# A Novel Iterative Approach to Top- $k$ Planning 

Michael Katz, Shirin Sohrabi, Octavian Udrea<br>IBM Research<br>Yorktown Heights, NY, USA

Dominik Winterer*<br>University of Freiburg<br>Germany


#### Abstract

While cost-optimal planning aims at finding one best quality plan, top-k planning deals with finding a set of solutions, such that no better quality solution exists outside that set. We propose a novel iterative approach to top- k planning, exploiting any cost-optimal planner and reformulating a planning task to forbid exactly the given set of solutions. In addition, to compare to existing approaches to finding top-k solutions, we implement the $K^{*}$ algorithm in an existing PDDL planner, creating the first $K^{*}$ based solver for PDDL planning tasks. We empirically show that the iterative approach performs better for up to a large required size solution sets (thousands), while $K^{*}$ based approach excels on extremely large ones.


## Introduction

Cost-optimal planning is the problem of finding one goalachieving sequence of actions or a plan of a minimal summed up cost. Such plans are of interest in many applications, where the quality of solution is of extreme importance. In particular, in problems where preferences or likelihoods are encoded as cost of the actions (e.g., (Keyder and Geffner 2009; Sohrabi, Udrea, and Riabov 2013)). In many cases, one best solution is not sufficient, and the desire is to obtain a set of solutions of a high quality. Top- $k$ planning is one way of obtaining such a set. It is the problem of finding a set of solutions of size $k$, such that no better solution exists outside the set. While finding a set of plans is motivated by several applications including plan repair, often within the context of diverse planning (e.g., (Fox et al. 2006; Bryce 2014; Coman and Muñoz-Avila 2011)), other applications including risk management, hypothesis and explanation generation require to focus on high-quality plans instead, as the underlying approach of plan-recognition-as-planning works best with such plans (Sohrabi, Udrea, and Riabov 2013; Sohrabi, Riabov, and Udrea 2016; Sohrabi et al. 2018).

Recently, Sohrabi et al. (2016) proposed the use of a $k$-shortest path algorithm called $K^{*}$ by Aljazzar and Leue (2011) to address the top- $k$ planning problem. Their experimental results have shown that the planning time to compute top-k plans is comparable to finding the optimal plan; in addition, the quality of the solutions found is much

[^0]higher compared to the use of the diverse version of the LPG planner (Nguyen et al. 2012) which computes a set of diverse plans. This is not surprising, as the focus of diverse planning is not on plan quality, making top-k planning the more suitable approach to the applications above.

One major limitation of $K^{*}$ is the consistency requirement for the heuristic function. This poses a restriction on the usability of the approach, disallowing the use of one of the most successful heuristics to date (Helmert and Domshlak 2009). Further, all plans are found almost simultaneously, requiring long time until the first solution is found. Moreover, the existing implementation of $K^{*}$ for planning is in a planner that supports Stream Processing Planning Language (SPPL) (Riabov and Liu 2005), and not in PDDL. Thus, there are no experimental results for $K^{*}$ on the standard planning benchmarks. To alleviate these problems, Riabov, Sohrabi, and Udrea (2014) suggest an iterative approach to top-k planning, given a solution to a planning task, encode a set of new planning tasks which, cumulatively, preserve all solutions of the original task, except for the given one. Then, a search is performed on a tree of reformulations, invoking an existing planner in each node. As the number of successors of each node is the number of actions in the found plan, the clear down side of such an approach is the large number of invocations of the underlying planner. On the positive side, the approach exhibits an anytime behavior, with the first plan found rather quickly.

In this work, we propose an alternative iterative approach to top-k planning, finding additional solutions by reformulating the planning task at hand into a single task, preserving all solutions except for the given one. To this end, we formally define such reformulations and present one such instance. We then suggest additional ways of deriving solutions from the previously found ones, further reducing the number of invocations of the underlying cost-optimal planner. For that, we extend our reformulation to forbid multiple plans at once, alleviating the increase in task formulation size. Additionally, we implement the $K^{*}$ algorithm on top of the Fast Downward planning system (Helmert 2006), allowing us to perform the long overdue experimental evaluation on International Planning Competition (IPC) benchmarks, comparing our new iterative approach to $K^{*}$. Our results show that $K^{*}$ works better for extremely large values of $k$, while the iterative approach excels otherwise.

## Background

We consider classical planning tasks as captured by the wellknown SAS ${ }^{+}$formalism (Bäckström and Nebel 1995), extended with action costs. In such a planning task $\Pi=$ $\left\langle\mathcal{V}, \mathcal{O}, s_{0}, s_{\star}, \operatorname{cost}\right\rangle, \mathcal{V}$ is a finite set of finite-domain state variables. Each variable $v \in \mathcal{V}$ is associated with a finite domain $\mathcal{D}(v)$ of variable values. A partial assignment $p$ maps a subset of variables $\operatorname{var} s(p) \subseteq \mathcal{V}$ to values in their domains. For a variable $v \in \mathcal{V}$ and partial assignment $p$, the value of $v$ in $p$ is denoted by $p[v]$ if $v \in \operatorname{vars}(p)$ and we say $p[v]$ is undefined if $v \notin \operatorname{vars}(p)$. A partial assignment $s$ with $\operatorname{var} s(s)=\mathcal{V}$, is called a state. State $s$ is consistent with partial assignment $p$ if they agree on all variables in $\operatorname{vars}(p)$, shortly denoted by $p \subseteq s$. The product $\mathcal{S}=\prod_{v \in \mathcal{V}} \mathcal{D}(v)$ is called the state space of planning task $\Pi$. The state $s_{0}$ is called initial state of $\Pi$ and the partial assignment $s_{\star}$ is called the goal of $\Pi$. A state $s$ is called a goal state if $s_{\star} \subseteq s$ and the set of all goal states is denoted by $\mathcal{S}_{s_{\star}}$. The finite set $\mathcal{O}$ is a set of actions, each action is a pair $\langle p r e$, eff $\rangle$ where pre is a partial assignment called precondition and eff is a partial assignment called effect. Further, each action $o$ has an associated natural number cost (o), called cost. An action $o=\langle p r e, e f f\rangle$ is applicable in state $s$ if pre $\subseteq s$. Applying action $o$ in state $s$ results in a state denoted by $s \llbracket o \rrbracket$ where $s \llbracket o \rrbracket[v]=e f f[v]$ for all $v \in \operatorname{vars}(e f f)$ and $=s \llbracket o \rrbracket[v]=s[v]$ for all other variables. An action sequence $\pi=\left\langle o_{1}, \cdots, o_{n}\right\rangle$ is applicable in state $s$ if there are states $s_{0}, \cdots, s_{n}$ such that $o_{i}$ is applicable in $s_{i-1}$ and $s_{i-1} \llbracket o_{i} \rrbracket=s_{i}$ for $0 \leq i \leq n$. We denote $s_{n}$ by $s \llbracket \pi \rrbracket$. For convenience we often write $o_{1}, \cdots, o_{n}$ instead of $\left\langle o_{1}, \cdots, o_{n}\right\rangle$. An action sequence with $s \llbracket \pi \rrbracket \in \mathcal{S}_{s_{\star}}$ is called a plan. The cost of a plan $\pi$, denoted by $\operatorname{cost}(\pi)$ is the summed cost of the actions in the plan. For a planning task $\Pi=\left\langle\mathcal{V}, \mathcal{O}, s_{0}, s_{\star}\right.$, cost $\rangle$, the set of all plans is denoted by $\mathcal{P}_{\Pi}$. A plan $\pi$ is optimal if its cost is minimal among all plans in $\mathcal{P}_{\Pi}$. Let $\Pi, \Pi^{\prime}$ be two planning tasks with actions $\mathcal{O}$ and $\mathcal{O}^{\prime}$ respectively. A total function $r: \mathcal{O}^{\prime} \mapsto \mathcal{O}$ is called an action mapping. We call the mapping $r^{\prime}: \mathcal{O}^{\prime n} \mapsto \mathcal{O}^{n}$ for a non-negative integer $n$, the extension of action mapping $r$ if for every action sequence $\pi^{\prime}=\left\langle o_{1}^{\prime}, \cdots, o_{n}^{\prime}\right\rangle$ with $o_{i} \in \mathcal{O}^{\prime}$ for $0 \leq i \leq n$ there is an action sequence $\pi=\left\langle r\left(o_{1}^{\prime}\right), \cdots, r\left(o_{n}^{\prime}\right)\right\rangle$ and $r^{\prime}\left(\pi^{\prime}\right)=\pi$.

Having provided the necessary background on classical planning, let us define the top-k planning problem (Sohrabi et al. 2016) that is the subject of interest in this paper.

Definition 1 (top- $k$ planning problem) Given a planning task $\Pi=\left\langle\mathcal{V}, \mathcal{O}, s_{0}, s_{\star}\right.$, cost $\rangle$ and a natural number $k$, find a set of plans $P \subseteq \mathcal{P}_{\Pi}$ such that:
(i) for all plans $\pi \in P$, if there exists a plan $\pi^{\prime}$ for $\Pi$ with $\operatorname{cost}\left(\pi^{\prime}\right)<\operatorname{cost}(\pi)$, then $\pi^{\prime} \in P$, and
(ii) $|P| \leq k$, with $|P|<k$ implying $P=\mathcal{P}_{\Pi}$.

An instance of the top- $k$ planning problem $\langle\Pi, k\rangle$, is called solvable if $|P|=k$ and unsolvable if $|P|<k$.

The objective of top- $k$ planning is finding the $k$-plans of lowest costs for a planning task $\Pi$ and thus optimal planning is the special case of top-1 planning. Our approach to solving the top- $k$ planning problem is the topic of the next section.

