

Beyond Stars – Generalized Topologies for Decoupled Search

Daniel Gnad^{1,2}, Álvaro Torralba³, Daniel Fišer^{1,4}

¹Saarland University, Saarland Informatics Campus, Saarbrücken, Germany

²Linköping University, Sweden

³Aalborg University, Denmark

⁴Czech Technical University in Prague, Faculty of Electrical Engineering, Czech Republic
daniel.gnad@liu.se; alto@cs.aau.dk; danfis@danfis.cz

Abstract

Decoupled search decomposes a classical planning task by partitioning its variables such that the dependencies between the resulting factors form a star topology. In this topology, a single center factor can interact arbitrarily with a set of leaf factors. The leaves, however, can interact with each other only indirectly via the center. In this work, we generalize this structural requirement and allow arbitrary topologies. The components must not overlap, i. e., each state variable is assigned to exactly one factor, but the interaction between factors is not restricted. We show how this generalization is connected to star topologies, which implies the correctness of decoupled search with this novel type of decomposition. We introduce factoring methods that automatically identify these topologies on a given planning task. Empirically, the generalized factorings lead to increased applicability of decoupled search on standard IPC benchmarks, as well as to superior performance compared to known factoring methods.

Introduction

Star-topology decoupled state-space search, decoupled search for short, tackles the state explosion problem by exploiting the dependency structure of the given model. In classical planning, decoupled search decomposes planning tasks by partitioning the state variables such that the dependencies between the resulting factors form a star topology (Gnad and Hoffmann 2018). Here, a single *center factor* C can interact arbitrarily with the remaining set of *leaf factors* $\{L_1, \dots, L_n\}$. Thereby, decoupled search exploits a form of conditional independence of the leaves; given a sequence of *center actions*, i. e., actions that have an effect on the center, the leaves are independent. This allows for an enumeration of the *compliant leaf paths* for each leaf separately. The search is then performed over center actions only, where each sequence of center actions π^C ends in a *decoupled state* consisting of a single *center state* (an assignment to C) and a non-empty set of *leaf states* (assignments to each L_i). The leaf states reached for an L_i are exactly those L_i -assignments reachable by any sequence of L_i -actions that can be executed along π^C , i. e., that is *compliant* with π^C . Thus, a decoupled state compactly represents exponentially

many explicit states, which share the same center state and result from all combinations of leaf states across leaf factors.

Application of decoupled search requires finding a suitable decomposition of the task into factors. Existing work has explored several methods to automatically find such a factoring. They partition the variables into factors by analyzing their causal dependencies (Gnad, Poser, and Hoffmann 2017; Gnad and Hoffmann 2018), or using integer linear programming (ILP) (Schmitt, Gnad, and Hoffmann 2019). All known approaches focus specifically on star-topology factorings with a designated center factor, where cross-leaf interactions are limited to actions affecting the center.

In this work, we generalize the concept of star factorings to *generalized factorings*, which allow arbitrary cross-factor interactions. In fact, we even allow factorings *without* center factor. Our main motivation for the generalization is that very few actions (even a single one) can render star factorings impossible if such actions lead to direct cross-leaf interactions. By relaxing the requirements of factorings, such actions can simply be made center actions, which are considered by the main search. We do so by extending the notion of center actions to those that share precondition or effect variables with more than one leaf factor. We prove the correctness of this novel form of decomposition by connecting it to the known star topologies, extending the possibilities of decoupled search while keeping all its desirable properties.

We illustrate the advantages of generalized factorings on a collaborative robotics task and devise a new ILP encoding with different optimization criteria to decompose planning tasks. Our evaluation in optimal and satisficing settings shows that generalized factorings open up a wide range of possibilities, enabling the use of decoupled search on planning tasks that could not be tackled before. While there is no single best strategy for every domain or setting, the factorings from our new strategies can significantly improve performance over prior methods and explicit-state search.

Some full proofs, details of our ILP encoding, and additional results are provided in a technical report (TR) (Gnad, Torralba, and Fišer 2022).

Background

A *planning task* (Bäckström and Nebel 1995) is a tuple $\Pi = \langle \mathcal{V}, \mathcal{A}, I, G \rangle$, where \mathcal{V} is a finite set of *variables*, and each variable $v \in \mathcal{V}$ has a finite domain $\mathcal{D}(v)$. \mathcal{A}

is a finite set of *actions*. Each action $a \in \mathcal{A}$ is a triple $\langle \text{pre}(a), \text{eff}(a), \text{cost}(a) \rangle$, where preconditions $\text{pre}(a)$ and effects $\text{eff}(a)$ are partial assignments to \mathcal{V} , and the cost $\text{cost}(a) \in \mathbb{R}^{0+}$ is a non-negative real number. A state is a complete assignment to \mathcal{V} , I is the *initial state*, and the *goal* G is a partial assignment to \mathcal{V} . For a partial assignment p , we denote by $\text{vars}(p) \subseteq \mathcal{V}$ the subset of variables on which p is defined. For $V' \subseteq \mathcal{V}$, we denote the restriction of p onto V' by $p[V']$, i. e., the assignment to $V' \cap \text{vars}(p)$ by p . We identify (partial) states with sets of variable/value pairs. An action a is *applicable* in a state s if $\text{pre}(a) \subseteq s$. Applying a in a (partial) state s changes the value of all $v \in \text{vars}(\text{eff}(a)) \cap \text{vars}(s)$ to $\text{eff}(a)[v]$, and leaves s unchanged elsewhere. The outcome state is denoted $s[a]$. A *plan* for Π is an action sequence π applicable in I that ends in a state $s_G \supseteq G$; it is *optimal* if its summed-up action cost, denoted $\text{cost}(\pi)$, is minimal among all plans for Π .

Decoupled search is a technique developed to avoid the combinatorial explosion of having to enumerate all possible variable assignments of causally independent parts of a planning task. It does so by partitioning the state variables \mathcal{V} into a *star factoring*, whose elements are called *factors*.

Definition 1 (Star Factoring) Let $\Pi = \langle \mathcal{V}, \mathcal{A}, I, G \rangle$ be a planning task. A factoring $\mathcal{F} \subset 2^{\mathcal{V}}$ for Π is a partition of \mathcal{V} .

A pair $\mathcal{F}_s = \langle C, \mathcal{L} \rangle$ is a star factoring for Π , if $\{C\} \cup \mathcal{L}$ is a factoring and for all actions $a \in \mathcal{A}$ either there exists an $L \in \mathcal{L}$ such that $\text{vars}(\text{pre}(a)) \subseteq C \cup L$ and $\text{vars}(\text{eff}(a)) \subseteq L$, or $\text{vars}(\text{eff}(a)) \cap C \neq \emptyset$. C is called the *center factor* of \mathcal{F}_s , and \mathcal{L} are its *leaf factors*.

