K^* and Partial Order Reduction for Top-quality Planning

Anonymous submission

Abstract

Partial order reduction techniques are successfully used for various settings in planning, such as classical planning with A^* search or with decoupled search, fully-observable non-deterministic planning with LAO^*, planning with resources, or even goal recognition design. Here, we continue this trend and show that partial order reduction can be used for top-quality planning with K^* search. We discuss the possible pitfalls of using stubborn sets for top-quality planning and the guarantees provided. We perform an empirical evaluation that shows the proposed approach to significantly improve over the current state of the art in unordered top-quality planning.

Introduction

The need for producing a set of top-quality plans is well established by many real-life applications, including plan recognition (Sohrabi, Riabov, and Udrea 2016), malware detection (Boddy et al. 2005), business process automation (Chakraborti et al. 2020), and automated machine learning (Katz et al. 2020). Top-quality planning deals with supplying the demand, generating all high-quality plans up to a certain bound (Katz, Sohrabi, and Udrea 2020). If some plans are considered equivalent from an application perspective, as is often the case with plans that are equivalent if operator orderings are ignored, it can be sufficient to cover the equivalence classes instead of all plans (Katz, Sohrabi, and Udrea 2020; Katz and Sohrabi 2022). This so-called unordered top-quality planning serves, among other, as a basis for solving additional computational problems, such as quality-aware diverse planning (Nguyen et al. 2012; Vadlamudi and Kambhampati 2016; Katz, Sohrabi, and Udrea 2022). While this is clearly an important computational problem of high practical value, the selection of available planners is quite limited. The two existing approaches are based on Forbid Iterative (Katz et al. 2018), an approach that iteratively solves a modified planning problem with classical planners, where problem modifications are intended to decrease the set of solutions by previously found ones, and SymK (Speck, Mattmüller, and Nebel 2020), a symbolic search based approach. While in top-k planning, a K^* search-based approach (Aljazzar and Leue 2011) was recently shown to produce competitive results (Anonymous 2023b), it was not applied to unordered top-quality planning.

In this work, we propose a competitive unordered top-quality planner based on K^* search and partial order reduction. Partial order reduction techniques are successfully used for various settings in planning, such as classical planning with A^* search or with decoupled search (Gnad, Hoffmann, and Wehrle 2019), fully-observable non-deterministic planning with LAO^* (Winterer et al. 2017), planning with resources (Wilhelm, Steinmetz, and Hoffmann 2018), or even goal recognition design (Keren, Gal, and Karpas 2018). They were never, however, applied in the setting of unordered top-quality planning, arguably among the more naturally fitting problem, where different orderings of the same plan are of no added value. Focusing on stubborn sets (Valmari 1989), we show that K^* search over the reduced with strong stubborn sets (Wehrle and Helmert 2012; Alkhazraji et al. 2012; Wehrle and Helmert 2014) space is safe to use for unordered top-quality planning. We show that no adaptation to these techniques is needed if a single-goal planning task transformation is used. Such transformations are required for K^* search and must be performed anyway. We point out the weaknesses of the transformation used in previous implementations of K^* (Katz et al. 2018) and suggest using a recently proposed transformation for symmetry pruning variant (Anonymous 2023a), which preserves the pruning power of the original planning task. We perform an experimental evaluation that shows a significant increase in performance over the current state of the art and establishes our approach as the new state of the art for unordered top-quality planning.

Background

In this section, we introduce the necessary concepts in top-quality planning, K^* search, as well as stubborn sets for partial order reduction.

Top-quality Planning

We consider classical planning tasks in the well-known SAS$^+$ formalism (Bäckström and Nebel 1995), extended with action costs. Such *planning tasks* $\Pi = \langle \mathcal{V}, \mathcal{O}, s_0, s_\star \rangle$ consist of a finite set of finite-domain *state variables* $\mathcal{V}$, a finite set of *actions* $\mathcal{O}$, an *initial state* s_0, and the *goal* $s_\star$. Each variable $v \in \mathcal{V}$ is associated with a finite domain $dom(v)$ of variable values. A *partial assignment* p maps a subset of