## Repeatedly Forbidding Plans

We now introduce an algorithmic scheme for solving a top- $k$ planning problem $\langle\Pi, k\rangle$. The key idea of this scheme is an iteration of the following steps: (1) Find an optimal plan $\pi$ for planning task $\Pi$. (2) Reformulate $\Pi$ to a planning task $\Pi^{\prime}$ with the same set of plans but excluding $\pi$. (3) Repeat (1) with $\Pi=\Pi^{\prime}$ and $\pi=\pi^{\prime}$ unless either $k$ solutions have been found or the $\Pi^{\prime}$ is provably unsolvable. The scheme is summarized below.

```
Algorithm 1: IterativeTopK \((\langle\Pi, k\rangle)\)
    \(P \leftarrow \emptyset, \quad \pi \leftarrow \emptyset, \quad r \leftarrow i d\)
    while \(|P|<k\) and no failure occurs do
        \(\Pi, r^{\prime} \leftarrow \operatorname{PLANFORBIDREFORMULATION}(\Pi, \pi)\)
        \(r \leftarrow r \circ r^{\prime}\)
        \(\pi \leftarrow\) GetOptimalPlan \((\Pi)\)
        if GetOptimalPlan \((\Pi)\) reports unsolvable
        then
            return \(P\)
        end
        \(P \leftarrow P \cup\{r(\pi)\}\)
    end
    return \(P\)
```

The key challenge of turning this algorithmic scheme into a concrete algorithm lies in the reformulation. More concretely, in how to forbid exactly one plan while at the same time preserving every other plan. Next, we state a necessary condition for such a reformulation.

Definition 2 (plan forbidding reformulation) Let $\Pi$ be a planning task over actions $\mathcal{O}$ and $\pi$ be a plan for $\Pi$. A planning task $\Pi^{\prime}$ over actions $\mathcal{O}^{\prime}$ is a plan forbidding reformulation of $\Pi$ if there exists an action mapping $r: \mathcal{O}^{\prime} \mapsto \mathcal{O}$ and its extension $r^{\prime}$ such that $\pi^{\prime}$ is a plan for $\Pi^{\prime}$ iff $r^{\prime}\left(\pi^{\prime}\right)$ is a plan for $\Pi$ with $\operatorname{cost}{ }^{\prime}\left(\pi^{\prime}\right)=\operatorname{cost}\left(r^{\prime}\left(\pi^{\prime}\right)\right)$ and $r^{\prime}\left(\pi^{\prime}\right) \neq \pi$.

We proceed by proving soundness and completeness of our scheme given some plan forbidding reformulation. Later we will show how to construct a specific such reformulation.

Theorem 1 The algorithm IterativeTopK is sound and complete.

Proof: Let $P$ be the set of plans returned by the algorithm, $\pi_{1} \ldots \pi_{m}$ be the ordering in which the plans were found and let $\Pi_{1} \ldots \Pi_{m}$ be the sequence of task reformulations constructed by the algorithm such that $\pi_{i}$ is the optimal plan found for $\Pi_{i}$. Then $\operatorname{cost}\left(\pi_{1}\right) \leq \operatorname{cost}\left(\pi_{2}\right) \leq$ $\cdots \leq \operatorname{cost}\left(\pi_{m}\right)$. If there exists a plan $\pi$ for $\Pi$ such that $\operatorname{cost}(\pi)<\operatorname{cost}\left(\pi_{i}\right)$ for some (assume W.L.O.G. smallest such) $i$, and $\pi \neq \pi_{j}$ for $j<i$, then from Definition 2, there exists a plan $\pi^{\prime}$ for $\Pi_{i}$ such that $r\left(\pi^{\prime}\right)=\pi$ and $\operatorname{cost}\left(\pi^{\prime}\right)=\operatorname{cost}(\pi)<\operatorname{cost}\left(\pi_{i}\right)$, contradicting the optimality of the plan $\pi_{i}$ for $\Pi_{i}$. If $m<k$, then the planning task $\Pi_{m+1}$ is unsolvable, and therefore no other solution exists for $\Pi_{m}$ and thus for $\Pi$.

Theorem 1 opens the door to a novel family of algorithms for top- $k$ planning. Yet, what is missing is to provide a concrete plan forbidding reformulation $\Pi_{\pi}^{-}$for planning task $\Pi$ and a plan $\pi \in \mathcal{P}_{\Pi}$. As mentioned earlier the idea here is that planning task $\Pi_{\pi}^{-}$forbids the sequence of actions $\pi$ from being a plan, accepting all other plans of $\Pi$.

Definition 3 Let $\Pi=\left\langle\mathcal{V}, \mathcal{O}, s_{0}, s_{\star}\right.$, cost $\rangle$ be a planning task and $\pi=o_{1} \ldots o_{n}$ be some plan. The task $\Pi_{\pi}^{-}=$ $\left\langle\mathcal{V}^{\prime}, \mathcal{O}^{\prime}, s_{0}^{\prime}, s_{\star}^{\prime}\right.$, cost $\rangle$ is defined as follows.

- $\mathcal{V}^{\prime}=\mathcal{V} \cup\left\{\bar{v}, \bar{v}_{0}, \ldots \bar{v}_{n}\right\}$, with $\bar{v}_{i}$ being binary variables,
- $\mathcal{O}^{\prime}=\left\{o^{e} \mid o \in \mathcal{O}, o \notin \pi\right\} \cup\left\{o^{1}, o^{2} \mid o \in \pi\right\} \cup \bigcup_{i=1}^{n}\left\{o_{i}^{3}\right\}$, where

$$
\begin{aligned}
& o^{e}=\langle\text { pre }(o), \text { eff }(o) \cup\{\langle\bar{v}, 0\rangle\}\rangle, \\
& o^{1}=\langle\operatorname{pre}(o) \cup\{\langle\bar{v}, 0\rangle\}, \text { eff }(o)\rangle, \\
& o^{2}=\left\langle\operatorname{pre}(o) \cup\{\langle\bar{v}, 1\rangle\} \cup \bigcup_{o=o_{i}, 1 \leq i \leq n}\left\{\left\langle\bar{v}_{i-1}, 0\right\rangle\right\},\right. \\
& \operatorname{eff}(o) \cup\{\langle\bar{v}, 0\rangle\}\rangle, \\
& o_{i}^{3}=\left\langle\operatorname{pre}\left(o_{i}\right) \cup\left\{\langle\bar{v}, 1\rangle,\left\langle\bar{v}_{i-1}, 1\right\rangle\right\},\right. \\
& \left.\qquad \quad \text { eff }\left(o_{i}\right) \cup\left\{\left\langle\bar{v}_{i-1}, 0\right\rangle,\left\langle\bar{v}_{i}, 1\right\rangle\right\}\right\rangle, \text { and } \\
& \\
& \operatorname{cost}^{\prime}\left(o^{e}\right)=\operatorname{cost}^{\prime}\left(o^{1}\right)=\operatorname{cost}^{\prime}\left(o^{2}\right)=\operatorname{cost}^{\prime}\left(o^{3}\right)=\operatorname{cost}(o), \\
& \bullet \quad s_{0}^{\prime}[v]=s_{0}[v] \text { for all } v \in \mathcal{V}, s_{0}^{\prime}[\bar{v}]=1, s_{0}^{\prime}\left[\bar{v}_{0}\right]=1, \text { and } \\
& s_{0}^{\prime}\left[\bar{v}_{i}\right]=0 \text { for all } 1 \leq i \leq n, \text { and } \\
& \bullet \\
& s_{\star}^{\prime}[v]=s_{\star}[v] \text { for all } v \in \mathcal{V} \text { s.t. } s_{\star}[v] \text { defined, and } s_{\star}^{\prime}[\bar{v}]=0 .
\end{aligned}
$$

Let us explain the semantics of the reformulation in Definition 3. The variable $\bar{v}$ starts from the value 1 and switches to 0 when the sequence of actions applied is not a prefix of $\pi$. Once a value 0 is reached indicating a deviating from plan $\pi$, it cannot be switched back to 1 . Variables $\bar{v}_{0}, \ldots, \bar{v}_{n}$ encode the progress along the plan $\pi$, before deviating from it. The actions $o^{1}$ and $o^{2}$ are the copies of actions in $\pi$ for the cases when $\pi$ is already discarded from consideration (variable $\bar{v}$ has switched its value to 0 ) and for discarding $\pi$ from consideration (switching $\bar{v}$ to 0 ). In case there are multiple appearances of an action $o$ on a plan, $o^{2}$ requires all the corresponding variables $\bar{v}_{i-1}$ for each $o_{i}=o$ to have the value 0 . Thus, $o^{2}$ will be applicable only in states that are different from the states resulting from applying a prefix of $\pi$ up to some $o$ 's representative in the initial state. $o_{i}^{3}$ are copies of actions along $\pi$, with a separate copy for each appearance of the same action $o$. These actions are responsible for following the sequence $\pi$ and are applicable only while the sequence is still followed.

In the following theorem, we show that the planning task in Definitions 3 is indeed a plan-forbidding reformulation.

Theorem 2 Let $\Pi$ be a planning task and $\pi$ be its plan. The task $\Pi_{\pi}^{-}$is a plan forbidding reformulation of $\Pi$ and $\pi$.

Proof: Let $r: \mathcal{O}^{\prime} \mapsto \mathcal{O}$ be the mapping defined by $r\left(o^{e}\right)=$ $o$ and $r\left(o^{1}\right)=r\left(o^{2}\right)=r\left(o_{i}^{3}\right)=o$ for all $1 \leq i \leq n$. Note that $\Pi_{\pi}^{-}$restricted to the variables $\mathcal{V}$ equals to the task $\Pi$, modulo the three equal instances of the actions in $\pi$. Thus, for each plan $\pi^{\prime}$ for $\Pi_{\pi}^{-}, r(\pi)$ is a plan for $\Pi$.

