*$  by a *root heuristic*, and uses a fork heuristic merely for upper-bound pruning. The search is drawn to cheap root paths, disregarding leaf costs, so to guarantee optimality we must exhaust the open list. Without early termination, this would be dominated by  $\text{FDA}^*$  because  $\text{AFRA}^*$  would then have to expand at least all  $N[s]$  where  $g(N) + h^{\mathcal{F}}(s)$  is less than optimal solution cost. *With* early termination, that is not so because in the best case we have to exhaust only those  $N[s]$  where  $g(N) + h^R(s)$  is less than optimal *root* solution cost.

**Theorem 5**  $\text{AFRA}^*$  is complete, and is optimal for root-admissible  $h^R$  and fork-admissible  $h^{\mathcal{F}}$ .  $\text{FRA}^*$  is optimal for root-admissible  $h^R$ .

**Proof:** Completeness of  $\text{AFRA}^*$  is obvious with Theorem 1. Towards proving optimality, observe: (1) Without early termination and upper-bound pruning,  $\text{AFRA}^*$  generates (in particular) every fork-optimal fork-decoupled plan  $\pi^{\mathcal{F}}$ . (2) With root-admissible  $h^R$ , if  $\text{AFRA}^*$  generates  $\pi_1^{\mathcal{F}}$  before it generates  $\pi_2^{\mathcal{F}}$ , then  $\text{cost}(\pi_1^{\mathcal{F}}) \leq \text{cost}(\pi_2^{\mathcal{F}})$ . (3) Upper-bound pruning trivially preserves (2), and it preserves (1) with fork-admissible  $h^{\mathcal{F}}$  because the pruned nodes  $N$  cannot lead to fork-optimal solutions. (4) If early termination fires for  $N[s_G]$ , then  $\text{Gprice}(s_G)$  is minimal in  $S_G^{\mathcal{F}}$ .

Optimality of  $\text{AFRA}^*$  without early termination holds by (1) and (3) with Theorem 1. Say that early termination fires for  $N[s_G]$ , let  $\pi_1^{\mathcal{F}}$  be the fork-decoupled plan leading to  $N$ , and let  $\pi_2^{\mathcal{F}}$  be any fork-decoupled plan that would be generated later if we continued the search. By (4)  $\pi_2^{\mathcal{F}}$  pays at least as high a goal price as  $\pi_1^{\mathcal{F}}$ , and by (2)  $\text{cost}(\pi_1^{\mathcal{F}}) \leq \text{cost}(\pi_2^{\mathcal{F}})$ , so  $\pi_2^{\mathcal{F}}$  cannot have better global cost than  $\pi_1^{\mathcal{F}}$ . With (1),  $\text{AFRA}^*$  already found a fork-optimal plan, which must be stored in  $\pi_U^{\mathcal{F}}$ .

To see optimality of  $\text{FRA}^*$ , just note that, by (2), the first  $\pi^{\mathcal{F}}$  generated by  $\text{AFRA}^*$  has minimal  $\text{cost}(\pi^{\mathcal{F}})$  among all fork-decoupled plans. With (4) and Theorem 1, the claim follows, concluding the proof.  $\square$

$\text{FRA}^*$  is not complete because we force it to prove optimality. In practice, one could return the plan anyway, flagging it as “proved optimal” only if  $\text{Gprice}(s_G) =$