variables $vars(p) \subseteq \mathcal{V}$ 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 $v \in vars(p)$ and we say $p[v]$ is *undefined* if $v \notin vars(p)$. A partial assignment s with $vars(s) = \mathcal{V}$, is called a *state*. State s is *consistent* with partial assignment p if they agree on all variables in $vars(p)$, denoted by $p \subseteq s$. s_0 is a state and $s_\star$ is a partial assignment. A state s is called a *goal state* if $s_\star \subseteq s$ and the set of all goal states is denoted by $\mathcal{S}_{s_\star}$. Each action o in $\mathcal{O}$ is a pair $\langle pre(o), \mathit{eff}(o)\rangle$ where $pre(o)$ and $\mathit{eff}(o)$ are partial assignments called *precondition* and *effect*, respectively. Further, o has an associated *cost* $C(o) \in \mathbb{R}^{0+}$. An action o is applicable in state s if $pre(o) \subseteq s$. Applying o in s results in a state denoted by $s[\![o]\!]$ where $s[\![o]\!][v] = \mathit{eff}(o)[v]$ for all $v \in vars(\mathit{eff})$ and $= s[\![o]\!][v] = s[v]$ for all other variables. An action sequence $\pi = \langle o_1, \cdots, o_n\rangle$ 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 $s_{i-1}[\![o_i]\!] = s_i$ for $0 \leq i \leq n$. We denote s_n by $s[\![\pi]\!]$. An action sequence with $s_0[\![\pi]\!] \in \mathcal{S}_{s_\star}$ is called a *plan*. The cost of a plan π, denoted by $C(\pi)$ is the summed cost of the actions in the plan. The set of all plans is denoted by $\mathcal{P}_\Pi$. A plan is *optimal* if its cost is minimal among all plans in $\mathcal{P}_\Pi$. Cost-optimal planning deals with finding an optimal plan or proving that no plan exists. In the latter case we say that the task is *unsolvable*.

Extending cost-optimal planning, top-quality planning (Katz, Sohrabi, and Udrea 2020) deals with finding *all* plans of up to a specified cost. Formally, the **top-quality** planning problem is as follows. Given a planning task Π and a natural number q, find the set of plans $P = \{\pi \in \mathcal{P}_\Pi \mid cost(\pi) \leq q\}$. In some cases, an equivalence between plans can be specified, allowing to possibly skip some plans, if equivalent plans are found. The corresponding problem is called **quotient top-quality** planning and it is formally specified as follows. Given a planning task Π, an equivalence relation N over its set of plans $\mathcal{P}_\Pi$, and a natural number q, find a set of plans $P \subseteq \mathcal{P}_\Pi$ such that $\bigcup_{\pi \in P} N[\pi]$ is the solution to the top-quality planning problem. The most common case of such an equivalence relation is when the order of actions in a valid plan is not significant from the application perspective. In other words, when you can reorder some of the actions in a plan and still get a valid plan. The corresponding problem is called **unordered top-quality** planning and is formally specified as follows. Given a planning task Π and a natural number q, find a set of plans $P \subseteq \mathcal{P}_\Pi$ such that P is a solution to the quotient top-quality planning problem under the equivalence relation $\mathrm{U}_\Pi = \{(\pi, \pi') \mid \pi, \pi' \in \mathcal{P}_\Pi, \mathrm{MS}(\pi) = \mathrm{MS}(\pi')\}$.

K^* Search for Top-quality Planning

Given a top-quality planning problem $\langle \Pi, q\rangle$, K^* first performs A^* until the cost bound is reached, and then applies Eppstein's Algorithm (EA) to the search graph revealed by A^*. We omit here the details of EA as they are not necessary for this paper and refer a curious reader to Aljazzar and Leue (2011) and Eppstein (1998) for details. One important property that we emphasize is that EA can enumerate all paths in the search graph in the order of their costs.