By imposing a structural requirement on the interaction between the factors, namely a *star topology*, decoupled search can efficiently handle cross-factor dependencies. Here, the center C can interact arbitrarily with the leaves \mathcal{L} , but interaction between leaves is allowed only if the center is affected at the same time. Actions affecting C , i. e., with an effect on a variable in C , are called *center actions*, denoted \mathcal{A}^C , and those affecting a leaf are called *leaf actions*, denoted \mathcal{A}^L . The actions that affect a particular leaf $L \in \mathcal{L}$ are denoted \mathcal{A}^L . We define the set of *leaf-only actions* of a leaf L as $\mathcal{A}_G^L := \mathcal{A}^L \setminus \mathcal{A}^C$. A sequence of center actions applicable in I in the projection onto C is a *center path*, a sequence of \mathcal{A}^L -actions applicable in I in the projection onto L , is a *leaf path*. A complete assignment to C , or to an $L \in \mathcal{L}$, is called a *center state*, or *leaf state*, respectively. S^C is the set of all leaf states and that of a particular leaf L is denoted S^L .

A *decoupled state* $s^{\mathcal{F}}$ is a pair $\langle \text{center}(s^{\mathcal{F}}), \text{prices}(s^{\mathcal{F}}) \rangle$ where $\text{center}(s^{\mathcal{F}})$ is a center state, and $\text{prices}(s^{\mathcal{F}}) : S^{\mathcal{L}} \mapsto \mathbb{R}^{0+} \cup \{\infty\}$ is a *pricing function*, mapping each leaf state to a non-negative *price*. By $\pi^C(s^{\mathcal{F}})$ we denote the center path that starts in the initial decoupled state $I^{\mathcal{F}}$ and ends in $s^{\mathcal{F}}$. The pricing function is maintained during decoupled search in a way so that the price of a leaf state s^L is the cost of a cheapest leaf path that ends in s^L and that is *compliant* with $\pi^C(s^{\mathcal{F}})$, i. e., that can be scheduled alongside the center path executed up to $s^{\mathcal{F}}$. A decoupled state $s^{\mathcal{F}}$ satisfies a condition p , a partial state, denoted $s^{\mathcal{F}} \models p$, iff (i) $p[C] \subseteq \text{center}(s^{\mathcal{F}})$ and (ii) for every $L \in \mathcal{L}$ there exists an $s^L \in S^L$ s.t. $p[L] \subseteq s^L$ and $\text{prices}(s^{\mathcal{F}})[s^L] < \infty$. We

define the set of leaf actions *enabled* by a center state s^C as $\mathcal{A}^C|_{s^C} := \{a^L \mid a^L \in \mathcal{A}^L \wedge \text{pre}(a^L)[C] \subseteq s^C\}$. For a center state s^C and a pair of leaf states $s_1^L, s_2^L \in S^L$, by $c_{s^C}(s_1^L, s_2^L)$ we define the cost of a cheapest path of $\mathcal{A}_{\emptyset}^L|_{s^C}$ actions from s_1^L to s_2^L . If no such path exists $c_{s^C}(s_1^L, s_2^L) = \infty$. A decoupled state $s^{\mathcal{F}}$ represents a set of explicit states, its *member states*, namely those states s where $s^{\mathcal{F}} \models s$.

Definition 2 (Decoupled State Space) Let Π be a planning task, and $\mathcal{F} = \langle C, \mathcal{L} \rangle$ a star factoring for Π . The decoupled state space is a labeled transition system $\Theta_{\Pi}^{\mathcal{F}} = \langle S^{\mathcal{F}}, \mathcal{A}^C, \text{cost}|_{\mathcal{A}^C}, \mathcal{T}^{\mathcal{F}}, I^{\mathcal{F}}, S_G^{\mathcal{F}} \rangle$ as follows:

- (i) $S^{\mathcal{F}}$ is the set of all decoupled states.
- (ii) The transition labels are the center actions \mathcal{A}^C .
- (iii) The cost function is that of Π , restricted to \mathcal{A}^C .
- (iv) $\mathcal{T}^{\mathcal{F}}$ contains a transition $s^{\mathcal{F}} \xrightarrow{a^C} t^{\mathcal{F}} \in \mathcal{T}^{\mathcal{F}}$ whenever $a^C \in \mathcal{A}^C$ and $s^{\mathcal{F}}, t^{\mathcal{F}} \in S^{\mathcal{F}}$ are such that:
 1. $\pi^C(s^{\mathcal{F}}) \circ \langle a^C \rangle = \pi^C(t^{\mathcal{F}})$,
 2. $s^{\mathcal{F}} \models \text{pre}(a^C)$,
 3. $\text{center}(s^{\mathcal{F}})[a^C] = \text{center}(t^{\mathcal{F}})$,
 4. for every leaf $L \in \mathcal{L}$ and leaf state $s^L \in S^L$, if $s^L \models \text{pre}(a^C)[L]$, then $\text{prices}(t^{\mathcal{F}})[s^L][a^C] = \text{prices}(s^{\mathcal{F}})[s^L]$. Additionally, $\text{prices}(t^{\mathcal{F}})[s^L] = \min_{t^L \in S^L} \text{prices}(t^{\mathcal{F}})[t^L] + c_{\text{center}(t^{\mathcal{F}})}(t^L, s^L)$.
- (v) $I^{\mathcal{F}}$ is the decoupled initial state, where $\text{center}(I^{\mathcal{F}}) := I[C]$, $\pi^C(I^{\mathcal{F}}) := \langle \rangle$, and, for every leaf $L \in \mathcal{L}$, $\text{prices}(I^{\mathcal{F}})[I[L]] = 0$ and for all other leaf states $s^L \in S^L$, $\text{prices}(I^{\mathcal{F}})[s^L] = c_{\text{center}(I^{\mathcal{F}})}(I[L], s^L)$.
- (vi) $S_G^{\mathcal{F}} = \{s_G^{\mathcal{F}} \mid s_G^{\mathcal{F}} \models G\}$ are the decoupled goal states.

Decoupled search runs a search over center actions only, enumerating, for each leaf separately, the set of leaf states that can be reached in the form of the pricing function. Every search algorithm and heuristic function can be employed on the decoupled state space (Gnad and Hoffmann 2018).

Beyond Star Topologies

In this section, we generalize the concept of star factorings by introducing *generalized factorings*. The key novelty is that we no longer impose any structural restriction regarding the dependencies between factors, i. e., any partition of the state variables is a generalized factoring.

Definition 3 (Generalized Factoring) Let Π be a planning task. A pair $\mathcal{F}_g = \langle C, \mathcal{L} \rangle$ is a generalized factoring for Π , if either \mathcal{L} or $\{C\} \cup \mathcal{L}$ is a factoring for Π . C is the (possibly empty) center factor of \mathcal{F}_g , and \mathcal{L} are its leaf factors.

In addition to allowing arbitrary cross-factor interactions, we also consider empty center factors, i. e., factorings where all components are leaves. Note that every star factoring is also a generalized factoring, but not vice versa. Next, we resolve complications with respect to the interaction between factors by introducing the notion of *global actions*, which replace the center actions from star factorings.

Definition 4 (Global Action) Let $\mathcal{F}_g = \langle C, \mathcal{L} \rangle$ be a generalized factoring for the planning task Π . An action $a \in \mathcal{A}$ is a global action iff there does not exist an $L \in \mathcal{L}$ such that

$\text{vars}(\text{pre}(a)) \subseteq C \cup L$ and $\text{vars}(\text{eff}(a)) \subseteq L$. The set of all global actions is denoted \mathcal{A}^G .