**Algorithm** *Anytime Fork-Root A\* (AFRA\*)*:  
**Input:** FDR planning task  $\Pi$ , fork factoring  $\mathcal{F}$   
 Root heuristic  $h^R$ , fork heuristic  $h^F$   
**Output:** An optimal plan for  $\Pi$ , or “unsolvable”  
 Let  $U := \infty$  /\* best known upper bound \*/  
 Let  $\pi_U^F := \perp$  /\* corresponding plan \*/  
 Run  $A^*$  with  $h^R$  on  $\Theta^F$ , with these modifications:  
 Continue search until the open list is empty  
 Whenever a goal vertex node  $N[s_G]$  is expanded:  
 If  $g(N) + \text{Gprice}(s_G) < U$   
 let  $U := g(N) + \text{Gprice}(s_G)$   
 let  $\pi_U^F :=$  the fork-decoupled plan leading to  $N$   
 If  $\text{Gprice}(s_G) = \text{MINGprice}$   
 return  $\text{GlobalPlan}(\pi_U^F)$  /\* early termination \*/  
 Whenever a node  $N[s]$  is generated, and  $U \neq \infty$ :  
 If  $g(N) + h^F(s) \geq U$   
 discard  $N$  /\* upper-bound pruning \*/  
 If  $\pi_U^F \neq \perp$  return  $\text{GlobalPlan}(\pi_U^F)$  else return “unsolvable”  
**Algorithm** *Fork-Root A\* (FRA\*)*:  
**Input:** FDR planning task  $\Pi$ , fork factoring  $\mathcal{F}$   
 Root heuristic  $h^R$   
**Output:** An optimal plan for  $\Pi$ , or “sorry no luck”  
 Run  $AFRA^*$  with  $h^R$  on  $\Theta^F$ , with these modifications:  
 Do not use upper-bound pruning  
 When the first goal vertex node  $N[s_G]$  is expanded, do:  
 If  $\text{Gprice}(s_G) = \text{MINGprice}$   
 let  $\pi^F :=$  the fork-decoupled plan leading to  $N$   
 return  $\text{GlobalPlan}(\pi^F)$   
 else return “sorry no luck”

Figure 5: New search algorithms. Search nodes are notated  $N[s]$  where  $s$  is the state and  $N$  the node itself.  $\text{MINGprice}$  is the sum, over the leaf factors  $F^L \in \mathcal{F}^L$ , of optimal plan cost for the *projection* of  $\Pi$  onto  $F^L$ .

$\text{MINGprice}$ . From a theory perspective, being *optimal but not complete* is a rather unique profile. It can be viewed as an effect of target-profile factoring, which is intrinsically geared at special-case structure (in our case, the leaf costs).

## Experiments

Our implementation is in FD (Helmert 2006). The pricing function is maintained enumeratively for each leaf factor. For plan extraction, each leaf state stores a backward pointer to its predecessor on a cheapest compliant path. We employ *subsumption pruning*, where  $s$  subsumes  $t$  if  $\text{rt}(s) = \text{rt}(t)$  and, for every leaf state  $s^L$ ,  $\text{prices}(s)[s^L] \leq \text{prices}(t)[s^L]$ . We implemented fork and root heuristics via Definitions 6 and 7 for  $h^{\max}$ , LM-cut, and  $h^{\text{FF}}$  (Bonet and Geffner 2001; Helmert and Domshlak 2009; Hoffmann and Nebel 2001). We also experiment with FD’s “blind heuristic”, returning 0 on goal states and  $\min_{a \in A} \text{cost}(a)$  elsewhere.

We ran all IPC STRIPS benchmarks (1998 – 2014), on a cluster of Intel E5-2660 machines running at 2.20 GHz, with time (memory) cut-offs of 30 minutes (4 GB).

## Factoring Strategy

We proceed as follows. If the input task has causal graph leaf variables, define  $\mathcal{F}_0$  as the fork factoring whose leaf factors are exactly those variables; else  $\mathcal{F}_0$  is undefined. Initialize  $D := \text{IG}(\mathcal{F}^{\text{SCC}})$  to be the DAG interaction graph over the

	#	State spaces				Leaf factors (Mean stats)					
		Built		Mean size (Common)		Common		All			
		Std	Fork	Std	$ S_R $	Fork	$ S_R^F $	Nr	MaxSize	Nr	MaxSize
Driverlog	20	5	8	3005640.0	27385.2	3.4	5.0	8.4	12.2		
Logistics	63	12	23	422507.6	2614.2	4.9	7.5	11.2	71.2		
Miconic	145	45	30	48544.0	6572.3	4.5	4.0	16.0	4.0		
NoMystery	40	11	26	3561168.0	1290.4	5.0	6.7	9.0	10.5		
Pathways	29	3	3	1167551.7	236471.7	3.0	2.0	21.1	2.0		
Rovers	40	5	5	3925444.8	138813.4	3.8	2.0	20.9	2.0		
Satellite	36	4	4	170808.0	109268.8	4.8	2.0	54.8	2.0		
TPP	24	5	11	10029509.4	21.6	4.0	7.8	8.3	954.4		
Woodwork	61	9	9	8204333.2	122433.9	2.8	160.1	3.9	43.6		
Zenotravel	20	7	7	1965822.0	41887.1	3.7	5.1	10.0	11.7		