For the second direction, since for each $o \in \pi$ at most one of the actions $o^{1}, o^{2}, o^{3}$ is applicable in each state $s$ of $\Pi_{\pi}^{-}$, given a sequence of actions $\rho$ applicable in the initial state of $\Pi$, it can be mapped to an applicable in the initial state of $\Pi_{\pi}^{-}$sequence of action $\rho^{\prime}$ such that $r\left(\rho^{\prime}\right)=\rho$, by choosing in each state the relevant representative out of $o^{e}, o^{1}, o^{2}$, and $o^{3}$. In other words, $r$ restricted to applicable in the initial state sequences of actions is invertible, and we denote its inverse mapping described above by $r^{-1}$.

First, let $\pi^{\prime}=r^{-1}(\pi)$ be the inverse of the plan $\pi=$ $o_{1} \ldots o_{n}$ for $\Pi$. Then $\pi^{\prime}=o_{1}^{3} \ldots o_{n}^{3}$, since at step $i$ we have $\bar{v}_{0}=1$ and $\bar{v}_{i-1}=1$. Thus, after applying $\pi^{\prime}$, the value of the variable $\bar{v}$ remains 1 , and thus $\pi^{\prime}$ is not a plan for $\Pi_{\pi}^{-}$.

Now, let $\rho$ be a plan for $\Pi$ such that $\rho \neq \pi$. Let $o$ be the first action on $\rho$ that differs from the corresponding action of $\pi$. In other words, there exists a prefix $\rho^{\prime}=o_{1} \ldots o_{m}$ of $\rho$ such that (i) $\rho^{\prime} o$ is a prefix of $\rho$, (ii) $\rho^{\prime}$ is a prefix of $\pi$, and (iii) $\rho^{\prime} o$ is not a prefix of $\pi$. Then we have $r^{-1}\left(\rho^{\prime}\right)=$ $o_{1}^{3} \ldots o_{m}^{3}$, and since $o \neq o_{m+1}$, the next action on $r^{-1}(\rho)$ will not be $o_{m+1}^{3}$. If $o \in \pi$, then the next action will be $o^{2}$ (which we next show applicability of), and otherwise it will be $o^{e}$, in both cases setting the value of $\bar{v}$ to 0 . Thus, all the following actions $o^{\prime}$ are mapped to either $o^{\prime e}$ or $o^{\prime 1}$, and the preconditions of these actions are restricted to $\mathcal{V}$ and $\bar{v}=0$, the sequence of actions $r^{-1}(\rho)$ achieves the goal values on $\mathcal{V}$ and thus is a plan.

We finalize the proof by showing the applicability of $o^{2}$ in the state $s_{m}^{\prime}:=s_{0}^{\prime} \llbracket o_{1}^{3} \ldots o_{m}^{3} \rrbracket$ for $o \in \pi$ such that $o \neq$ $o_{m+1}$. Naturally, pre (o) holds in $s_{m}^{\prime}$. Further, since $s_{0}^{\prime}[\bar{v}]=$ $1, s_{0}^{\prime}\left[\bar{v}_{0}\right]=1$, and $s_{0}^{\prime}\left[\bar{v}_{i}\right]=0$ for all $1 \leq i \leq n$, after applying $o_{1}^{3} \ldots o_{m}^{3}$ we have $s_{m}^{\prime}[\bar{v}]=1, s_{m}^{\prime}\left[\bar{v}_{m}\right]=1$, and $s_{m}^{\prime}\left[\bar{v}_{i}\right]=0$ for all $0 \leq i \leq n, i \neq m$. Since $o \neq o_{m+1}$, for each $1 \leq i \leq n$ such that $o=o_{i}$, we have $s_{m}^{\prime}\left[\bar{v}_{i-1}\right]=0$, and thus $o^{2}$ is applicable in $s_{m}^{\prime}$.

## Devising Additional Plans

The top- $k$ planning approach proposed in the previous section reformulates a planning task of a solvable top- $k$ planning instance exactly $k$ times. While in each iteration, the reformulated planning task grows only linearly in the plan size, for larger $k$, such an approach is prohibitively expensive for anything but small tasks. How can we bypass this problem? Given an optimal plan $\pi$ for $\Pi$, it is often possible to infer additional optimal plans for $\Pi$ from the structure of the planning task $\Pi$. Our idea is to forbid in each iteration a set of plans instead of a single plan, decreasing the number of reformulations needed. Before we introduce an algorithm for this approach, we focus on characteristics of a graph $G(P)$ representing such a set of plans.
First, given two plans $\pi_{1}$ and $\pi_{2}$, if these plans intersect, i.e., pass through the same state $s$, then additional plans may be devised out of these two by following one of the plans until the state $s$ and the other plan from the state $s$ onwards. In general, a set of plans $P$ induces a directed graph $G(P)$ over the states of $\Pi$ with edges annotated by the actions on the plans. Each path in $G(P)$ from the initial state to some goal state is a plan for $\Pi$. Formally, $G(P)=(N, E)$, where $N=\left\{s \in \mathcal{S} \mid o_{1} \ldots o_{n} \in P, s=s_{0} \llbracket o_{1} \ldots o_{i} \rrbracket, 0 \leq i \leq n\right\}$


Figure 1: An example task with one variable and four actions changing its value from the initial value $p$ to $q$ and to the goal value $r$. Edges are labelled with actions, costs.
are all the states encountered by the plans in $P$ and $E=$ $\left\{(s, t) \mid s, t \in \mathcal{S}, o_{1} \ldots o_{n} \in P, s=s_{0} \llbracket o_{1} \ldots o_{i-1} \rrbracket, t=\right.$ $\left.s \llbracket o_{i} \rrbracket, 1 \leq i \leq n\right\}$ its edges. Each edge $(s, t)$ is labelled with the action $o \in \pi \in P$ inducing it.
$G(P)$ can be viewed as a compact representation for a set of plans $P$ of a planning task $\Pi$. Hence, often more plans are represented by $G(P)$ as compared to $P$. Let us now proof the correspondence of paths in $G(P)$ and plans for $\Pi$.

Lemma 1 Let $\Pi$ be a planning task and $P$ be a set of plans for $\Pi$. Then, any path in $G(P)$ from $s_{0}$ to some goal state of $\Pi$ corresponds to a plan for $\Pi$.

Proof: Let $s_{0}, s_{1}, \ldots, s_{n}$ with $s_{n} \in \mathcal{S}_{s_{\star}}$ be some path in $G(P)$. Each edge $\left(s_{i-1}, s_{i}\right)$ corresponds to some action $o_{i}$ on a plan in $P$, and thus $o_{i}$ is applicable in $s_{i-1}$, giving us $o_{1} \ldots o_{n}$ being a plan for $\Pi$.

Theorem 3 Let $\Pi$ be a planning task and $P$ be a set of optimal plans for $\Pi$. Then, any path in $G(P)$ from $s_{0}$ to some goal state of $\Pi$ corresponds to an optimal plan for $\Pi$.

Proof: Let $s_{0}, s_{1}, \ldots, s_{n}$ with $s_{n} \in \mathcal{S}_{s_{\star}}$ be some path in $G(P)$. From Lemma 1 we have that it corresponds to some plan $o_{1} \ldots o_{n}$ for $\Pi$, where each edge $\left(s_{i-1}, s_{i}\right)$ corresponds to an action $o_{i}$ on some optimal plan in $P$. Therefore, $h^{*}\left(s_{i-1}\right)=h^{*}\left(s_{i}\right)+\operatorname{cost}\left(o_{i}\right)$ or $\operatorname{cost}\left(o_{i}\right)=$ $h^{*}\left(s_{i-1}\right)-h^{*}\left(s_{i}\right)$. Summing over the actions in the plan we $\operatorname{get} \sum_{i=1}^{n} \operatorname{cost}\left(o_{i}\right)=\sum_{i=1}^{n} h^{*}\left(s_{i-1}\right)-h^{*}\left(s_{i}\right)=h^{*}\left(s_{0}\right)$.

If not all plans in $P$ are optimal, we may get plans from $G(P)$ with costs larger than of any plan in $P$. A simple example for that is described in Figure 1. There is one optimal plan $\pi_{1}=o_{1} o_{3}$ with the cost 2 , there are two plans with the cost $3, \pi_{2}=o_{1} o_{4}$ and $\pi_{3}=o_{2} o_{3}$. If $P=\left\{\pi_{1}, \pi_{2}, \pi_{3}\right\}$, then there is also a path in $G(P)$ that corresponds to a plan $o_{2} O_{4}$ with the cost 4 , which is strictly larger than of any of the plans in $P$. An algorithm for extracting plans out of $G(P)$ must therefore, be able to extract paths of a bounded cost from $G(P)$. One approach is a simple traversal of $G(P)$, starting from the node $s_{0}$. Given a bound $b$ on the total plan cost, a plan can be incrementally constructed by moving along an edge $(x, y)$ that corresponds to an action $o$ only if for the plan prefix constructed so far $\pi$ that leads from $s_{0}$ to $x$ holds $\operatorname{cost}(\pi)+\operatorname{cost}(o)+w^{*}(y) \leq b$, where $w^{*}(y)$ is the cost of the cheapest path from $y$ to some goal node $t \in \mathcal{S}_{s_{\star}}$. In other words, we can extend $\pi$ with $o$ if there is a possibility to reach the goal under the bound.


Figure 2: An order reduction of the example plan for the GRIPPER task with 4 balls.