While leaf and leaf-only actions are defined as before, we need to adapt center actions to be able to handle more complex interactions. We call the new type of action *global* to make the distinction clear, since global actions not necessarily affect the center. Essentially, an action a is global if it is not a leaf-only action of any leaf. In particular, actions with preconditions or effects on more than one leaf, but without center effect become global actions.

The basic concept of decoupled search remains intact. In Definition 2, we replace mentions to center actions \mathcal{A}^C by global actions \mathcal{A}^G . The search then branches over global actions, and all leaf states reachable via leaf-only actions are enumerated for each leaf separately. The center state associated with a decoupled state can now be *empty* (if $C = \emptyset$). Hence, while we generalize the notion of factorings, the underlying idea of the decomposition is preserved. In particular, even though it might seem we adopt the flexibility of traditional factored planning approaches—and thus their weaknesses—our search still strictly distinguishes between center/global transitions where components interact, and leaf-only transitions that do not affect another leaf.

Example 1 We show the advantages of generalized factorings on a collaborative robotics task where n battery-powered agents (a_i variables) move on a map. Moving consumes battery (b_i variables), and an agent can charge another one by sharing its battery charge. A single agent has a non-empty battery initially, all agents need to reach a goal location. Formally, the task is defined with variables $\mathcal{V} = \{a_1, b_1, \dots, a_n, b_n\}$, domains $\mathcal{D}(a_i) = \{1, \dots, l_m\}$, $\mathcal{D}(b_i) = \{0, \dots, B\}$, actions $\mathcal{A} = \{\text{move}(a_i, l_x, l_y, k), \text{charge}(a_i, a_j, l, x, y)\}$, where $i \neq j$, initial state $I = \{a_1 = l_1, b_1 = B\} \cup \{a_i = l_i, b_i = 0 \mid i > 1\}$, and goal $G = \{a_i = l_1\}$. The actions are defined as follows: $\text{pre}(\text{move}(a_i, l_x, l_y, k)) = \{a_i = l_x, b_i = k\}$, $\text{eff}(\text{move}(a_i, l_x, l_y, k)) = \{a_i = l_y, b_i = k - 1\}$, and $\text{pre}(\text{charge}(a_i, a_j, l, x, y)) = \{a_i = l, a_j = l, b_i = x, b_j = y\}$, $\text{eff}(\text{charge}(a_i, a_j, l, x, y)) = \{b_i = x - 1, b_j = y + 1\}$.

With generalized factorings, we can have every agent with its battery in a leaf $L_i = \{a_i, b_i\}$, the center is empty, resulting in a factoring where the number of leaves scales with the number of agents. The search then branches over the global charge actions, the move actions are leaf-only actions.

Due to the charge actions, however, by which any pair of agents can interact, in a star factoring we can only place a subset of the agents (and possibly their battery) into one leaf, and the remaining variables in the center. Thus, only a linear state-space reduction can be achieved by decoupled search. We remark that there actually exists a scaling star factoring (i. e., where $|\mathcal{L}| = n$), namely placing all battery variables in the center, and each agent in a separate leaf. As we will show in Proposition 1, though, this does not lead to any state-space reduction, since all actions are center actions.

Correctness

We prove the correctness of decoupled search with generalized factorings using a polynomial mapping from a task Π

and a generalized factoring \mathcal{F}_g to a modified task Π_s and a star factoring \mathcal{F}_s such that the decoupled state spaces $\Theta_{\Pi}^{\mathcal{F}_g}$ for Π and \mathcal{F}_g and $\Theta_{\Pi_s}^{\mathcal{F}_s}$ for Π_s and \mathcal{F}_s are the same.

Definition 5 (Star-Mapping) Let $\mathcal{F}_g = \langle C, \mathcal{L} \rangle$ be a generalized factoring for $\Pi = \langle \mathcal{V}, \mathcal{A}, I, G \rangle$, and \mathcal{A}^G the global actions for Π and \mathcal{F}_g . $S(\Pi, \mathcal{F}_g) = \langle \Pi_s, \mathcal{F}_s \rangle$ is a star mapping constructed as follows: $\Pi_s := \langle \mathcal{V}_s, \mathcal{A}_s, I_s, G \rangle$ is the star-mapped planning task and $\mathcal{F}_s := \langle C_s, \mathcal{L} \rangle$ the star-mapped factoring for Π_s given Π and \mathcal{F}_g , where:

- $\mathcal{V}_s := \mathcal{V} \cup \{x\}$, $x \notin \mathcal{V}$, with $\mathcal{D}(x) := \{0\}$,
- $\mathcal{A}_s := \{a_s \mid a \in \mathcal{A}^G, \text{pre}(a_s) := \text{pre}(a), \text{cost}(a_s) := \text{cost}(a), \text{eff}(a_s) := \text{eff}(a) \cup \{x = 0\}\} \cup (\mathcal{A} \setminus \mathcal{A}^G)$,
- $I_s := I \cup \{x = 0\}$, and
- $C_s := C \cup \{x\}$.

In words, we map a generalized factoring \mathcal{F}_g for a task Π to a star factoring \mathcal{F}_s for a modified task Π_s by introducing a new state variable x with unary domain. Then x is added to the center factor C_s of \mathcal{F}_s and we add an auxiliary effect $\{x = 0\}$ to all global actions. It is easy to see that \mathcal{F}_s is indeed a valid star factoring for Π_s , where leaf-only actions $\mathcal{A}_{\emptyset}^L$ of a leaf $L \in \mathcal{L}$ affect only L and are preconditioned by $C \cup L$, as for the generalized factoring. All other actions, namely the global actions of \mathcal{F}_g , are made center actions for \mathcal{F}_s by adding the effect on x . Note that the set of plans for Π is exactly the same as for Π_s , i. e., solutions are not affected.

Given a decoupled state space $\Theta_{\Pi}^{\mathcal{F}}$ and a variable x , $\Theta_{\Pi}^{\mathcal{F}} \setminus x$ denotes the same decoupled state space except the variable x and its assignments are removed from all states and labels.

Theorem 1 (Correspondence) Let \mathcal{F}_g be a generalized factoring for Π , let $S(\Pi, \mathcal{F}_g) = \langle \Pi_s, \mathcal{F}_s \rangle$, and let $\Theta_{\Pi}^{\mathcal{F}_g}$ and $\Theta_{\Pi_s}^{\mathcal{F}_s}$ be decoupled state spaces of \mathcal{F}_g and \mathcal{F}_s , respectively. Then $\Theta_{\Pi}^{\mathcal{F}_g} = \Theta_{\Pi_s}^{\mathcal{F}_s} \setminus x$.

Proof Sketch: The key observations are that the set of global actions for \mathcal{F}_g coincides with the set of center actions for \mathcal{F}_s . So, the search branches over exactly the same action sequences. Next, the set of leaf-only actions for each $L \in \mathcal{L}$ is also the same for both factorings, so the pricing functions are exactly the same. Finally, the auxiliary center variable x and the additional effect $x = 0$ of center actions does not affect the behaviour of decoupled search in any way. \square

With Theorem 1, we can perform decoupled search with generalized factorings like with star factorings, and all properties of decoupled search with star factorings are preserved.

Corollary 1 Decoupled search with generalized factorings is sound, complete, and preserves optimality.