Table 1: Reachable state space size and leaf factor statistics, on IPC instances where our factoring method does not abstain (see text). Where there are several IPC instance suites of a domain, we take their union here. “Built”: reachable state space could be fully built; “Common”: those instances where both state spaces were built; “Nr”: (mean) number of leaf factors; “MaxSize”: (mean) maximum leaf factor size, where the “size” of  $F^L \in \mathcal{F}^L$  is the number of reachable states in the projection of  $\Pi$  onto  $F^L$ .

input task’s SCCs. Set  $B_0 := \emptyset$  and  $i := 1$ . Then iterate the following steps: Take the *candidates* to be those leaf SCCs  $L$  in  $D$  where, with  $B'_L := B_{i-1} \cup L$  and  $\mathcal{W}'_L$  being the weakly connected components within  $B'_L$ , all  $W' \in \mathcal{W}'_L$  satisfy  $\prod_{v \in W'} |\mathcal{D}(v)| \leq 2^{32}$ . If there are no candidates, terminate the iteration. Else, select a candidate  $L$  with minimal  $\prod_{v \in L} |\mathcal{D}(v)|$ . Set  $B_i := B'_L$ , define  $\mathcal{F}_i$  by  $\mathcal{F}_i^R := V \setminus B_i$  and  $\mathcal{F}_i^L := \mathcal{W}'_L$ , remove  $L$  from  $D$ , increment  $i$ , and iterate. Upon termination, consider the set of defined  $\mathcal{F}_i$ . If this set is empty, or no factoring in this set has more than a single leaf, then we *abstain* from solving  $\Pi$ . Otherwise, we select the factoring  $\mathcal{F}_i$  whose number of leaf factors is maximal, and whose index  $i$  is minimal among these factorings.

In words, we run a greedy pass through the lattice of fork factorings, imposing for sanity that no leaf factor has estimated size greater than  $\text{MAXINT}$  (for 32 bits). We select the factoring with the maximal number of leaves, keeping these as small as possible. This maximizes the amount of conditional independence, while minimizing the overhead of maintaining pricing functions. It takes negligible runtime effort as the graph  $\text{IG}(\mathcal{F}^{\text{SCC}})$  is reliably very small. The runtime is (rounded to) 0.00 in 97% of all IPC STRIPS tasks, and is  $\leq 0.1$  seconds in all but a single task, where we take 0.19 seconds but FD’s pre-process takes several minutes.

We also ran alternative strategies where we selected the *last* factoring generated when imposing a bound  $M$  on the sum  $\sum_{F^L \in \mathcal{F}_i^L} \prod_{v \in F^L} |\mathcal{D}(v)|$  of estimated leaf sizes. In almost all domains, this either hardly changed the factoring at all, or changed it only for large  $M$  (1 million and more), and then only on the smallest instances where we obtained a single “leaf” containing most of the state space, see next.

We abstain from single-leaf factorings because these nearly always arise from a factoring grouping most of the state variables into the leaf. In transportation domains with capacity constraints, for example, the “horizontal line” is drawn between the vehicles and the entire rest of the task.



Such factorings *can* sometimes be useful; e. g., in IPC’08-optimal Transport instance p04, using LM-cut evaluations go down from 95281 for  $A^*$  to 17606 for  $FDA^*$ , and runtime goes down from 56.8 to 17.6 seconds. But the huge single leaf is typically manageable only on the smallest instances of a domain. An interesting topic for future work is to handle large leaves using a symbolic, rather than enumerative, representation. For now, we concentrate on exploiting conditional independence, across  $> 1$  leaves.

Consider Table 1. One might wonder whether the radical state space size reduction is pathological (if the entire task is a single “leaf” then  $|S_R^F| = 1$  but “MaxSize” =  $|S_R|$ ), yet the right-hand side of the table shows that this is not so. The only two domains where leaves do get “large” are TPP and Woodworking, but compared to the state space reduction they are still tiny. Observe, though, that the state space reduction does not pay off in Miconic, where it does not outweigh the runtime/memory overhead and more standard state spaces are built. The reduction tends to be less if the leaves are “too small”, as in Miconic, Rovers, Satellite, and Pathways. This makes sense as the leaves then cannot “do much on their own”, so, relatively speaking, less world states are condensed into each single fork-decoupled state.