## Partial Orders on Plans

One simple way to derive additional plans from an existing one is by reordering the actions along the plan. A sequential plan corresponds to a total order over the actions in the plan. It can be reduced to a partial order, which corresponds to multiple total orders, and thus to multiple plans. A minimal partial order can be obtained in polynomial time (Bäckström 1998). The suggested procedure greedily attempts to remove an ordering between two plan actions, while maintaining a valid partial order plan ${ }^{1}$. However, the best order of removal is not clear, and thus particular implementations may vary significantly in their performance.

In what follows, we propose a slightly different approach, exploiting the notion of independence between actions (Wehrle and Helmert 2012). The procedure simply follows the order of the actions in the plan, gathering (pairwise) independent actions into the set as long as possible. A new set is started every time a non-independent action is met. In what follows, we refer to this procedure as Indep. The complexity of this procedure is exponential in the size of the largest independent set of actions (Wehrle and Helmert 2012). Figure 2 depicts the graph constructed from a plan $\pi_{0}=\operatorname{pick}\left(b_{1}, L\right) \operatorname{pick}\left(b_{2}, R\right) \operatorname{move}(A, B) \operatorname{drop}\left(b_{1}, L\right)$ $\operatorname{drop}\left(b_{2}, R\right) \quad \operatorname{move}(B, A) \quad \operatorname{pick}\left(b_{3}, L\right) \quad \operatorname{pick}\left(b_{4}, R\right) \quad \operatorname{move}(A, B)$ $\operatorname{drop}\left(b_{3}, L\right) \operatorname{drop}\left(b_{4}, R\right)$ for the GRIPPER planning task with 4 balls, when reducing orders. The graph encodes 16 different plans obtained by reducing the order between $\operatorname{pick}\left(b_{1}, L\right)$ and $\operatorname{pick}\left(b_{2}, R\right)$, between $\operatorname{drop}\left(b_{1}, L\right)$ and $\operatorname{drop}\left(b_{2}, R\right)$, between $\operatorname{pick}\left(b_{3}, L\right)$ and $\operatorname{pick}\left(b_{4}, R\right)$, and between $\operatorname{drop}\left(b_{3}, L\right)$ and drop $\left(b_{4}, R\right)$.

Note that some other valid plan reorderings, such as $\operatorname{pick}\left(b_{3}, L\right) \operatorname{pick}\left(b_{4}, R\right) \operatorname{move}(A, B) \operatorname{drop}\left(b_{3}, L\right) \operatorname{drop}\left(b_{4}, R\right)$ $\operatorname{move}(B, A) \quad \operatorname{pick}\left(b_{1}, L\right) \quad \operatorname{pick}\left(b_{2}, R\right) \quad \operatorname{move}(A, B) \quad \operatorname{drop}\left(b_{1}, L\right)$ $\operatorname{drop}\left(b_{2}, R\right)$ are not obtained this way. In order to obtain all possible reorderings of a plan, one can use, e.g., a traversal procedure. In such a case, for a set $P$ of partial order plans, the graph $G(P)$ is induced by possible sequentializations of the plans in $P$. Note that this can be exponential in the plan length and thus might turn out too expensive in practice. Since we are interested in a bounded number of solutions, which might be significantly smaller than the number of valid reorderings (e.g., in GRIPPER domain), in our experiments a depth first search traversal without duplicate detection was chosen. The algorithm is bound to apply each of the actions on the plan exactly once, and therefore complete. In addition, a cycle detection is performed on each path. The number of times the algorithm

[^1]

Figure 3: An extension with symmetries of the example plan for the GRIPPER task with 4 balls (a fragment).
has reached the goal is then used in the stopping criteria. In what follows, we refer to this procedure as Naive.

## Symmetric Plans

One limitation of deriving plans based on reducing orders between actions is the restriction to the actions of the original plan. In what follows, we show how such a restriction can be alleviated. An additional way of deriving plans from already existing plans is by extending the graph $G(P)$, adding its symmetric counterparts in the state transition graph. This can be done using structural symmetries (Shleyfman et al. 2015). Structural symmetries are permutations of facts and actions that induce automorphisms of the state transition graph. Here, we present the definition of structural symmetries for $\mathrm{SAS}^{+}$as was given by Sievers et al. (2017).

Definition 4 (structural symmetry) For a $\mathrm{SAS}^{+}$planning task $\Pi=\left\langle\mathcal{V}, \mathcal{O}, s_{0}, s_{\star}\right.$, cost $\rangle$, let $F$ be the set of $\Pi$ 's facts, i.e., pairs $\langle v, d\rangle$ with $v \in \mathcal{V}, d \in \mathcal{D}(v)$. A structural symmetry for $\Pi$ is a permutation $\sigma: \mathcal{V} \cup F \cup \mathcal{O} \rightarrow \mathcal{V} \cup F \cup \mathcal{O}$, where:

1. $\sigma(\mathcal{V})=\mathcal{V}$ and $\sigma(F)=F$ such that $\sigma(\langle v, d\rangle)=\left\langle v^{\prime}, d^{\prime}\right\rangle$ implies $v^{\prime}=\sigma(v)$;
2. $\sigma(\mathcal{O})=\mathcal{O}$ such that for $o \in \mathcal{O}, \sigma(\operatorname{pre}(o))=\operatorname{pre}(\sigma(o))$, $\sigma(e f f(o))=\operatorname{eff}(\sigma(o)), \operatorname{cost}(\sigma(o))=\operatorname{cost}(o)$;
3. $\sigma\left(s_{\star}\right)=s_{\star}$;
where $\sigma\left(\left\{x_{1}, \ldots, x_{n}\right\}\right):=\left\{\sigma\left(x_{1}\right), \ldots, \sigma\left(x_{n}\right)\right\}$, and for $a$ partial state $s, s^{\prime}:=\sigma(s)$ is the partial state obtained from $s$ s.t. for all $v \in \mathcal{V}(s), \sigma(\langle v, s[v]\rangle)=\left\langle v^{\prime}, d^{\prime}\right\rangle$ implies $s^{\prime}\left[v^{\prime}\right]=d^{\prime}$.

We say that a structural symmetry $\sigma$ stabilizes a state $s$ if $\sigma(s)=s$. The set of all structural symmetries $\Gamma_{\Pi}$ of a planning task $\Pi$ forms a group under the composition operation. In what follows, for simplicity, by a symmetry group $\Gamma$ we refer to a subgroup of the symmetry group $\Gamma_{\Pi}$ of the planning task $\Pi$. In practice, a set of structural symmetries of a planning task (that stabilize a given state) can be efficiently computed using off-the-shelf tools for discovery of automorphisms in explicit graphs (Shleyfman et al. 2015).

We extend structural symmetries to sequences of actions in a natural way. Let $\sigma$ be a structural symmetry that stabilizes the initial state $s_{0}$. Then, for a plan $\pi$ of $\Pi$ we have $\sigma(\pi)$ being a plan of $\Pi$ with $\operatorname{cost}(\sigma(\pi))=\operatorname{cost}(\pi)$. Вy $P^{\Gamma}$ we denote the closure of $P$ under $\Gamma$, i.e., $P^{\Gamma}=\{\sigma(\pi) \mid \sigma \in$ $\Gamma, \pi \in P\}$.

Theorem 4 Let $\Pi$ be a planning task, $\Gamma$ be a symmetry group of $\Pi$ and $P$ be a set of $\Pi$ 's plans. If $P$ is a solution to
the top- $|P|$ planning problem, then $P^{\Gamma}$ is a solution to the top- $\left|P^{\Gamma}\right|$ planning problem.
Proof: Let $\pi \in P^{\Gamma}$ be some plan for $\Pi$. Since $\Gamma$ is a group, there exists $\sigma \in \Gamma$ such that $\sigma(\pi) \in P$. Since $P$ is a solution to the top- $|P|$ planning problem, if there exists a plan $\pi^{\prime}$ for $\Pi$ with $\operatorname{cost}\left(\pi^{\prime}\right)<\operatorname{cost}(\pi)=\operatorname{cost}(\sigma(\pi))$, then $\pi^{\prime} \in P$ and thus $\pi^{\prime} \in P^{\Gamma}$.

Similarly, but more generally, we define the closure of a graph $G=(N, E)$ under $\Gamma$ as a graph $G^{\Gamma}=\left(N^{\prime}, E^{\prime}\right)$, where $N^{\prime}=N^{\Gamma}=\{\sigma(n) \mid \sigma \in \Gamma, n \in N\}$ the closure of $N$ under $\Gamma$ and $E^{\prime}=\left\{\left(\sigma(n), \sigma\left(n^{\prime}\right)\right) \mid \sigma \in \Gamma,\left(n, n^{\prime}\right) \in E\right\}$.

Theorem 5 Let $\Pi$ be a planning task, $\Gamma$ be a symmetry group of $\Pi$ and $P$ be a set of $\Pi$ 's plans. Then $G(P)^{\Gamma}=$ $G\left(P^{\Gamma}\right)$.