Characteristics of Generalized Factorings

Gnad, Poser, and Hoffmann (2017) observed that the so called *mobility* is an important property of factorings, closely related to the state-space reduction of the decoupled search. For a generalized factoring \mathcal{F}_g , a leaf $L \in \mathcal{L}$ is *mobile* if there exists a leaf-only action for L , i. e., $|\mathcal{A}_{\emptyset}^L| > 0$. A factoring is mobile if all its leaves are. Next, we show that a state-space reduction is possible *only* with mobile leaves.

Proposition 1 (Non-mobile Leaves) *Let Π be a planning task and $\mathcal{F}_g = \langle C, \mathcal{L} \rangle$ a generalized factoring for Π . If a leaf $L \in \mathcal{L}$ is not mobile, then in all decoupled states $s^{\mathcal{F}}$ reachable from the initial decoupled state $I^{\mathcal{F}}$ there exists exactly one $s^L \in S^L$ where $\text{prices}(s^{\mathcal{F}})[s^L] < \infty$.*

Proof: The claim is true in the initial state, since there exists no leaf-only action for L , so $\text{prices}(I^{\mathcal{F}})[I[L]] = 0$ and $\text{prices}(I^{\mathcal{F}})[s^L] = \infty$ for all $s^L \neq I[L]$. Let $s^{\mathcal{F}}$ be a successor of $I^{\mathcal{F}}$ via global action a^G . Then the leaf state $s^L = I[L][a^G]$ has a price of 0 in $s^{\mathcal{F}}$. Again, because there are no leaf-only actions of L , no other leaf state of L is reached in $s^{\mathcal{F}}$. This argument applies inductively to all decoupled states reachable from $I^{\mathcal{F}}$. \square

Consequently, a leaf $L \in \mathcal{L}$ for which there can only ever be a single reached leaf state can be merged into a new center factor $C \cup L$ without affecting the set of member states of any reachable decoupled state. In the extreme case where no leaf is mobile, every decoupled state has exactly one member state, so decoupled search degrades to explicit-state search.

While mobility can be seen as a qualitative property, prior work also tried to quantify the “amount of work a leaf can do on its own”. We adopt the definition of mobility of a factoring as the sum of the number of leaf-only actions $|\mathcal{A}_{\mathcal{L}}^L|$ of its leaves $L \in \mathcal{L}$. This definition allows us to reason about the reduction power of a factoring more accurately, since the number of leaf-only actions provides an estimate on how many leaf states can be reached in a single decoupled state.

Although the state-space reduction is exponential in the number of leaves, sometimes we can obtain stronger reduction by moving leaf variables into the center.

Proposition 2 *Let $\mathcal{F}_g = \langle C, \mathcal{L} \rangle$ be a generalized factoring for Π with mobility M . If there exist non-mobile leaves $\{L_1, \dots, L_n\} \subseteq \mathcal{L}$, then there exists a mobile factoring with $|\mathcal{F}_g| - n$ leaves with mobility $M' \geq M$.*

Proof: The mobility of a leaf $L \in \mathcal{L}$ equals the number of leaf-only actions of L , namely $|\mathcal{A}_{\mathcal{L}}^L|$. Thus, as L_1, \dots, L_n are not mobile, the factoring $\mathcal{F}'_g = \langle C \cup L_1 \cup \dots \cup L_n, \mathcal{L} \setminus \{L_1, \dots, L_n\} \rangle$ has at least the mobility of \mathcal{F}_g .

We get a strictly higher mobility if there exists a leaf action $a \in \mathcal{A}^L$ of an $L \in \mathcal{L}$ where, w.l.o.g. for L_1 , we have $\text{vars}(\text{eff}(a)) \subseteq L$ and $\text{vars}(\text{pre}(a)) \subseteq L \cup L_1$. Then a was a global action for \mathcal{F}_g , but is a leaf-only action for \mathcal{F}'_g . \square

Given the construction in the proof, it can even make sense to sacrifice a leaf with low mobility and merge it into the center, since this can increase the number of leaf-only actions for potentially many other leaves.

Relation to Other Forms of Factored Planning

Decoupled search on star topologies differs conceptually from other factored planning approaches by requiring a specific structure to be present in a planning task. With generalized factorings, this is no longer the case. Like other factored planning methods, this allows arbitrary variable partitions.

In localized factored planning (e.g. Amir and Engelhardt 2003; Fabre et al. 2010; Brafman and Domshlak 2006, 2008,

2013), the planning process performs a local search for the individual factors. Solutions to these sub-problems are coordinated by resolving cross-factor interactions using a global constraint-satisfaction problem. Hierarchical factored planning (e.g. Knoblock 1994; Kelareva et al. 2007; Wang and Williams 2015) refines top-level solutions for an increasing number of factors while moving down a hierarchy of decreasing level of abstraction, resolving inconsistencies by backtracking to higher levels. The crucial difference to decoupled search is that keeping global (center) and leaf actions separate allows to perform a monolithic search, just over the more complex decoupled state space. Thus, there is never the need to coordinate sub-solutions, or backtrack from partial plans that cannot be refined. This also allows us to use any standard search algorithm, planning heuristic, or pruning method, which are orthogonal to the decomposition performed by decoupled search (Torralba et al. 2016; Gnad et al. 2017; Gnad, Hoffmann, and Wehrle 2019).

Obtaining Generalized Factorings

Following prior work (Schmitt, Gnad, and Hoffmann 2019), we formulate the problem of finding a generalized factoring as an integer linear program (ILP) and describe several objective functions targeting different factoring properties.

First, we choose a set of candidates for leaves called *potential leaves*. Then we set constraints so that the resulting partition of variables is indeed a generalized factoring. Finally, we add more constraints ensuring that all leaves are mobile, as we have shown in Proposition 1 and 2 that non-mobile leaves do not contribute to the state-space reduction.

Before getting to the specifics of the ILP encoding, we define the notion of *action variable schemas* expressing causal dependencies between the center and leaves.

Definition 6 *Given an action $a \in \mathcal{A}$, an action variable schema (or a-schema for short), denoted by A_a , is a tuple $A_a = \langle \text{pre}(A_a), \text{eff}(A_a) \rangle$, where $\text{pre}(A_a) = \text{vars}(\text{pre}(a))$, and $\text{eff}(A_a) = \text{vars}(\text{eff}(a))$. The set of all a-schemas is denoted by $\mathcal{K} = \{A_a \mid a \in \mathcal{A}\}$. Given a set of variables $L \subseteq \mathcal{V}$, $\mathcal{K}(L) = \{A \in \mathcal{K} \mid \text{eff}(A) \subseteq L\}$ denotes a set of a-schemas affecting L . And given an a-schema $A \in \mathcal{K}$, $|A|$ denotes the number of actions with a-schema A .*

Besides the planning task Π , the ILP encoding depends on the choice of a set of potential leaves, $\Lambda \subseteq 2^{\mathcal{V}}$. We choose potential leaves as the sets of variables affected by at least one action, i.e., $\Lambda = \{\text{eff}(A) \mid A \in \mathcal{K}\}$. Clearly, a set of variables V such that $\text{eff}(A) \not\subseteq V$ for every $A \in \mathcal{K}$ can never be a mobile leaf, which makes our selection of Λ the set of minimal possible variable sets that can become mobile leaves. We experimented with supersets of these potential leaves, e.g., by combining pairs of intersecting or causally related leaves. However, this led to a significant increase in the size of the encoding and never paid off in practice.