In all except the domains from Table 1, we abstain from every instance. While this may be construed as a weakness of fork factoring, arguably it is, at least as much, a strength. Almost all known planning/search techniques work well only on a limited domain subset. For example, state-of-the-art partial-order reduction (POR) (Wehrle and Helmert 2014), a long-standing technique into which intense effort was invested, improves coverage with LM-cut on 7 IPC domains, substantially ( $> 1$  instance) on 5. These numbers are almost the same, actually slightly better (7 and 6, Table 2), for fork factoring. However, in difference to POR and basically all other known techniques, fork factoring *abstains* outside its scope: The structure it requires to work well is explicit and easily testable. This is perfect for portfolios (e. g. (Xu et al. 2008; Cenamor, de la Rosa, and Fernández 2014)), combining the strengths of *different* basic solvers, which is necessary for effective general off-the-shelf tools anyhow.

## Planner Performance Results

For satisficing planning, the domains within our scope are often too easy for fork factoring to be beneficial. The most notable exception is transportation with fuel consumption. Denote with GBFS FD’s lazy-greedy best-first search with a second open list for preferred operators. On the IPC’11 satisficing NoMystery benchmarks, while GBFS solves 9 instances with  $h^{FF}$ , FDGBFS solves 19. Figure 6 examines NoMystery more closely, on Nakhost et al.’s (2012) controlled benchmarks. We ran the respective best satisficing and optimal planners from Nakhost et al.’s experiments, as well as IPC’14 optimal-track winner *Symba* (Torralba and Alcázar 2013). FDGBFS solves *all* instances, vastly outperforming its satisficing competitors when resources are scarce.  $FDA^*$  with LM-cut solves more than 90% of the instances in “small”, vastly outperforming merge-and-shrink (M&S). It solves all instances in “large”, in stark contrast to the optimal planners run by Nakhost et al., as well as M&S




Figure 6: % instances covered in Nakhost et al.’s (2012) NoMystery test suites, over *resource constrainedness* (factor by which the available fuel exceeds the minimum required).

in our own runs here, none of which managed to solve *any* of these instances. *Symba* handles “large” much better, though clearly worse than  $FDA^*$ ; in “small”, *Symba* has only 44% coverage in total (data not included for readability).

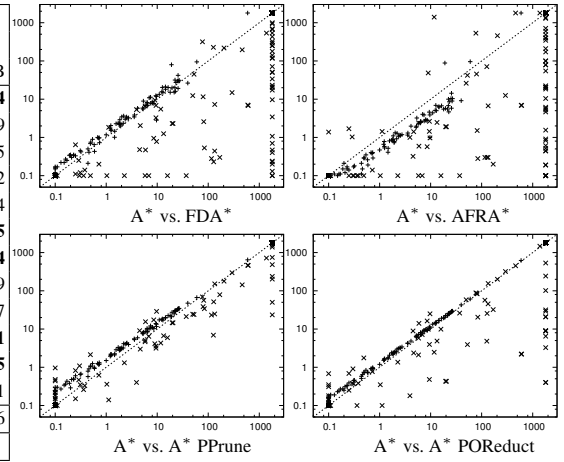
Table 2 gives full results for optimal planning. To clarify the competition: Almost all works on factored planning have notions of local- or step-optimality, and could be rendered globally optimal only by impractical searches over bounds on the number of (local) steps. Fabre et al. (2010) do guarantee global cost-optimality, but report themselves that their approach is impractical (except on manually reformulated versions of the IPC’04 Promela domains). We compare against partition pruning in  $A^*$  (Nissim, Apsel, and Brafman 2012; Nissim and Brafman 2014), the only other work exploiting some form of decomposition in globally cost-optimal heuristic search. We compare against POR (Wehrle and Helmert 2012; Wehrle et al. 2013; Wehrle and Helmert 2014), in its best-performing (on our non-abstained benchmarks) variant, as a representation of the state of the art in admissible search space pruning.

Fork-decoupled state space search clearly has the edge in overall coverage. For each of LM-cut,  $h^{\max}$ , and the blind heuristic,  $FDA^*$  and  $AFRA^*$  coverage dominates that of  $A^*$  in all domains except Miconic (as well as 2 instances in Zenotravel for  $AFRA^*$  with LM-cut). Our techniques yield amazing benefits in Logistics, NoMystery, and TPP; strong benefits in Woodworking and Rovers (the latter especially when using the alternative factoring strategy described above); and minor to significant benefits in various other cases. In Logistics, NoMystery, and TPP, as well as Rovers with the alternative factoring strategy, even  $h^{\max}$  and blind search with fork-decoupled search have higher coverage than LM-cut with standard search. Indeed, *when ignoring the oversize Miconic domain,  $h^{\max}$  and blind search outperform LM-cut in summed-up coverage* (109–112 vs. 97).