Proof: First, we show that $G(P)^{\Gamma} \subseteq G\left(P^{\Gamma}\right)$. Let $(u, v)$ be some edge in $G(P)^{\Gamma}$. Thus $u=\sigma\left(u^{\prime}\right)$ and $v=\sigma\left(v^{\prime}\right)$ for some $\sigma \in \Gamma$ and edge $\left(u^{\prime}, v^{\prime}\right) \in G(P)$. Let $\pi^{\prime} \in P$ be some plan that traverses the edge $\left(u^{\prime}, v^{\prime}\right)$. Then $\pi=\sigma\left(\pi^{\prime}\right)$ is a plan, and it traverses $(u, v)$. Further, $\pi$ is in $P^{\Gamma}$, and thus $(u, v) \in G\left(P^{\Gamma}\right)$. For the other direction, let $(u, v)$ be some edge in $G\left(P^{\Gamma}\right)$. Then, $(u, v)$ is traversed by some plan $\pi \in P^{\Gamma}$. Thus, there exist $\sigma \in \Gamma$ and $\pi^{\prime} \in P$ such that $\pi=\sigma\left(\pi^{\prime}\right)$. Thus, $\left(\sigma^{-1}(u), \sigma^{-1}(v)\right)$ is traversed by the plan $\pi^{\prime}$, giving us $\left(\sigma^{-1}(u), \sigma^{-1}(v)\right) \in G(P)$. Therefore, by the definition of $G(P)^{\Gamma}$ we have $(u, v) \in G(P)^{\Gamma}$.

Theorems 4 and 5 above allow us to safely extend the graph $G(P)$ with symmetries by taking its closure under the symmetries that stabilize the initial state. That way, methods extending a set of plans can include a reduction of plan actions order or a symmetry based extension or both.

Figure 3 depicts (part of) the graph obtained by extending the example gripper task plan with structural symmetries. The structural symmetries found on this task are between balls and between arms, i.e., all balls are symmetric to each other, and the two arms are symmetric. Note though, that not all optimal plans are symmetric to each other. For instance, the example plan $\pi_{0}=\operatorname{pick}\left(b_{1}, L\right) \operatorname{pick}\left(b_{2}, R\right) \operatorname{move}(A, B)$ $\operatorname{drop}\left(b_{1}, L\right) \quad \operatorname{drop}\left(b_{2}, R\right) \quad \operatorname{move}(B, A) \quad \operatorname{pick}\left(b_{3}, L\right) \quad \operatorname{pick}\left(b_{4}, R\right)$ $\operatorname{move}(A, B) \operatorname{drop}\left(b_{3}, L\right) \operatorname{drop}\left(b_{4}, R\right)$ is not symmetric to the plan $\pi_{1}=\operatorname{pick}\left(b_{1}, L\right) \operatorname{pick}\left(b_{2}, R\right) \operatorname{move}(A, B) \operatorname{drop}\left(b_{1}, L\right) \operatorname{drop}\left(b_{2}, R\right)$ $\operatorname{move}(B, A) \quad \operatorname{pick}\left(b_{3}, L\right) \quad \operatorname{pick}\left(b_{4}, R\right) \quad \operatorname{move}(A, B) \quad \operatorname{drop}\left(b_{4}, R\right)$ $\operatorname{drop}\left(b_{3}, L\right)$ (the last two actions reordered), since there is no composition of symmetries that permute balls and arms that can map between the two plans. However, when permuting
ball3 with ball4 and arms together, we obtain the plan $\pi_{2}=\operatorname{pick}\left(b_{1}, R\right) \operatorname{pick}\left(b_{2}, L\right) \operatorname{move}(A, B) \operatorname{drop}\left(b_{1}, R\right) \operatorname{drop}\left(b_{2}, L\right)$ $\operatorname{move}(B, A) \quad \operatorname{pick}\left(b_{4}, R\right) \quad \operatorname{pick}\left(b_{3}, L\right) \quad \operatorname{move}(A, B) \quad \operatorname{drop}\left(b_{4}, R\right)$ $\operatorname{drop}\left(b_{3}, L\right)$. Both this $\pi_{2}$ and $\pi_{0}$, if followed through up to the last two actions, end up in the same state. Thus, adding both $\pi_{0}$ and $\pi_{2}$ to $G(P)$, allows us to extract $\pi_{1}$, as it follows $\pi_{0}$ for the first nine actions and $\pi_{2}$ for the last two. All optimal plans for the example task can be obtained this way and therefore the graph encodes all optimal plans for this task.

## Repeatedly Forbidding Multiple Plans

Having introduced the graph $G(P)$ as compact representation of a set of plans $P$ and having clarified how to extract plans from $G(P)$, we now devise a reformulation that forbids all plans represented by $G(P)$.

Definition $5(G(P)$-forbidding reformulation) Let $\Pi$ be a planning task over actions $\mathcal{O}$ and $P$ be some set of plans for $\Pi$. A task $\Pi^{\prime}$ over actions $\mathcal{O}^{\prime}$ is a $G(P)$-forbidding reformulation of $\Pi$ if there exists a mapping of actions $r: \mathcal{O}^{\prime} \mapsto \mathcal{O}$, and its extension $r^{\prime}$ such that $\pi^{\prime}$ is a plan for $\Pi^{\prime}$ iff $r^{\prime}\left(\pi^{\prime}\right)$ is a plan for $\Pi$ with $\operatorname{cost}^{\prime}\left(\pi^{\prime}\right)=\operatorname{cost}\left(r^{\prime}\left(\pi^{\prime}\right)\right)$ and $r^{\prime}\left(\pi^{\prime}\right) \notin G(P)$.

In what follows, we sometimes abuse the notation and treat $G(P)$ as the set of plans that can be derived from $G(P)$, denoting such plans by $\pi \in G(P)$. Further, by $\mathcal{O}(G(P))$ we denote the set of all actions in $P$. These are exactly the actions labeling the edges of $G(P)$. Finally, by $E^{o}$ we denote the subset of edges that are induced by the action $o$, and each action instance that corresponds to the edge $(s, t)$ is denoted by $o_{(s, t)}$. We now extend the plan forbidding reformulation presented in the previous section to forbid all plans in $G(P)$.

Definition 6 Let $\Pi=\left\langle\mathcal{V}, \mathcal{O}, s_{0}, s_{\star}\right.$, cost $\rangle$ be a planning task, $P$ be some set of plans, and $G:=G(P)=(N, E)$. The task $\Pi_{G}^{-}=\left\langle\mathcal{V}^{\prime}, \mathcal{O}^{\prime}, s_{0}^{\prime}, s_{\star}^{\prime}\right.$, cost $\rangle$ is defined as follows.

- $\mathcal{V}^{\prime}=\mathcal{V} \cup\{\bar{v}\} \cup\left\{\bar{v}_{s} \mid s \in N\right\}$, with $\bar{v}_{s}$ being binary variables,
- $\mathcal{O}^{\prime}=\left\{o^{e} \mid o \in \mathcal{O} \backslash \mathcal{O}(G)\right\} \cup \bigcup_{o \in(G)}\left\{o^{1}, o^{2}\right\} \cup \bigcup \bigcup\left\{o_{(s, t)}^{3}\right\}$,
where

$$
\begin{gathered}
o^{e}=\langle\operatorname{pre}(o), \text { eff }(o) \cup\{\langle\bar{v}, 0\rangle\}\rangle \\
o^{1}=\langle\operatorname{pre}(o) \cup\{\langle\bar{v}, 0\rangle\}, \text { eff }(o)\rangle \\
o^{2}=\left\langle\operatorname{pre}(o) \cup\{\langle\bar{v}, 1\rangle\} \cup\left\{\left\langle\bar{v}_{s}, 0\right\rangle \mid(s, t) \in E^{o}\right\},\right. \\
\operatorname{eff}(o) \cup\{\langle\bar{v}, 0\rangle\}\rangle \\
o_{(s, t)}^{3}=\left\langle\operatorname{pre}\left(o_{(s, t)}\right) \cup\left\{\langle\bar{v}, 1\rangle,\left\langle\bar{v}_{s}, 1\right\rangle\right\},\right. \\
\left.\operatorname{eff}\left(o_{(s, t)}\right) \cup\left\{\left\langle\bar{v}_{s}, 0\right\rangle,\left\langle\bar{v}_{t}, 1\right\rangle\right\}\right\rangle, \text { and } \\
\operatorname{cost}^{\prime}\left(o^{e}\right)=\operatorname{cost}^{\prime}\left(o^{1}\right)=\operatorname{cost}^{\prime}\left(o^{2}\right)=\operatorname{cost}^{\prime}\left(o^{3}\right)=\operatorname{cost}(o), \\
\bullet s_{0}^{\prime}[v]=s_{0}[v] \text { for all } v \in \mathcal{V}, s_{0}^{\prime}[\bar{v}]=1, s_{0}^{\prime}\left[\bar{v}_{s_{0}}\right]=1 \text {, and } \\
s_{0}^{\prime}\left[\bar{v}_{s}\right]=0 \text { for all } s \in N \backslash\left\{s_{0}\right\}, \text { and } \\
\bullet s_{\star}^{\prime}[v]=s_{\star}[v] \text { for all } v \in \mathcal{V} s . t . s_{\star}[v] \text { defined, and } s_{\star}^{\prime}[\bar{v}]=0 .
\end{gathered}
$$

The planning task $\Pi_{G}^{-}$forbids all sequences of actions $\pi \in G$ from being a plan, accepting all other plans of $\Pi$. Similarly to the case of a single plan, the variable $\bar{v}$ starts from the value 1 and switches to 0 when the sequence of actions applied is not a prefix of $\pi$. Once a value 0 is reached, it cannot be switched back to 1 . The actions $o^{1}$ and $o^{2}$ are the copies of actions in $\pi$ for the cases when $\pi$ is already discarded from consideration (variable $\bar{v}$ has already switched its value to 0 ) and for discarding $\pi$ from consideration (switching $\bar{v}$ to 0 ). In case there are multiple appearances of an action $o$ in $G, o^{2}$ requires all the corresponding variables $\bar{v}_{s}$ for each edge $(s, t)$ in $E$ that is induced by $o$ to have the value 0 . Thus, $o^{2}$ will be applicable only in states that are different from the states resulting from applying the prefix of some plan $\pi \in G$ up to $o$ in the initial state. $o_{(s, t)}^{3}$ are copies of the corresponding action $o$ along $G$, with a separate copy for each appearance of the same action $o$. These actions are responsible for following sequences $\pi \in G$ and are applicable while some such sequence is still followed.