Our ILP encoding uses the following binary variables:

- (i) X_L for every potential leaf $L \in \Lambda$; X_L is set to 1 when L becomes a leaf, and it is set to 0 otherwise.
- (ii) $Y_{L,A}$ for every potential leaf $L \in \Lambda$ and every a-schema $A \in \mathcal{K}(L)$. This variable is set to 1 when the actions $a \in \mathcal{A}$ with the a-schema A are mobile for L .

(iii) Z_v for every variable $v \in \mathcal{V}$. This variable is set to 1 if v becomes a center variable in the resulting factoring.

The X_L variables correspond to the actual choice of the variable partition that is made, $Y_{L,A}$ and Z_v are used to ensure that we obtain a proper factoring.

The following constraints ensure that the solution to our ILP is indeed a mobile generalized factoring. There are constraints $X_L + X_{L'} \leq 1$ for every $L, L' \in \Lambda$ s.t. $L \neq L'$ and $L \cap L' \neq \emptyset$, enforcing a partition of the leaf variables. Constraints $Z_v \leq 1 - X_L$ for every $L \in \Lambda$ and every $v \in L$, ensure that the center is disjoint with all leaves.

A necessary condition for a leaf L to be mobile is that there exists an action a such that $\text{vars}(\text{eff}(a)) \subseteq L$. So, it is enough to consider a-schemas A with $\text{eff}(A) \subseteq L$, i.e., the set $\mathcal{K}(L)$. For every potential leaf $L \in \Lambda$ we add the constraint:

$$X_L - \sum_{A \in \mathcal{K}(L)} Y_{L,A} \leq 0 \quad (1)$$

The next constraint serves two purposes. First, it ensures that any variable $Y_{L,A}$ is set to 1 only if the corresponding variable X_L is set to 1. Second, in combination with the previous constraint it enforces mobility of each leaf. The mobility of a leaf L requires that there exists at least one action a such that (a) $\text{vars}(\text{eff}(a)) \subseteq L$, and (b) every outside precondition of a (i.e., $\text{vars}(\text{pre}(a)) \setminus L$) is part of the center. As the constraint Eq. (1) takes care of (a), the following constraint ensures that (b) holds as well. For every potential leaf $L \in \Lambda$ and every action schema $A \in \mathcal{K}(L)$, we define the following constraint:

$$|\text{pre}(A) \setminus L| \cdot Y_{L,A} - \sum_{v \in \text{pre}(A) \setminus L} Z_v \leq X_L - Y_{L,A} \quad (2)$$

Finally, the following set of constraints makes sure that if a variable v is not part of any leaf, then it is part of the center: $1 - \sum_{L \in \Lambda, v \in L} X_L \leq Z_v$ for every variable $v \in \mathcal{V}$.

Every solution to the ILP described above is a generalized factoring with mobile leaves. Next, we discuss different optimization criteria that can be applied.

Towards Finding Good Factorings

Our main target is to obtain factorings that reduce the state-space the most, i. e., where each decoupled state has as many member states as possible. We consider several objective functions that correlate positively with this metric.

Maximize number of mobile leaves (L) This is desirable, as the number of member states is exponential in the number of leaves. In the ILP, the objective is to maximize $\sum X_L$.

Maximize leaf mobility (M) An increased leaf mobility contributes linearly to the state-space reduction. In the ILP, the objective is to maximize $\sum Y_{L,A} |A|$.

Similar to previous attempts to maximize mobility, though, this completely ignores how much each leaf contributes to the overall mobility. Ideally, we want a balanced contribution to maximize the search-space reduction. We capture this with the new notion of balanced leaf mobility.

Maximize balanced leaf mobility We define balanced leaf mobility as $\prod_{L \in \mathcal{L}} |\mathcal{A}_\varnothing^L|$ or equivalently $\sum_{L \in \mathcal{L}} \log(|\mathcal{A}_\varnothing^L|)$. That is, we multiply the number of mobile actions per leaf so that factorings with a balanced number of leaf-only actions per leaf are preferred (e.g., we prefer two leaves with 5 leaf-only actions each over one with 9 and another with only 1). In the ILP encoding, we can achieve this by introducing a binary variable $W_{L,C}$ for every combination $C \in 2^{\mathcal{K}(L)}$ of a-schemas from $\mathcal{K}(L)$ that can be mobile for a potential leaf L . We add constraints such that at most one $W_{L,C}$ can be set to 1 for every L . The objective is to maximize $\sum_{L \in \Lambda} \log(\sum_{A_i \in C} |A_i|) W_{L,C}$.

This encoding is not practical, though, as its size is exponential in the number of a-schemas of a potential leaf. Empirically, we observed an increase in the ILP size by one to two orders of magnitude on average, making it prohibitively large (statistics are available in the TR). Thus, we introduce an approximation that does not require additional variables.

Maximize (approximate) balanced leaf mobility (bM)

For every potential leaf L , we divide the set of affecting actions $\mathcal{A}^L := \{a \in \mathcal{A} \mid \text{vars}(\text{eff}(a)) \subseteq L\}$ into those whose precondition is contained in L , $\mathcal{A}_\top^L := \{a \in \mathcal{A}^L \mid \text{vars}(\text{pre}(a)) \subseteq L\}$, and the rest, $\mathcal{A}_{pre}^L := \mathcal{A}^L \setminus \mathcal{A}_\top^L$. Note that, if L is chosen as leaf factor, all actions in \mathcal{A}_\top^L will be leaf-only actions, regardless of what variables are in the center. The remaining actions \mathcal{A}_{pre}^L will be leaf-only actions iff $\text{vars}(\text{pre}(a)) \subseteq L \cup C$. Hence, the minimum and maximum balanced mobility of L are $M_{min}^L = \log(|\mathcal{A}_\top^L|)$ and $M_{max}^L = \log(|\mathcal{A}^L|)$, respectively. Moreover, let $\mathcal{A}_{pre}^L|_A$ denote the actions that have a-schema A and are in \mathcal{A}_{pre}^L . Our objective function maximizes the following expression:

$$\sum_{L \in \Lambda} M_{min}^L X_L + \sum_{A \in \mathcal{K}(L)} Y_{L,A} (M_{max}^L - M_{min}^L) \frac{|\mathcal{A}_{pre}^L|_A|}{|\mathcal{A}_{pre}^L|}$$

The intuition is that the approximation is correct at the extremes, whenever none or all actions in \mathcal{A}_{pre}^L are leaf-only actions. This is always the case if $|\mathcal{K}(L)| \leq 1$, which happens in many domains. In the middle, the function grows with the number of leaf-only actions, so that factorings with more leaf-only actions are always preferred.

Maximize average leaf fact flexibility (F) Prior definitions of mobility attempt to maximize the number of leaf-only actions, but ignore their effects. Yet, this is important because the number of member states of a decoupled state is exponential only in the number of leaves with more than one reached leaf state. If all leaf-only actions have exactly the same effect, then the number of leaf states with finite price is limited. Thus, we aim to maximize the percentage of actions affecting each leaf fact.