One of the limitations of K^* is its restriction to graphs with a single goal state. In planning, however, tasks can have many goal states. In cases when the partial assignment $s_\star$ is not a full state, it is possible to transform Π into a planning task with a single goal state. One way to achieve that is to add one binary variable v_g to indicate whether a goal was reached. Further, we add one zero-cost action o_g with precondition $pre(o_g) = s_\star \cup \{v_g = 0\}$ and effect $\mathit{eff}(o_g) = \{v_i = t[v_i] | v_i \in vars(t)\} \cup \{v_g = 1\}$ for an arbitrary full state t over the original variables. Additionally, each original action precondition, as well as the initial state are extended with $\{v_g = 0\}$. Finally, the goal is set to $\mathit{eff}(o_g)$, making it a full variable assignment. For a planning task Π, we denote its single goal transformation by Π_g. In words, the additional goal-achieving zero-cost action can be applied once (and only once) the original goal was achieved, changing the state to the new goal state. No action is applicable in the new goal state, and therefore there is one-to-one correspondence between the plans of Π and those of Π_g. The transformation was used, albeit not described by Katz et al. (2018), in their implementation of K^*. It is worth noting that domain independent heuristics for planning can be quite sensitive to such transformations.

Partial Order Reduction

A central to partial order reduction techniques is the notion of *safe* successor pruning (Wehrle and Helmert 2014).

Definition 1 (safe) *Let succ be a successor pruning function for a planning task Π. We say that succ is safe if for every state s, the cost of an optimal solution for s is the same when using the pruned state space induced by succ as when using the full state space.*

When using safe successor pruning, it is possible to search the pruned state space instead when searching for cost-optimal plans. Stubborn sets (Wehrle and Helmert 2012; Alkhazraji et al. 2012) induce safe successor pruning functions by helping identifying actions that can safely be ignored at node expansion. It is done by specifying a set, such that if an applicable action is not in the set, it can be safely ignored (e.g., Wehrle and Helmert 2014).

At the core of these partial order reduction techniques is the idea that, for each non-goal state s, if a goal is reachable from s, then at least one *strongly optimal* (an optimal plan with a minimal number of 0-cost actions among all optimal plans) is preserved in the pruned state space.

Two main notions in stubborn sets are *interference* and *necessary enabling sets*. Interference dictates whether two actions disable each other or conflict. Necessary enabling sets for an action o and a set of paths from the initial state is a set of actions that appear on the paths that include o before its first appearance.

Stubborn Sets for Top-quality Planning

In this section, we show that K^* search can be used for solving the unordered top-quality planning problem. To do that, we need to ensure that the stubborn sets are well-defined for all states expanded by the search algorithm. For non-goal states, this is true. For goal states, however, stubborn sets

may prune all plans. Therefore, it is crucial that the search algorithm does not expand goal states. For K^* this is indeed the case, since it solves the single-goal transformation Π_g of the input planning task Π, and in the single goal transformation of Katz et al. (2018), the (only) goal state does not have successors, so plans do not traverse through the goal state. Therefore, from now on, we focus on non-goal states only.

We start with an observation that existing work (e.g., Alkhazraji et al. 2012) proves a stronger property than preserving at least one optimal plan. Let us state the property explicitly. Let Π be a planning task, s be some state and $\pi_s = o_1, \ldots, o_n$ be some plan for s. Let $f : \mathcal{S} \mapsto P(O)$ be a strong stubborn set and $O_f(s) := O(s) \cap f(s)$ be the corresponding subset of applicable actions. If $o_1 \notin O_f(s)$, let i be the smallest index such that $o_i \in O_f(s)$. Then, $\pi'_s = o_i, o_1, \ldots, o_{i-1}, o_{i+1}, \ldots, o_n$ obtained from π_s by moving the action o_i to the front, is also a plan for s.

It is worth noting that this property is also strongly related to the *operator shifting property* (Sievers and Wehrle 2021).

We can now use this observation to prove the following theorem.

Theorem 1 *Let $\mathcal{T}$ be the transition system of Π_g, $f : \mathcal{S} \mapsto P(O)$ be a strong stubborn set, and $\mathcal{T}_f$ be the corresponding reduced transition system. For each plan π of Π_g, there exists a reordering π' that corresponds to a goal path in $\mathcal{T}_f$.*

Proof: We construct a plan π' from $\pi = o_1, \ldots, o_n$ by iterative application of the observation above. We start with π' being an empty sequence. Let s be the current state, starting with the initial state s_0 and let π_s be a plan for the current state, starting with π.