Theorem 6 Let $\Pi$ be a planning task, $P$ be some set of plans, and $G:=G(P)=(N, E)$. The task $\Pi_{G}^{-}$is a $G$ forbidding reformulation of $\Pi$.

Proof: Let $r: \mathcal{O}^{\prime} \mapsto \mathcal{O}$ be the mapping defined by $r\left(o^{e}\right)=$ $o$ and $r\left(o^{1}\right)=r\left(o^{2}\right)=r\left(o^{3}\right)=o$. Note that $\Pi_{G}^{-}$restricted to the variables $\mathcal{V}$ equals to the task $\Pi$, modulo the three equal instances of the actions in $G$. Thus, for each plan $\pi$ for $\Pi_{G}^{-}, r(\pi)$ is a plan for $\Pi$.
For the second direction, since for each $o \in G$ at most one of the actions $o^{1}, o^{2}, o^{3}$ is applicable in each state $s$ of $\Pi_{G}^{-}$, given a sequence of actions $\rho$ applicable in the initial state of $\Pi$, it can be mapped to an applicable in the initial state of $\Pi_{G}^{-}$sequence of action $\rho^{\prime}$ such that $r\left(\rho^{\prime}\right)=\rho$, by choosing in each state the relevant representative out of $o^{e}$, $o^{1}, o^{2}$, and $o^{3}$. In other words, $r$ restricted to applicable in ${ }^{\prime}$ the initial state sequences of actions is invertible, and we denote its inverse mapping described above by $r^{-1}$.
First, let $\pi=o_{1} \ldots o_{n}$ be some plan in $G$ and let $\pi^{\prime}=$ $r^{-1}(\pi)$ be the inverse of $\pi$. Then $\pi^{\prime}=s_{0}^{\prime} o_{1}^{3} s_{1} \ldots s_{n-1} o_{n}^{3} s_{n}$, since at step $i$ we have $\bar{v}=1$ and $\bar{v}_{s_{i-1}}=1$. Thus, after applying $\pi^{\prime}$, the value of the variable $\bar{v}$ in $s_{n}$ remains 1 , and thus $\pi^{\prime}$ is not a plan for $\Pi_{G}^{-}$.

Now, let $\rho$ be a plan for $\Pi$ such that $\rho \notin G$. Let $o$ be the first action on $\rho$ that does not follow a path in $G$. In other words, there exists a prefix $\rho^{\prime}=s_{0} o_{1} s_{1} \ldots s_{m-1} o_{m} s_{m}$ of $\rho$ such that (i) $\rho^{\prime} o$ is a prefix of $\rho$, (ii) $\rho^{\prime}$ is a path in $G$, and (iii) $\rho^{\prime} o$ is not a path in $G$. Then we have $r^{-1}\left(\rho^{\prime}\right)=$ $o_{1}^{3} \ldots o_{m}^{3}$, and since $o \neq o_{m+1}$, the next action on $r^{-1}(\rho)$ will not be $o_{m+1}^{3}$. If $o \in G$, then the next action will be $o^{2}$ (which we next show applicability of), and otherwise it will be $o^{e}$, in both cases setting the value of $\bar{v}$ to 0 . Thus, all the following actions $o^{\prime}$ are mapped to either $o^{\prime e}$ or $o^{\prime 1}$, and the preconditions of these actions are restricted to $\mathcal{V}$ and $\bar{v}=0$, the sequence of actions $r^{-1}(\rho)$ achieves the goal values on $\mathcal{V}$ and thus is a plan.

We finalize the proof by showing the applicability of $o^{2}$ in the state $s_{m}^{\prime}:=s_{0}^{\prime} \llbracket o_{1}^{3} \ldots o_{m}^{3} \rrbracket$ for $o \in G$ such that $o$ does not correspond to any edge $\left(s_{m}, t\right)$. Naturally, pre $(o)$ holds in

```
Algorithm 2: IterativeTopKMultiple \((\langle\Pi, k\rangle)\)
    \(P \leftarrow \emptyset, \quad T \leftarrow \emptyset, \quad B \leftarrow \emptyset, \quad r \leftarrow i d\)
    while \(|T|<k\) and no failure occurs do
        \(\Pi, r^{\prime} \leftarrow \operatorname{GraphFORBIDREFORMULATION}(\Pi, G(P))\)
        \(r \leftarrow r \circ r^{\prime}\)
        \(\pi \leftarrow\) GetOptimalPlan \((\Pi)\)
        if GetOptimalPlan \((\Pi)\) reports UnSolvable
            then
                return SetTopElements \((T \cup B, k)\)
        end
        \(P \leftarrow \operatorname{ExtendPlan}(\Pi, \pi)\)
        \(T_{P} \leftarrow\left\{r\left(\pi^{\prime}\right) \mid \pi^{\prime} \in G(P), \operatorname{cost}\left(\pi^{\prime}\right)=\operatorname{cost}(\pi)\right\}\)
        \(B_{P} \leftarrow\left\{r\left(\pi^{\prime}\right) \mid \pi^{\prime} \in G(P), \operatorname{cost}\left(\pi^{\prime}\right)>\operatorname{cost}(\pi)\right\}\)
        \(T \leftarrow T \cup T_{P} \cup\left\{\pi^{\prime} \in B \mid \operatorname{cost}\left(\pi^{\prime}\right)=\operatorname{cost}(\pi)\right\}\)
        \(B \leftarrow B \backslash T \cup B_{P}\)
    end
    return SetTopElements \((T, k)\)
```

$s_{m}^{\prime}$. Further, since $s_{0}^{\prime}[\bar{v}]=1, s_{0}^{\prime}\left[\bar{v}_{s_{0}}\right]=1$, and $s_{0}^{\prime}\left[\bar{v}_{s}\right]=0$ for all $s \in N \backslash\left\{s_{0}\right\}$, after applying $o_{1}^{3} \ldots o_{m}^{3}$ we have $s_{m}^{\prime}[\bar{v}]=$ $1, s_{m}^{\prime}\left[\bar{v}_{s_{m}}\right]=1$, and $s_{m}^{\prime}\left[\bar{v}_{s}\right]=0$ for all $s \in N \backslash\left\{s_{m}\right\}$. Since $o$ does not correspond to any edge $\left(s_{m}, t\right)$, for each $(s, t) \in E^{o}$ we have $s_{m}^{\prime}\left[\bar{v}_{s}\right]=0$, and thus $o^{2}$ is applicable in $s_{m}^{\prime}$.

The scheme, depicted in Algorithm 2, works as follows. Once a plan is found, it is extended to a set of plans $P$ and then to the graph $G(P)$, which is forbidden in the next iteration. Further, the plans encoded by $G(P)$ are extracted and partitioned into two sets, optimal plans $T$ and non-optimal ones $B$. In the next iterations, the set $T$ is extended with optimal plans $T^{\prime}$ from that iteration, as well as all plans of the same cost as those in $T^{\prime}$ from the set $B$. The algorithm is thus iterating until the set $T$ consists of at least $k$ plans or no more plans exist. In practice, however, the algorithm can be simplified if the graph $G(P)$ encodes only plans of the same cost.

## Experimental Evaluation

In order to empirically evaluate the feasibility of our approach to finding top- $k$ plans, we implemented Algorithm 2 on top of the Fast Downward planning system (Helmert 2006), extended with the support for structural symmetries and the orbit space search algorithm (Domshlak, Katz, and Shleyfman 2015; Alkhazraji et al. 2014). As an underlying classical planner we used an orbit space search, not stabilizing the initial state, with the LM-cut heuristic (Helmert and Domshlak 2009), a state-of-the-art heuristic search planner. A natural candidate as a baseline for the comparison would be the $K^{*}$ algorithm. Unfortunately, the only existing (to our knowledge) implementation of the $K^{*}$ algorithm for classical planning supports only the SPPL language. Therefore, in order to be able to compare to an existing method for deriving top-k plans, we also implemented $K^{*}$ search within the Fast Downward planning system. As $K^{*}$ requires a consistent heuristic, we used the iPDB heuristic (Haslum et al. 2007). Since iPDB may require a long pre-search computation time, we also experimented with the blind heuristic.


Figure 4: The $k$-coverage as a function of the restriction on the number of required plans, for $k \leq 10000$.

Our benchmark suite includes all 1667 tasks from the optimal IPC STRIPS benchmarks. The experiments were performed on Intel(R) Xeon(R) CPU E7-8837 2.67GHz, with time (memory) limit of 30 min (2GB).

To measure the effect of various techniques for extending the set of existing plans, we switched the extension by symmetry on and off (adding " $S$ " to the configuration name if symmetry is switched on), and experimented with the two aforementioned approaches for plan reordering (Indep and Naive), as well as not reordering at all (None), giving us in total six configurations for the iterative approach and two configurations for the $K^{*}$ approach. The experiments were performed with a large $k$ value, namely $k=10000$, aiming at checking the feasibility of the suggested approaches for finding a large set of plans.

For a given task and a natural number $k$, the $k$-coverage is a value in $\{0,1\}$ assigning the value 1 if a solution to the top- $k$ planning problem was found (or proven unsolvable), according to Definition 1, and 0 otherwise. In order to measure the effect of the number of required plans $k$, we computed the $k$-coverage for any value of $k$ in $[1,10000]$, without rerunning the experiments for that value of $k$. We also experimented with a small value of $k=10$, obtaining similar results for all configurations. We also note that for the iterative configurations, there are only up to 10 tasks for each configuration that fail on memory, with most failures being due to timeouts. For $K^{*}$ based configurations, the opposite is true, with 12 and 68 timeouts for blind heuristic and iPDB, respectively.