This can be encoded in the ILP by introducing a real variable $W_{v=d}$ with domain $[0, 1]$ for every fact $v = d$ of Π , where $v \in \mathcal{V}$ and $d \in \mathcal{D}(v)$. The value of each $W_{v=d}$ corresponds to the percentage of actions with effect $v = d$ that are leaf-only actions. To set the value of $W_{v=d}$, we introduce a constraint $W_{v=d} = (\sum_{L \in \Lambda, A \in \mathcal{K}(L)} |A_{v=d}| Y_{L,A}) / |\mathcal{A}_{v=d}|$, where $A_{v=d}$ is the subset of actions a with a-schema A

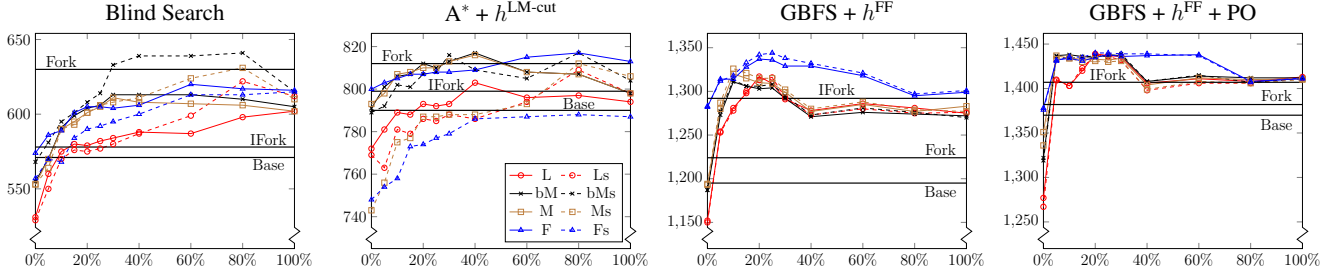


Figure 1: Coverage (on y -axis) of our new factoring strategies, with fallback to Base if a method abstains, as a function of the minimum leaf flexibility (x -axis). We also show the coverage of the baselines Base, Fork, and IFork.

where $\text{eff}(a)[v] = d$, and $\mathcal{A}_{v=d}$ is the set of all actions with that effect. The objective is to maximize $\sum W_{v=d}$. We remark that this ignores the distribution of facts across leaves.

Constraint on minimum leaf flexibility While the previous objectives aim to maximize overall mobility across all leaves, here we target a minimum amount of mobility per leaf. Given a parameter f_{\min} , we restrict the minimum flexibility of a potential leaf L by introducing an additional constraint per candidate leaf $L \in \Lambda$:

$$f_{\min} \cdot X_L \leq \sum_{A \in \mathcal{K}(L)} \frac{|A|Y_{L,A}}{|\{a \in \mathcal{A} \mid \text{eff}(a) \cap L \neq \emptyset\}|} \quad (3)$$

Here, to select L as a leaf, the ratio of leaf-only actions affecting L must be at least f_{\min} . As soon as $f_{\min} > 0$, this is strictly more constrained than Equation 1, requiring not only that the leaf is mobile but that at least some minimum percentage of actions affecting it are leaf-only.

Polynomial Test of Existence

The ILP can become very large for some tasks, so it is beneficial to check in advance if a non-trivial mobile factoring exists. There is an exact check that is quadratic in $|\mathcal{A}|$:

Proposition 3 (Existence of Mobile 2-Leaf Factoring)

There exists a mobile generalized factoring with at least two leaves iff there exist two distinct actions a_1, a_2 such that $\text{vars}(\text{eff}(a_1)) \cap \text{vars}(\text{eff}(a_2)) = \emptyset$, $\text{vars}(\text{eff}(a_1)) \cap \text{vars}(\text{pre}(a_2)) = \emptyset$, and $\text{vars}(\text{pre}(a_1)) \cap \text{vars}(\text{eff}(a_2)) = \emptyset$.

Proof Sketch: “ \Rightarrow ”: If such two actions exist, the following factoring fulfills the conditions: $\mathcal{F}_g = \langle C, \{L_1, L_2\} \rangle$, $L_1 = \text{vars}(\text{eff}(a_1))$, $L_2 = \text{vars}(\text{eff}(a_2))$, and $C = \mathcal{V} \setminus (L_1 \cup L_2)$.

“ \Leftarrow ”: Given a mobile factoring with leaves L_1, L_2 , for each of the two there exists an action a_i where $\text{vars}(\text{eff}(a_i)) \subseteq L_i$ and $\text{vars}(\text{pre}(a_i)) \subseteq L_i \cup C$. \square

With Proposition 3, we have an efficient way to check if a mobile 2-leaf factoring exists at all. We just need to iterate over all pairs of actions and check the stated conditions.

Experimental Evaluation

We implemented our factoring strategies in the decoupled search planner of Gnad and Hoffmann (2018), which is based on Fast Downward (Helmert 2006). Our experiments

were conducted using Lab (Seipp et al. 2017). We used all benchmarks of the International Planning Competitions (IPC) from 1998-2018 in the optimal as well as satisficing tracks, eliminating duplicate instances that appeared in several IPC iterations. For optimal planning, we report results for blind search and A^* with $h^{\text{LM-cut}}$ (Helmert and Domshlak 2009); in satisficing planning, we use greedy best-first search (GBFS) with the h^{FF} heuristic (Hoffmann and Nebel 2001), with and without preferred operator pruning (PO) using the common dual-queue approach (Richter and Helmert 2009). To compute these heuristics for decoupled states, prior work has introduced a task compilation that enables, in principle, the use of any heuristic. Via Theorem 1, this compilation is directly applicable to generalized factorings. The experiments were performed on a cluster of Intel E5-2660 machines running at 2.20 GHz with the runtime/memory limits of 30min/4GiB. Our code and data are publicly available (Gnad, Torralba, and Fišer 2022).

In our evaluation, we include the four ILP encodings that maximize the number of mobile leaves (L), the leaf mobility (M), the approximated balanced leaf mobility (bM), and the leaf-fact flexibility (F). For all these strategies, we run two variants. Our basic variant takes the set of all a-schema effects as potential leaves, the second one additionally considers the set of all strongly-connected components (SCC) in the causal graph (if there are at least two). We indicate the latter configurations with a postfix “s”, e. g., Ls considers all a-schema effects and SCCs. We perform a systematic study of the influence of the minimum leaf flexibility on the performance of decoupled search. To do so, we run each of the aforementioned eight variants with increasing f_{\min} : 0% – 30%, in steps of 5, and 40%, 60%, 80%, 100%. For non-zero flexibility, we indicate the used percentage in the configuration alias, e. g., bM60 maximizes balanced leaf mobility and enforces a minimum leaf flexibility of 60%.