In almost all cases, as evidenced by the coverage of  $FRA^*$ , the first plan found by  $AFRA^*$  is proved optimal by early termination already. We need to continue search only sometimes in Driverlog, Woodworking, and Zenotravel.

The search space reduction obtained by fork decoupling is typically much larger than that for partition pruning and POR, although there are exceptions (Pathways, Satellite, Woodworking). Similarly for runtime. The runtime factors are sometimes skewed by the small size of commonly solved instances, as well as FD’s rounding; e. g., in TPP, *all* runs of  $FDA^*$  and  $AFRA^*$  on commonly solved instances are rounded to 0.1. The scatter plots shed additional light on

		(A): Coverage																		
		A*			FDA*			AFRA*			FRA*			PPrune			POReduce			SB
#		BL	HM	LM	BL	HM	LM	BL	HM	LM	BL	HM	LM	BL	HM	LM	BL	HM	LM	
Driver	20	7	9	13	<i>11</i>	<i>11</i>	13	<i>11</i>	<i>11</i>	13	9	10	10	7	9	13	7	9	13	<b>14</b>
Log00	28	10	12	20	22	22	<b>28</b>	22	22	25	22	22	25	<i>11</i>	12	20	10	12	20	19
Log98	35	2	2	<b>6</b>	4	3	<b>6</b>	4	4	<b>6</b>	4	4	<b>6</b>	3	3	<b>6</b>	2	3	<b>6</b>	5
Mico	145	50	50	<b>136</b>	36	36	135	36	36	135	36	36	135	45	45	<b>136</b>	45	45	<b>136</b>	102
NoMy	20	8	8	14	<i>17</i>	<i>18</i>	<b>20</b>	<i>17</i>	<i>19</i>	<b>20</b>	<i>17</i>	<i>18</i>	<b>20</b>	8	8	14	8	8	14	14
Pathw	29	3	3	4	3	3	4	3	3	4	3	3	4	3	3	4	3	3	4	<b>5</b>
Rovers	40	6	6	7	6	6	8	6	7	9	6	7	9	6	6	<i>10</i>	7	8	9	<b>14</b>
Satell	36	6	6	7	6	6	7	6	6	8	6	6	8	7	7	<b>12</b>	6	7	<i>11</i>	9
TPP	27	5	5	5	<b>23</b>	22	<i>18</i>	<b>23</b>	22	<b>23</b>	<b>23</b>	22	<b>23</b>	5	5	5	5	5	5	7
Woo08	13	4	5	6	5	6	10	5	6	<b>11</b>	1	2	7	4	5	6	6	8	<b>11</b>	<b>11</b>
Woo11	5	0	1	2	<i>1</i>	<i>1</i>	4	<i>1</i>	<i>1</i>	<b>5</b>	1	1	<b>5</b>	0	1	2	1	3	<b>5</b>	<b>5</b>
Zeno	20	8	8	<b>13</b>	<i>11</i>	<i>11</i>	<b>13</b>	<i>11</i>	<i>11</i>	11	9	9	8	8	8	<b>13</b>	8	8	<b>13</b>	11
$\Sigma$	418	109	115	233	<i>145</i>	<i>145</i>	266	<i>145</i>	<i>148</i>	<b>270</b>	<i>137</i>	<i>140</i>	260	107	112	241	108	119	247	216
Rovers $\mathcal{F}'$					9	9	9	9	9	9	7	7	7							