Figure 4 depicts the results for all configurations, summed over all tasks in our benchmark suite. The horizontal lines correspond to $K^{*}$ with iPDB heuristic (top) and the blind heuristic (bottom). Both $K^{*}$ configurations found all required plans within a short time window, and thus always either fail to find any plan or find all required ones. In contrast, the iterative approaches exhibit an anytime behaviour. Further, since the iterative configurations exploit a better underlying classical planner, the $k$-coverage for smaller values of $k$ is significantly higher for the iterative approaches than

|  | $\mathrm{k}=10$ | $\mathrm{k}=100$ |  | $\mathrm{k}=1000$ | $\mathrm{k}=10000$ |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Coverage | IndepS NaiveS | IndepS | NaiveS | IndepS NaiveS | IndepS | NaiveS |  | $K_{\mathrm{iPDB}}^{*}$ |
| airport | 28 (2) 28 (1) | 22 (18) | 26 (1) | 14 (1) 26 (1) | 11 (1) | 26 (1) | 16 | 11 |
| barman11 | 8 (1) 8 (1) | 8 (1) | 8 (1) | 1 (1) 8 (1) | 0 (0) | 0 (0) | 4 | 4 |
| barman 14 | 3 (1) 3 (1) | 3 (1) | 3 (1) | 3 (1) 3 (1) | 0 (0) | 3 (1) | 0 | 0 |
| blocks | 22 (9) 27 (5) | 13 (99) | 16 (78) | 0 (0) 0 (0) | 0 (0) | 0 (0) | 18 | 18 |
| childsnack14 | 6 (1) 6 (1) | 6 (1) | 6 (1) | 6 (1) 6 (1) | 6 (1) | 6 (1) | 0 | 0 |
| depot | 9 (1) 9 (1) | 6 (20) | 9 (9) | 3 (1) 8 (34) | 2 (1) | 7 (1) | 3 | 3 |
| driverlog | 13 (1) 13 (1) | 13 (15) | 13 (6) | 2 (1) 11 (2) | 0 (0) | 7 (1) | 4 | 4 |
| elevators08 | 22 (2) 22 (1) | 8 (18) | 19 (6) | 3 (2) 18 (1) | 0 (0) | 14 (1) | 4 | 4 |
| elevators 11 | 18 (2) 18 (1) | 6 (7) | 15 (1) | 3 (2) 15 (1) | 0 (0) | 12 (1) | 2 | 2 |
| floortile11 | 8 (1) 8 (1) | 8 (1) | 8 (1) | 4 (1) 8 (1) | 0 (0) | 8 (1) | 0 | 0 |
| floortile14 | 8 (1) 8 (1) | 8 (1) | 8 (1) | 6 (1) 8 (1) | 0 (0) | 8 (1) | 0 | 0 |
| freecell | 13 (2) 14 (1) | 6 (75) | 11 (36) | 0 (0) 2 (1) | 0 (0) | 0 (0) | 13 | 13 |
| ged14 | 13 (8) 13 (8) | 5 (96) | 5 (94) | 0 (0) 0 (0) | 0 (0) | 0 (0) | 10 | 10 |
| grid | 1 (9) 1 (9) | 1 (99) | 1 (99) | 0 (0) 0 (0) | 0 (0) | 0 (0) | 0 | 0 |
| gripper | 20 (1) 20 (1) | 20 (1) | 20 (1) | 20 (1) 20 (1) | 20 (2) | 20 (1) | 6 | 6 |
| hiking 14 | 13 (1) 13 (1) | 13 (1) | 13 (1) | 13 (8) 13 (1) | 6 (1) | 11 (1) | 6 | 6 |
| logistics00 | 20 (1) 20 (1) | 20 (5) | 20 (1) | 17 (1) 20 (11) | 16 (1) | 17 (1) | 10 | 10 |
| logistics98 | 6 (1) 6 (1) | 6 (13) | 6 (1) | 2 (1) 6 (1) | 2 (1) | 5 (1) | 1 | 1 |
| miconic | 142 (1) 142 (1) | 142 (13) | 141 (13) | 102 (6) 79 (5) | 78 (2) | 37 (4) | 40 | 40 |
| movie | 30 (1) 30 (1) | 30 (1) | 30 (1) | 30 (4) 30 (1) | 30 (76) | 30 (6) | 30 | 30 |
| mprime | 22 (7) 22 (6) | 13 (88) | 13 (81) | 0 (0) 0 (0) | 0 (0) | 0 (0) | 1 | 1 |
| mystery | 19 (6) 19 (6) | 15 (61) | 15 (58) | 4 (1) 4 (1) | 4 (1) | 4 (1) | 5 | 8 |
| nomystery 11 | 15 (1) 15 (1) | 15 (14) | 15 (12) | 3 (26) 5 (6) | 1 (1) | 1 (1) | 7 | 9 |
| openstacks08 | 24 (1) 24 (1) | 24 (2) | 24 (1) | 23 (2) 23 (1) | 18 (1) | 22 (1) | 14 | 17 |
| openstacks11 | 19 (1) 19 (1) | 19 (1) | 19 (1) | 19 (1) 19 (1) | 16 (1) | 19 (1) | 9 | 12 |
| openstacks14 | 5 (1) 5 (1) | 5 (1) | 5 (1) | 4 (1) 5 (1) | 4 (1) | 5 (1) | 1 | 1 |
| openstacks | 5 (1) 7 (1) | 5 (1) | 6 (1) | 1 (10) 5 (1) | 0 (0) | 5 (1) | 7 | 7 |
| parcprinter08 | 18 (2) 18 (1) | 17 (30) | 18 (6) | 5 (2) 17 (1) | 4 (2) | 16 (2) | 7 | 13 |
| parcprinter 11 | 13 (2) 13 (1) | 12 (30) | 13 (1) | 3 (1) 13 (1) | 2 (1) | 12 (2) | 3 | 9 |
| parking 11 | 1 (9) 1 (9) | 0 (0) | 0 (0) | 0 (0) 0 (0) | 0 (0) | 0 (0) | 0 | 0 |
| parking 14 | 0 (0) 0 (0) | 0 (0) | 0 (0) | 0 (0) 0 (0) | 0 (0) | 0 (0) | 0 | 0 |
| pathways-nn | 5 (1) 5 (1) | 5 (14) | 5 (2) | 4 (42) 5 (28) | 2 (1) | 3 (1) | 3 | 3 |
| pegsol08 | 28 (7) 28 (4) | 15 (66) | 20 (47) | 5 (73) 7 (51) | 5 (73) | 5 (72) | 26 | 26 |
| pegsol11 | 18 (6) 18 (3) | 5 (44) | 9 (17) | 0 (0) 3 (1) | 0 (0) | 0 (0) | 16 | 16 |
| pipes-notank | 20 (2) 21 (1) | 17 (18) | 19 (8) | 1 (87) 12 (16) | 0 (0) | 6 (1) | 6 | 6 |
| pipes-tank | 16 (2) 16 (1) | 14 (16) | 15 (15) | 11 (10) 12 (8) | 5 (6) | 8 (6) | 6 | 6 |
| psr-small | 49 (3) 49 (2) | 43 (42) | 49 (28) | 7 (4) 21 (2) | 1 (1) | 11 (1) | 43 | 43 |
| rovers | 7 (1) 7 (1) | 6 (30) | 7 (3) | 0 (0) 5 (84) | 0 (0) | 3 (1) | 4 | 4 |
| satellite | 13 (1) 13 (1) | 11 (18) | 13 (12) | 4 (1) 8 (1) | 3 (1) | 8 (1) | 4 | 4 |
| scanalyzer08 | 16 (2) 16 (1) | 16 (21) | 16 (16) | 9 (1) 13 (5) | 7 (1) | 8 (1) | 6 | 6 |
| scanalyzer11 | 13 (1) 13 (1) | 13 (11) | 13 (7) | 8 (1) 12 (5) | 7 (1) | 8 (1) | 3 | 3 |
| sokoban08 | 14 (6) 29 (2) | 6 (75) | 23 (20) | 0 (0) 11 (1) | 0 (0) | 10 (1) | 15 | 21 |
| sokoban 11 | 9 (6) 19 (2) | 3 (52) | 14 (11) | 0 (0) 6 (1) | 0 (0) | 5 (1) | 12 | 17 |
| storage | 17 (2) 17 (2) | 17 (23) | 17 (20) | 3 (19) 7 (16) | 0 (0) | 5 (2) | 11 | 11 |
| tetris 14 | 8 (1) 8 (1) | 5 (22) | 8 (6) | 0 (0) 7 (1) | 0 (0) | 6 (1) | 5 | 1 |
| tidybot11 | 11 (1) 14 (1) | 3 (68) | 14 (5) | 0 (0) 10 (1) | 0 (0) | 0 (0) | 1 | 1 |
| tidybot14 | 4 (1) 9 (1) | 0 (0) | (1) | 0 (0) 6 (1) | 0 (0) | 1 (1) | 0 | 0 |
| tpp | 7 (2) 7 (2) | 7 (19) | 7(18) | 5 (7) 5 (8) | 4 (1) | 4 (1) | 5 | 5 |
| transport08 | $9(5) \quad 9(4)$ | 5 (58) | 7(32) | 0 (0) 3 (7) | 0 (0) | 1 (1) | 8 | 8 |
| transport11 | 4 (6) 5 (4) | 0 (0) | (1) | 0 (0) 3 (1) | 0 (0) | 1 (1) | 3 | 3 |
| transport14 | 2 (5) 4 (3) | 1 (22) | ) 2 (1) | 0 (0) 2 (1) | 0 (0) | 1 (1) | 3 | 3 |
| trucks | 12 (1) 12 (1) | 12 (14) | ) 12 (10) | 5 (1) 8 (1) | 1 (1) | 5 (1) | 4 | 4 |
| visitall11 | 9 (6) 9 (6) | 9 (68) | 9 (68) | 1 (124) 1 (124) | 0 (0) | 0 (0) | 8 | 8 |
| visitall14 | 4 (7) 4 (7) | 1 (13) | 1 (13) | 0 (0) 0 (0) | 0 (0) | 0 (0) | 2 | 2 |
| woodwork08 | 18 (1) 18 (1) | 18 (15) | 18 (3) | 8 (10) 16 (1) |  | 13 (1) | 4 | 6 |
| woodwork11 | 12 (1) 12 (1) | 12 (4) | 12 (1) | 7 (11) 12 (1) | 3 (1) | 12 (1) | 0 | 1 |
| zenotravel | 13 (3) 13 (3) | 12 (43) | ) 13 (39) | 4 (57) 7 (33) | 1 (1) | 4 (1) | 5 | 5 |
| Sum | $882 \quad 927$ | 725 | - 841 | 373 574 | 263 | 423 | 424 | 452 |