We compare to explicit-state search (Base) and a set of existing factoring strategies for decoupled search as baselines. In particular, we run fork and inverted-fork factorings (Fork/IFork, in short Fo/IF), a strategy based on computing maximum independent sets of the causal graph (MIS), a greedy strategy (IA), and two ILP-based methods that produce *strict-star* (LPS), respectively general star factorings (LPG) (Gnad, Poser, and Hoffmann 2017; Schmitt, Gnad, and Hoffmann 2019). *Strict-star* factorings are a special case of star factorings where no action is allowed to touch, in pre-

		Optimal Planning: A* with LM-cut																	
	Base	Fo	IF	IA	MIS	LPS	LPG	O(o)	L	L80s	bM	bM40	M	M40	F	F80	O(n)	O(d)	O(a)
Tackled	1630	430	385	912	641	1125	1126	1243	1247	794	1233	847	1230	847	1236	846	1293	1321	1046
Not t. before	-	-	-	-	-	-	-	-	77	1	76	3	76	3	76	3	78	78	44
# fork	-	430	1	207	244	225	225	312	166	245	207	227	223	225	226	245	241	313	243
# inv-fork	-	0	384	136	46	137	57	194	30	193	30	213	35	215	85	229	107	178	171
# strict-star	-	0	0	569	351	763	274	620	40	196	35	171	22	153	26	158	79	531	418
# star	-	0	0	0	0	0	570	117	217	6	240	60	229	53	261	42	259	129	93
# generalized	-	0	0	0	0	0	0	0	794	154	721	176	721	201	638	172	607	170	121
$C = \emptyset$	-	1	1	0	0	0	0	0	1	39	1	0	1	0	3	0	2	0	0
Solved	790	263	132	471	397	549	568	637	584	415	596	470	598	470	606	471	673	681	851
+Base	-	812	799	791	802	812	811	836	772	809	791	817	793	816	800	817	838	840	851
Solved\Base	-	28	13	30	21	35	39	55	19	36	35	35	37	34	37	35	60	61	61
Base\Solved	-	6	4	29	9	13	18	9	37	17	34	8	34	8	27	8	12	11	0
		Satisficing Planning: GBFS using FF and preferred-operator pruning																	
	Base	Fo	IF	IA	MIS	LPS	LPG	O(o)	L	L20	bM	bM30s	M	M5	F	F20s	O(n)	O(d)	O(a)
Tackled	1686	437	403	897	654	1173	1171	1309	1280	916	1270	896	1255	1093	1271	991	1335	1374	898
Not t. before	-	-	-	-	-	-	-	-	65	2	64	2	60	27	64	13	65	65	25
# fork	-	437	1	207	244	225	225	318	166	166	207	240	223	223	226	255	214	303	128
# inv-fork	-	0	402	96	46	140	41	219	9	143	19	191	26	121	97	156	166	183	169
# strict-star	-	0	0	594	364	808	277	587	32	222	51	176	23	120	31	137	156	435	251
# star	-	0	0	0	0	0	628	185	220	88	258	84	247	226	302	176	218	197	164
# generalized	-	0	0	0	0	0	0	0	853	297	735	205	736	403	615	267	581	256	186
$C = \emptyset$	-	1	1	0	0	0	0	1	1	0	1	2	1	0	1	2	0	1	0
Solved	1370	427	366	775	574	996	993	1167	1006	845	1046	832	1070	985	1102	899	1254	1279	1499
+Base	-	1382	1407	1383	1383	1361	1343	1440	1277	1439	1322	1439	1351	1437	1378	1440	1490	1493	1499
Solved\Base	-	12	37	35	19	43	79	104	35	82	66	81	71	98	90	91	128	129	129
Base\Solved	-	0	0	22	5	51	105	34	127	13	113	12	89	30	81	21	8	6	0

Table 1: Results summary on optimal and satisficing planning. See text for detailed explanations.

condition or effect, more than one leaf. For the ILP-based strategies, we include two variants in our evaluation, one that maximizes the number of mobile leaves, and another that maximizes leaf mobility. We only report data for the variants that maximize the mobility, since these give consistently better results. All baseline strategies (except the two that maximize mobility) have in common that they return the factoring with the maximum number of mobile leaves among the respective subset of considered factorings. For all factoring methods, we only consider candidate leaves with size (product of variable domain sizes) below 2^{32} .

Like prior work on decoupled search, our methods *abstain* from solving a task if the generated factoring has less than two leaves. We say that a method *tackles* an instance if it does not abstain. We impose a runtime limit of 30s on the factoring process. On abstained instances, we sometimes (in particular in Figure 1) report the performance of the explicit-state baseline, because otherwise the underlying instance set would vary for each method. In this case, we consider an instance solved if the factoring method abstains or times out, and the baseline solves the instance in the remaining time.

In Figure 1, we show results of our minimum-leaf-flexibility investigation. For each of the eight base variants, we report the coverage (number of solved instances) as a function of the minimum flexibility for all four search settings. In addition to our new strategies, we also include the coverage of Base, Fork, and IFork for comparison. In general, the graphs confirm that the minimum flexibility has a significant impact on performance. For all methods, it is ben-

eficial to impose at least a mild minimum of around 5%; and performance is worst without any restriction. Requiring a flexibility of 100% almost never results in the highest coverage, either. There usually is a sweet-spot between the two extremes that yields the best performance. For satisficing planning, the data indicates a maximum coverage between 20% and 40%, with a peak around 25% for most methods. Adding the causal-graph SCCs does not lead to significant changes in satisficing planning. The SCCs do not seem to increase the number of alternative factorings that result in better performance in many cases. For optimal planning, the best coverage is achieved with higher minimum flexibility, between 40% and 80% for all but one configuration: blind search with L has its maximum at 100%. Adding SCC has a mixed impact, it results in significantly higher coverage for some objective functions, but can be detrimental, too.

Comparing the objective functions, there is no clear winner. For blind search bMs performs best, for A* with h^{LM-cut} it is bM and F, and for both satisficing search types it is F. Still, in all settings there are several configurations that outperform even the best previous method (which is either Fork or IFork), and most methods are significantly stronger than Base at the minimum flexibility where their coverage peaks.

In Table 1 we summarize the results for a selection of our new factoring strategies for A* search with h^{LM-cut} , and GBFS with h^{FF} and PO. We show Base and all prior methods. From our strategies, we include the basic variant without enforced minimum flexibility and the best sub-configuration for all four objective functions. In addition,




Figure 2: Distribution of different factorings per instance found by the 8 old and 88 new factoring methods.

we show results for oracle configurations (Oracle(x), short $O(x)$), where $x \in \{old, new, dec, all\}$ denotes the set of configurations that are considered by the oracle: “old” are all previous strategies, “new” are the ones presented in this paper, “dec” considers both, and “all” includes Base as well. An oracle is a simulated best-case combination of all considered methods, which picks, per instance, the factoring that results in the lowest search time. For instances not solved by any method, it picks any of the methods that tackle the instance, and abstains only if all methods do.

The overall conclusions are similar for both settings. First, we observe that generalized factorings have opened up many new possibilities and they can tackle a lot more instances, up to 78 for $O(n)$. This is largely due to generalized factorings, as confirmed by the second part of the table, which reports the kind of factorings found by each method. While, e. g., Fork and IFork obviously always return fork/inverted-fork factorings, IA, MIS, LPS, and LPG produce a larger variety of star topologies. Our new strategies make great use of generalized no-star factorings, sometimes with empty center.

But are those new possibilities any good? Definitely, at least in some cases. Comparing the factorings preferred by the oracles, it turns out that strict-star factorings often yield the best performance. For both search settings these are by far preferred by Oracle(all), but generalized factorings are also performing better than all other factorings and the baseline in 121 instances for optimal and 186 for satisficing planning. The “+Base” row reveals that our best new methods outperform all prior ones, especially in satisficing planning. Some configurations, however, perform worse than Base. Indeed, the best performing variants do not tackle that many new instances, showing that minimum leaf flexibility is a good criterion to decide when to use decoupled search.