		(B): Evaluations with LM-cut: Improvement factor relative to A*									(C): Runtime with LM-cut: Improvement factor relative to A*															
		FDA*			AFRA*			PPrune			POReduce			FDA*			AFRA*			PPrune			POReduce			
#		$\Sigma D$	Gmea	max	$\Sigma D$	Gmea	max	$\Sigma D$	Gmea	max	$\Sigma D$	Gmea	max	$\Sigma D$	Gmea	max	$\Sigma D$	Gmea	max	$\Sigma D$	Gmea	max	$\Sigma D$	Gmea	max	
Driver	13	2.2	3.5	73.2	1.2	3.0	83.3	1.4	1.3	2.0	1.0	1.2	1.9	9	0.7	1.8	31.6	0.4	1.7	54.0	1.2	1.0	1.7	1.0	0.9	1.0
Log00	20	960.3	130.7	1407	562.3	106.5	1119	2.8	2.1	5.4	2.9	2.7	22.3	11	311.9	67.1	561.1	337.9	69.9	841.7	2.9	2.2	5.4	3.0	2.4	5.2
Log98	6	49.3	11.4	458.8	0.2	1.7	458.8	4.3	2.4	3.0	2.5	3.4	5.4	4	14.8	7.7	158.3	0.1	1.0	299.6	2.4	1.9	2.7	3.9	2.7	4.8
Mico	135	2.2	1.7	6.5	2.2	1.7	6.5	1.0	1.0	1.0	1.0	1.0	1.0	100	0.8	0.8	3.4	1.7	2.0	9.7	0.7	0.6	0.9	0.9	0.8	1.0
NoMy	14	20.4	8.8	22.2	8.4	5.0	9.3	1.0	1.0	1.0	1.0	1.0	1.0	12	17.9	4.5	19.7	20.5	4.8	23.4	0.9	0.4	1.0	0.8	0.4	0.9
Pathw	4	1.9	1.3	1.9	1.7	1.2	1.7	2.2	1.3	2.2	16.2	3.0	16.7	1	1.2	1.2	1.2	1.2	1.2	1.2	2.2	2.2	2.2	13.4	13.4	13.4
Rovers	7	4.8	2.5	8.9	7.8	3.4	15.5	2.4	1.6	3.2	4.4	2.3	19.9	3	1.9	2.1	2.9	4.3	4.3	6.5	2.4	2.3	3.0	3.5	4.5	19.6
Satell	7	1.0	1.1	1.4	2.2	1.6	5.1	16.3	3.6	16.9	21.5	5.6	27.8	3	0.6	0.6	0.7	1.8	1.8	3.9	17.6	11.3	18.2	25.7	18.6	26.2
TPP	5	2491	49.0	5657	2858	52.7	6869	1.1	0.9	1.6	1.0	1.0	1.0	1	68.2	68.2	68.2	68.2	68.2	68.2	1.1	1.1	1.1	1.0	1.0	1.0
Woo08	6	82.4	42.9	135.2	72.8	41.4	108.0	1.4	1.2	1.5	243.2	10.6	346.9	2	65.9	27.3	84.7	69.8	30.4	85.9	1.3	1.3	1.4	233.6	113.1	269.8
Woo11	2	82.4	55.4	135.2	72.8	49.6	108.0	1.4	1.4	1.5	250.3	173.7	346.9	2	68.7	27.5	89.1	71.6	30.3	88.5	1.3	1.3	1.4	237.7	112.1	275.1
Zeno	11	6.5	3.0	13.2	1.4	1.3	4.0	2.3	1.5	3.1	1.0	1.0	1.3	7	4.0	1.6	6.9	1.1	0.8	3.1	2.0	1.0	2.2	0.9	0.8	1.0

Table 2: Performance on IPC optimal-track instances where our factoring strategy does not abstain. Best coverage **bold**, better-than-baseline coverage *italic*. “BL” blind heuristic; “HM”  $h^{\max}$ ; “LM” LM-cut; “PPrune” partition pruning; “POReduce” partial-order reduction; “SB” Symba. The scatter plots show runtime with LM-cut ( $A^*$  on the  $x$ -axis). In (B) and (C), “ $\Sigma D$ ” is the factor over the per-domain sums, and “Gmea” (“max”) is the Geometric mean (the maximum) over the per-instance factors. Instances for (B) are commonly solved ones, for (C) the same but excluding ones commonly solved in  $\leq 0.1$  seconds (FD rounds runtimes  $\leq 0.1$  to 0.1). Rovers  $\mathcal{F}'$  uses our alternative factoring strategy with leaf size bound  $M = 10$  million.

this, showing that  $FDA^*$  and  $AFRA^*$  are highly beneficial here, at a slightly higher risk in the case of  $AFRA^*$ .

Of course, our observations must be qualified against the benchmark selection: fork factoring is great *within its scope*. Yet keep in mind that, outside this scope, we can after negligible runtime invoke whatever alternative planner we want.