Table 1: Per-domain $k$-coverage for selected values of $k$. Average number of iterations in parentheses.
for the $K^{*}$ ones. Looking at the overall $k$-coverage, we can see a clear dominance of NaiveS over all other configurations, up to $k=4320$. For $k>4320, K^{*}$ with iPDB has the best overall $k$-coverage. For specific domains, however, the picture might differ substantially.

Table 1 shows the per-domain $k$-coverage results on selected configurations, namely IndepS and NaiveS for the iterative approach and both $K^{*}$ configurations for three values of $k$, namely $k=10,100,1000,10000$. The first two columns depict the $k$-coverage of the two iterative configurations for $k=10$, the next two for $k=100$ and the next two for $k=1000$. The last block of four columns shows the $k$ coverage of all four selected configurations for $k=10000$. As mentioned above, the $K^{*}$ configurations (two rightmost columns) perform similarly with all values of $k \leq 10000$ and thus are shown once.

First, looking at the rightmost part, corresponding to $k=$ 10000 , note that there is no clear advantage to either of the approaches across the domains. $K^{*}$ achieves better performance in 23 domains out of 57, while the iterative approach performs better in 28 domains. Further, there is often a large difference in coverage between the two approaches, to one side and to another, making them complementary. Within each approach, there is a clear advantage to one of the configurations. For the iterative approach, with the exception of the Miconic domain, NaiveS performs at least as good as IndepS, performing strictly better in 37 domains. This dominance is preserved for smaller $k$ values, with strict dominance in 39 domains for $k=1000$, in 25 domains for $k=100$, and 10 for $k=10$. NaiveS loses to IndepS only in Miconic domain, for $k>10$. For the $K^{*}$ based approach, there are two domains where the blind search performs better than iPDB, namely Airport and Tetris, and performs strictly worse in 10 domains. Note that there are 45 domains with equal $k$-coverage for the two $K^{*}$ based approaches.

Moving on to smaller $k$ values, note that for $k=1000$ the dominance shifts further towards the iterative approach, with 34 domains vs. 17 domains for $K^{*}$. For $k=100$ it becomes 43 domains vs. 10 domains. For $k=10$, the iterative approach always performs at least as good as $K^{*}$, achieving better performance on 54 out of 57 domains.

Looking at some non-IPC domains (Sohrabi, Udrea, and Riabov 2013; Riabov, Sohrabi, and Udrea 2014; Sohrabi et al. 2018), we note that these were modelled with planner efficiency in mind. For the iterative approach, the number of plans found and forbidden per iteration is typically 1 , and therefore the number of plans found within the time and memory restrictions is typically low. $K^{*}$, on the other hand, even with the blind heuristic performs extremely well on these domains, solving all tasks for $k=10000$ in all but the risk management domain (Sohrabi et al. 2018). On this domain, it finds all 10000 requested top plans in 133 out of 200 tasks, failing to find any plans on the other 67 tasks.

## Conclusions and Future Work

We have presented a novel approach to the problem of topk planning, based on iterative computation of optimal solutions, exploiting existing optimal planners. We proposed two techniques for deriving additional solutions from previously found ones. To empirically evaluate the feasibility of our suggested approach, we have implemented an existing approach to top- $k$ planning, $K^{*}$, on top of a state-of-the-art planner. We show that these techniques are complementary in their performance, with $K^{*}$ being more beneficial for extremely large values of $k$.

In the future work, we intend to focus on both approaches. For the iterative approach, other methods for deriving additional plans given already existing plans can lead to considerable further performance improvement. For the $K^{*}$ based approach, we intend to explore several directions. First, to investigate the adaptation of the search algorithm to work with state-of-the-art non-consistent admissible heuristics, such as LM-cut. Second, we would like to explore the usage of search pruning techniques, such as partial order reduction and symmetry reduction for top- $k$ planning, ensuring
that relevant solutions are not pruned. Finally, some effort should be invested in exploring the high memory consumption of $K^{*}$ and ways to overcome it.

## References

Aljazzar, H., and Leue, S. 2011. K*: A heuristic search algorithm for finding the k shortest paths. Artificial Intelligence 175(18):2129-2154.
Alkhazraji, Y.; Katz, M.; Mattmüller, R.; Pommerening, F.; Shleyfman, A.; and Wehrle, M. 2014. Metis: Arming Fast Downward with pruning and incremental computation. In the 8th International Planning Competition (IPC-8): planner abstracts, 88-92.
Bäckström, C., and Nebel, B. 1995. Complexity results for $\mathrm{SAS}^{+}$planning. Computational Intelligence 11(4):625655.

Bäckström, C. 1998. Computational aspects of reordering plans. Journal of Artificial Intelligence Research 9:99-137.
Bryce, D. 2014. Landmark-based plan distance measures for diverse planning. In Proceedings of the 24th International Conference on Automated Planning and Scheduling (ICAPS), 56-64.
Coman, A., and Muñoz-Avila, H. 2011. Generating diverse plans using quantitative and qualitative plan distance metrics. In Proceedings of the 25th National Conference on Artificial Intelligence (AAAI), 946-951.
Domshlak, C.; Katz, M.; and Shleyfman, A. 2015. Symmetry breaking in deterministic planning as forward search: Orbit space search algorithm. Technical Report IS/IE-201503, Technion, Haifa.
Fox, M.; Gerevini, A.; Long, D.; and Serina, I. 2006. Plan stability: Replanning versus plan repair. In Proceedings of the 16th International Conference on Automated Planning and Scheduling (ICAPS), 212-221.
Haslum, P.; Botea, A.; Helmert, M.; Bonet, B.; and Koenig, S. 2007. Domain-independent construction of pattern database heuristics for cost-optimal planning. In Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI 2007), 1007-1012.
Helmert, M., and Domshlak, C. 2009. Landmarks, critical paths and abstractions: What's the difference anyway? In Proceedings of the 19th International Conference on Automated Planning and Scheduling (ICAPS 2009), 162-169.
Helmert, M. 2006. The Fast Downward planning system. Journal of Artificial Intelligence Research 26:191-246.
Keyder, E., and Geffner, H. 2009. Soft Goals Can Be Compiled Away. Journal of Artificial Intelligence Research 36:547-556.
Nguyen, T.; Do, M.; Gerevini, A.; Serina, I.; Srivastava, B.; and Kambhampati, S. 2012. Generating diverse plans to handle unknown and partially known user preferences. Artificial Intelligence 190:1-31.
Riabov, A., and Liu, Z. 2005. Planning for stream processing systems. In Proceedings of the 20th National Conference on Artificial Intelligence (AAAI), 1205-1210.

Riabov, A. V.; Sohrabi, S.; and Udrea, O. 2014. New algorithms for the top-k planning problem. In ICAPS 2014 Scheduling and Planning Applications woRKshop, 10-16.
Shleyfman, A.; Katz, M.; Helmert, M.; Sievers, S.; and Wehrle, M. 2015. Heuristics and symmetries in classical planning. In Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI 2015), 3371-3377.
Sievers, S.; Wehrle, M.; Helmert, M.; and Katz, M. 2017. Strengthening canonical pattern databases with structural symmetries. In Proceedings of the 10th Annual Symposium on Combinatorial Search (SoCS 2017), 91-99.
Sohrabi, S.; Riabov, A.; Udrea, O.; and Hassanzadeh, O. 2016. Finding diverse high-quality plans for hypothesis generation. In Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI), 1581-1582.
Sohrabi, S.; Riabov, A. V.; Katz, M.; and Udrea, O. 2018. An AI planning solution to scenario generation for enterprise risk management. In Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI 2018).
Sohrabi, S.; Riabov, A.; and Udrea, O. 2016. Plan recognition as planning revisited. In Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI), 3258-3264.
Sohrabi, S.; Udrea, O.; and Riabov, A. 2013. Hypothesis exploration for malware detection using planning. In Proceedings of the 27th National Conference on Artificial Intelligence (AAAI), 883-889.
Wehrle, M., and Helmert, M. 2012. About partial order reduction in planning and computer aided verification. In Proceedings of the 22nd International Conference on Automated Planning and Scheduling (ICAPS 2012), 297-305.


[^0]:    *The work was performed during an internship at IBM
    Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[^1]:    ${ }^{1}$ A partial order plan is valid if all its total order plans are valid.