Figure 2 shows that there are many planning tasks for which our methods return a significant number of different factorings. The main peak is at 2 factorings per instance, but there exist tasks for which up to 110 different factorings were generated. Across both benchmark sets, with a total of 2330 instances, there are only 494 for which no factoring was computed, out of which in 340 no factoring with two mobile leaves is provably possible by Proposition 3.

Finally, in Figure 3 we report per-instance runtime results, comparing Oracle(new) to Base and Oracle(old). Every point below the diagonal indicates an instance where one of the new factoring strategies performs better than Base or than *every* of the existing methods. For blind search, the results look worst. Here, Base is often better than the new

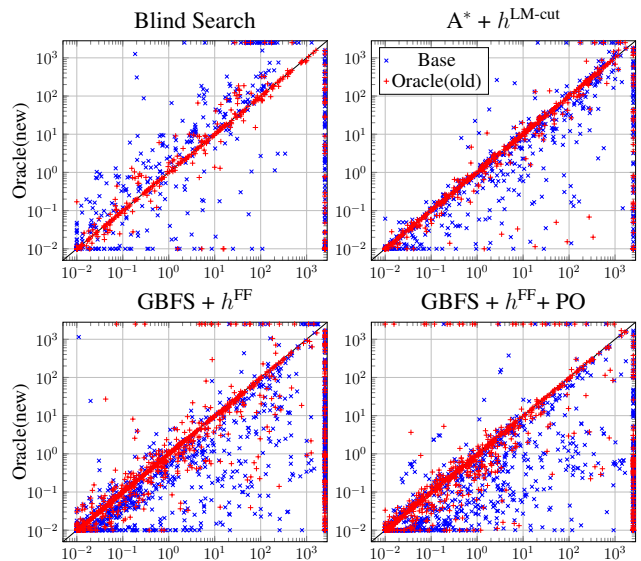


Figure 3: Per-instance runtime comparison for Oracle(new) (on y -axis) vs. Base and Oracle(old) on the x -axis.

methods, though Oracle(new) actually solves 59 more instances. The old strategies perform close to the new ones, indicating that the oracles probably often chose the same factorings. Still, there are instances where the new strategies result in a strong performance, sometimes several orders of magnitude faster than Base, and even than the best old method. For the remaining three search settings, the results look much more favorable for Oracle(new). There are many cases where the performance is a lot better than Base and Oracle(old), and it happens only very rarely that the best of the new strategies performs a lot worse than both.

Conclusion

We have overcome the structural requirements of decoupled search and extended its applicability to planning tasks that could not be tackled before. The former restriction to star topologies is obsolete, and we can now decompose planning tasks very flexibly by partitioning the state variables in an arbitrary way. We proved the correctness of the novel generalized factorings by making a connection to the existing star factorings. Thereby, all properties of decoupled search using star factorings are inherited, most importantly soundness, completeness, and optimality preservation.

We introduced a factoring strategy based on integer linear programming that is capable of producing generalized factorings, while being adaptable with respect to the properties of the factoring that should be optimized. Our experimental evaluation showed the effectiveness of our strategies and illustrated the benefit of the new freedom to decompose planning tasks for decoupled search. Given the vast space of possible factorings, however, the question of what the best decomposition for a given task is remains open. We believe this to be an important research topic for future work, that can further enhance the understanding, but also the performance of decoupled search.

Acknowledgements

Daniel Gnad was supported by the German Research Foundation (DFG), under grant HO 2169/6-2, “Star-Topology Decoupled State Space Search”. Jörg Hoffmann’s research group has received support by DFG grant 389792660 as part of TRR 248 (see perspicuous-computing.science).

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⁺ Planning. *Computational Intelligence*, 11(4): 625–655.
- Brafman, R.; 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.; and Domshlak, C. 2013. On the Complexity of Planning for Agent Teams and Its Implications for Single Agent Planning. *Artificial Intelligence*, 198: 52–71.
- Brafman, R. I.; and Domshlak, C. 2008. From One to Many: Planning for Loosely Coupled Multi-Agent Systems. In Rintanen, J.; Nebel, B.; Beck, J. C.; and Hansen, E., eds., *Proceedings of the 18th International Conference on Automated Planning and Scheduling (ICAPS’08)*, 28–35. AAAI Press.
- 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.
- Gnad, D.; and Hoffmann, J. 2018. Star-Topology Decoupled State Space Search. *Artificial Intelligence*, 257: 24 – 60.
- Gnad, D.; Hoffmann, J.; and Wehrle, M. 2019. Strong Stubborn Set Pruning for Star-Topology Decoupled State Space Search. *Journal of Artificial Intelligence Research*, 65: 343–392.
- Gnad, D.; Poser, V.; and Hoffmann, J. 2017. Beyond Forks: Finding and Ranking Star Factorings for Decoupled Search. In Sierra, C., ed., *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI’17)*, 4310–4316. AAAI Press/IJCAI.
- Gnad, D.; Torralba, Á.; and Fišer, D. 2022. Technical report and data of the ICAPS’22 paper “Beyond Stars - Generalized Topologies for Decoupled Search”. <https://doi.org/10.5281/zenodo.6384091>.
- Gnad, D.; Torralba, Á.; Shleyfman, A.; and Hoffmann, J. 2017. Symmetry Breaking in Star-Topology Decoupled Search. In *Proceedings of the 27th International Conference on Automated Planning and Scheduling (ICAPS’17)*, 125–134. AAAI Press.
- Helmert, M. 2006. The Fast Downward Planning System. *Journal of Artificial Intelligence Research*, 26: 191–246.
- 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.
- Hoffmann, J.; and Nebel, B. 2001. The FF Planning System: Fast Plan Generation Through Heuristic Search. *Journal of Artificial Intelligence Research*, 14: 253–302.
- 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.
- Richter, S.; and Helmert, M. 2009. Preferred Operators and Deferred Evaluation in Satisficing Planning. In Gerevini, A.; Howe, A.; Cesta, A.; and Refanidis, I., eds., *Proceedings of the 19th International Conference on Automated Planning and Scheduling (ICAPS’09)*, 273–280. AAAI Press.
- Schmitt, F.; Gnad, D.; and Hoffmann, J. 2019. Advanced Factoring Strategies for Decoupled Search using Linear Programming. In *Proceedings of the 29th International Conference on Automated Planning and Scheduling (ICAPS’19)*. AAAI Press.
- Seipp, J.; Pommerening, F.; Sievers, S.; and Helmert, M. 2017. Downward Lab. <https://doi.org/10.5281/zenodo.790461>.
- Torralba, Á.; Gnad, D.; Dubbert, P.; and Hoffmann, J. 2016. On State-Dominance Criteria in Fork-Decoupled Search. In Kambhampati, S., ed., *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI’16)*, 3265–3271. AAAI Press/IJCAI.
- Wang, D.; and Williams, B. C. 2015. tBurton: A Divide and Conquer Temporal Planner. In Bonet, B.; and Koenig, S., eds., *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI’15)*, 3409–3417. AAAI Press.