While π_s is not an empty sequence, we find the next action to add to π' as follows. Let o be the first action from π_s. If o is in $O_f(s)$, then we add o to π' and update the current state to be $s[\![o]\!]$. We remove the first action from π_s, maintaining π_s to be a plan for s. If o is not in $O_f(s)$, we use the observation to find $o_i \in O_f(s)$ and a plan π'_s for s that is a reordering of π_s that starts with o_i. We then add o_i to π' and update the current state to be $s[\![o_i]\!]$. We replace π_s with π'_s and then remove the first action from π_s, maintaining π_s to be a plan for s.

Every iteration, the length of π_s is reduced by 1 action, moving the action to π'. Thus, the concatenation of π' and π_s at each step is a plan, a reordering of the plan π. □

A corollary from Theorem 1 is that we can solve unordered top-quality planning by finding plans in the reduced transformed transition system, using strong stubborn sets for pruning the search.

Next, we would like to question whether the stubborn sets applied to the single-goal transformation have the same pruning power as if they were applied to the original task. Looking closer at the transformation, since the goal state may include variable values that are achieved by original actions, these actions (and therefore some other) may be added to the stubborn set, making it unnecessarily large. The issue can be avoided if a different transformation is used. Recently, such a transformation was proposed for a similar purpose, to preserve the pruning power of symmetries (Anonymous 2023a). Their transformation differs in the goal state (and consequently the effect of the goal achieving action), defining new goal values for all original variables. With that, the issue mentioned above disappears. The issue still persists in the states of Π_g that correspond to goal states of Π. There, the action o_g will be added to the stubborn set, and so will all other actions of Π_g, since they interfere with o_g. For any other non-goal state s of the transformed problem Π_g, such that a goal is reachable from s, let s' be the corresponding (non-goal) state of the original problem Π. Strong stubborn set algorithms applied to Π_g will produce the same *applicable* action set for s as the set obtained for s' when applied to Π. To see that, observe that o_g is the only action that achieves the goal in Π_g. It interferes with all other actions of Π_g and its precondition, except for on v_g agrees precisely with the goal of Π. All original actions interact with each other in Π_g in precisely the same way as in Π. Once any applicable action is added to the stubborn set of s, so will o_g (interference). Since o_g is not applicable, this will trigger adding the *necessary enabling set* for o_g. Assuming that the choice of finding an unsatisfied goal and of finding an unsatisfied precondition of o_g are resolved the same way, we get the same behavior.

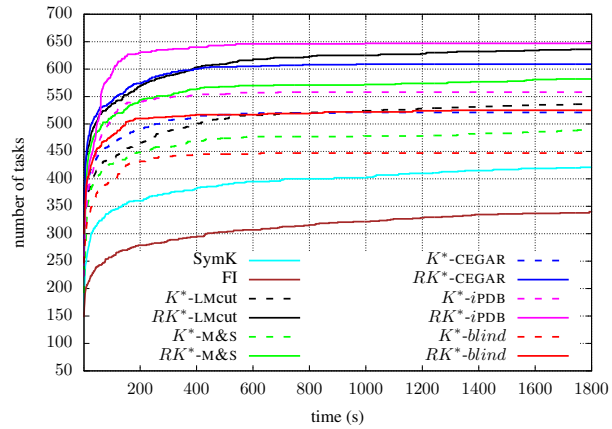

Figure 1: Anytime performance of tested configurations.

Experimental Evaluation

To empirically compare the suggested algorithm, we have integrated the existing implementation of partial order reduction techniques into an existing implementation of K^* algorithm (Anonymous 2023b), all built on top of the Fast Downward planning system (Helmert 2006). All experiments were performed on Intel(R) Xeon(R) Gold 6248 CPU @ 2.50GHz machines, with the timeout of 30 minutes and memory limit of 8GB per run. The benchmark set consists of all benchmarks from optimal tracks of International Planning Competitions 1998-2018, a total of 1827 tasks in 65 domains. We have experimented with four admissible heuristics, LMcut (Helmert and Domshlak 2009), merge-and-shrink abstraction (denoted by M&S) (Helmert, Haslum, and