### Future Research Directions

In preliminary work, we have determined that the scope of fork factoring can be extended to *star-topology factoring*: replace the “root” with a “center” and the leaves with “periphery factors”, allowing arbitrary interactions between the center and each periphery factor. This corresponds to a much generalized notion of “conditional independence”: For any fixed center path, the periphery factors can still be moved independently of each other, given that each complies with the preconditions requested and provided by the center, as well as co-occurring effects with the center. Note that this allows to address planning tasks whose causal graph is strongly connected. Indeed, if we are willing to tackle single-leaf factorings (see next), then *any* partition of the state variables into two subsets yields a star-topology factoring.

We currently enumerate leaf states explicitly, a major

caveat when dealing with large leaf factors. But one could represent the pricing functions *symbolically* instead, via ADDs (Bahar et al. 1997). This is true also of star-topology factoring, which then becomes a highly configurable form of mixing explicit-state search with symbolic search, including in particular the option to freely divide the state variables into ones (“center”) tackled by explicit search vs. ones (“single periphery factor”) tackled by symbolic search.

Beyond this, the road opens up to target-profile factoring more generally. Are there interesting structures other than “conditional independence” to exploit? In chain structures, e. g., what about restrictions on the “domain transition graphs” of the factors to guarantee that backtracking won’t be needed, abstaining if these restrictions are not satisfied? Can we combine different target profiles through hierarchical sub-structuring, where “factors” at a higher level are optimally solved via a target profile at a lower level? Viewing target-profile factoring as a highly parameterized way of reformulating the search, can we benefit from orthogonal enhancements like partial-order reduction, macro-actions (e. g. (Vidal 2004; Botea et al. 2005)), symmetries (e. g. (Domshlak, Katz, and Shleyfman 2012)))? Can we understand under which circumstances a reformulation is beneficial or detri-

mental? Can we predict which alternative target profile best matches the input, and design self-configuring portfolios?

Last but not least, it is worth mentioning that star topology is a classical system design paradigm in many areas of CS. Hence star-topology decoupled state space search could turn out to be useful for control, synthesis, and verification problems far beyond the realm of AI Planning.

**Acknowledgments.** We thank Carmel Domshlak for discussions, which contributed to many of the ideas for future research.

## References

- Amir, E., and Engelhardt, B. 2003. Factored planning. In Gottlob, G., ed., *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, 929–935. Acapulco, Mexico: Morgan Kaufmann.
- Bäckström, C., and Nebel, B. 1995. Complexity results for SAS<sup>+</sup> planning. *Computational Intelligence* 11(4):625–655.
- Bahar, R. I.; Frohm, E. A.; Gaona, C. M.; Hachtel, G. D.; Macii, E.; Pardo, A.; and Somenzi, F. 1997. Algebraic decision diagrams and their applications. *Formal Methods in System Design* 10(2/3):171–206.
- Bonet, B., and Geffner, H. 2001. Planning as heuristic search. *Artificial Intelligence* 129(1–2):5–33.
- 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.
- Botea, A.; Enzenberger, M.; Müller, M.; and Schaeffer, J. 2005. Macro-FF: Improving AI planning with automatically learned macro-operators. *Journal of Artificial Intelligence Research* 24:581–621.
- Brafman, R., and Domshlak, C. 2003. Structure and complexity in planning with unary operators. *Journal of Artificial Intelligence Research* 18:315–349.
- Brafman, R. I., and Domshlak, C. 2006. Factored planning: How, when, and when not. In Gil, Y., and Mooney, R. J., eds., *Proceedings of the 21st National Conference of the American Association for Artificial Intelligence (AAAI-06)*, 809–814. Boston, Massachusetts, USA: AAAI Press.
- Brafman, R. I., and Domshlak, C. 2008. From one to many: Planning for loosely coupled multi-agent systems. In Rintanen et al. (2008), 28–35.
- Brafman, R., and Domshlak, C. 2013. On the complexity of planning for agent teams and its implications for single agent planning. *Artificial Intelligence* 198:52–71.
- Cenamor, I.; de la Rosa, T.; and Fernández, F. 2014. IBACOP and IBACOP2 planner. In *IPC 2014 planner abstracts*, 35–38.
- Domshlak, C.; Hoffmann, J.; and Katz, M. 2015. Red-black planning: A new systematic approach to partial delete relaxation. *Artificial Intelligence* 221:73–114.
- Domshlak, C.; Katz, M.; and Shleyfman, A. 2012. Enhanced symmetry breaking in cost-optimal planning as forward search. In Bonet et al. (2012).
- Fabre, E.; Jezequel, L.; Haslum, P.; and Thiébaux, S. 2010. Cost-optimal factored planning: Promises and pitfalls. In Brafman, R. I.; Geffner, H.; Hoffmann, J.; and Kautz, H. A., eds., *Proceedings of the 20th International Conference on Automated Planning and Scheduling (ICAPS'10)*, 65–72. AAAI Press.
- 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 19th International Conference on Automated Planning and Scheduling (ICAPS'09)*, 162–169. AAAI Press.
- 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.
- Jonsson, P., and Bäckström, C. 1995. Incremental planning. In *European Workshop on Planning*.
- Katz, M., and Domshlak, C. 2008. Structural patterns heuristics via fork decomposition. In Rintanen et al. (2008), 182–189.
- Kelareva, E.; Buffet, O.; Huang, J.; and Thiébaux, S. 2007. Factored planning using decomposition trees. In Veloso, M., ed., *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI-07)*, 1942–1947. Hyderabad, India: Morgan Kaufmann.
- Knoblock, C. 1994. Automatically generating abstractions for planning. *Artificial Intelligence* 68(2):243–302.
- Laneky, A. L., and Getoor, L. 1995. Scope and abstraction: Two criteria for localized planning. In Mellish, S., ed., *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI-95)*, 1612–1619. Montreal, Canada: Morgan Kaufmann.
- McDermott, D. V. 1999. Using regression-match graphs to control search in planning. *Artificial Intelligence* 109(1–2):111–159.
- Nakhost, H.; Hoffmann, J.; and Müller, M. 2012. Resource-constrained planning: A monte carlo random walk approach. In Bonet et al. (2012), 181–189.
- Nissim, R., and Brafman, R. 2014. Distributed heuristic forward search for multi-agent planning. *Journal of Artificial Intelligence Research* 51:293–332.
- Nissim, R.; Apsel, U.; and Brafman, R. I. 2012. Tunneling and decomposition-based state reduction for optimal planning. In Raedt, L. D., ed., *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI'12)*, 624–629. Montpellier, France: IOS Press.
- Rintanen, J.; Nebel, B.; Beck, J. C.; and Hansen, E., eds. 2008. *Proceedings of the 18th International Conference on Automated Planning and Scheduling (ICAPS'08)*. AAAI Press.
- Sacerdoti, E. D. 1974. Planning in a hierarchy of abstraction spaces. *Artificial Intelligence* 5:115135.
- Torralba, A., and Alcázar, V. 2013. Constrained symbolic search: On mutexes, BDD minimization and more. In

Helmert, M., and Röger, G., eds., *Proceedings of the 6th Annual Symposium on Combinatorial Search (SOCS'13)*, 175–183. AAAI Press.

Vidal, V. 2004. A lookahead strategy for heuristic search planning. In Koenig, S.; Zilberstein, S.; and Koehler, J., eds., *Proceedings of the 14th International Conference on Automated Planning and Scheduling (ICAPS'04)*, 150–160. Whistler, Canada: Morgan Kaufmann.

Wehrle, M., and Helmert, M. 2012. About partial order reduction in planning and computer aided verification. In Bonet et al. (2012).

Wehrle, M., and Helmert, M. 2014. Efficient stubborn sets: Generalized algorithms and selection strategies. In Chien, S.; Do, M.; Fern, A.; and Ruml, W., eds., *Proceedings of the 24th International Conference on Automated Planning and Scheduling (ICAPS'14)*. AAAI Press.

Wehrle, M.; Helmert, M.; Alkhazraji, Y.; and Mattmüller, R. 2013. The relative pruning power of strong stubborn sets and expansion core. In Borrajo, D.; Fratini, S.; Kambhampati, S.; and Oddi, A., eds., *Proceedings of the 23rd International Conference on Automated Planning and Scheduling (ICAPS'13)*. Rome, Italy: AAAI Press.

Xu, L.; Hutter, F.; Hoos, H. H.; and Leyton-Brown, K. 2008. Satzilla: Portfolio-based algorithm selection for sat. *Journal of Artificial Intelligence Research* 32:565–606.