				LMcut		M&S		CEGAR		iPDB		blind	
		SymK	FI	K^*	RK^*	K^*	RK^*	K^*	RK^*	K^*	RK^*	K^*	RK^*
	SymK	0	**31**	5	5	7	7	6	6	3	3	17	16
	FI	12	0	6	2	7	3	7	3	7	3	7	3
LMcut	K^*	**39**	**43**	0	1	**14**	13	**12**	13	6	6	**31**	**29**
	RK^*	**41**	**48**	**12**	0	**25**	**18**	**23**	**18**	**18**	11	**39**	**34**
M&S	K^*	**35**	**41**	7	6	0	1	11	10	0	0	**27**	**25**
	RK^*	**37**	**45**	**14**	6	**10**	0	**19**	14	9	4	**34**	**30**
CEGAR	K^*	**32**	**40**	4	4	**12**	12	0	1	5	5	**27**	**25**
	RK^*	**34**	**45**	**14**	4	**22**	14	**10**	0	15	6	**34**	**29**
iPDB	K^*	**44**	**43**	**17**	15	**19**	**19**	**19**	**19**	0	0	**35**	**33**
	RK^*	**46**	**47**	**26**	**16**	**29**	**21**	**28**	**21**	**10**	0	**42**	**37**
blind	K^*	**23**	**38**	2	1	6	6	2	1	1	1	0	0
	RK^*	**25**	**42**	8	2	13	7	8	2	8	1	**10**	0
Overall Coverage		421	341	536	636	489	582	521	609	558	**647**	447	525

Table 1: Pair-wise domain level comparison of (unordered) top-quality planners. Each entry in the table represents the number of domains where the row configuration achieves better coverage than the column one. The last row depicts the overall coverage.

Hoffmann 2007), counterexample-guided Cartesian abstraction refinement (denoted by CEGAR) (Seipp and Helmert 2018), and pattern database heuristic iPDB (Haslum et al. 2007). As previously mentioned, planning heuristics can be sensitive to task formulation. Thus, Anonymous (2023b) propose to evaluate the heuristics on the original task instead, showing significant increase in search performance. Here, we follow their suggestion. The found plans are added to the solution set, checking duplicates as de-ordered multisets. Vanilla K^* is compared to K^*, pruned by atom-centric stubborn sets (Röger et al. 2020), denoted by RK^*. Additionally, we compare our approach to the existing unordered top-quality planners, ForbidIterative (FI) (Katz et al. 2018) and SymK (Speck, Mattmüller, and Nebel 2020). Importantly, all planners perform the same translation from PDDL to SAS$^+$, solving the same SAS$^+$ planning problem. In our experiments, we disabled the actual writing of plans to disk, excluding the write time from our evaluation. All experiments are using the cost multiplier $q = 1$, which corresponds to finding all cost-optimal plans, modulo reorderings.

We measure the coverage of each configuration, giving the score of 1 to a planner for each task where the unordered top-quality solution was found withing the 30 minutes time bound. Additionally, we measure the time it took to find the solution. Figure 1 depicts the any-time overall coverage performance of the tested configurations. The lines for the same heuristic are depicted with the same color, solid line for RK^* and dashed line for K^*. For very small timeouts, up to 4 seconds, the best performing overall configuration is RK^* with the LMcut heuristic. From 4 to 60 seconds, RK^* with CEGAR takes the lead. Starting from 60 seconds an onwards, RK^* with iPDB becomes the best performer. It is worth noting that it reaches its almost maximal coverage of 646 out of 647 tasks already at 540 seconds and the maximal coverage by 995 seconds.

Switching now to domain-level coverage, Table 1 shows the pairwise comparison of the tested approaches. Each entry denotes the number of domains where the row planner achieves a better summed coverage than the column planner. Additionally, the last row denotes the overall coverage for each planner. The winners are marked in bold. For the pairwise comparison, the value in (x,y) is bolded if it is larger than the value in (y,x), that is planner x excelled over planner y in more domains than planner y excelled over planner x. Observe that RK^* consistently outperforms K^*, with the difference in the number of domains with superior performance being between 9 and 11. RK^* with iPDB is the absolute per-domain winner, as well as overall winner, achieving the overall coverage of 647. Compared to existing unordered top-quality planners, FI (341) and SymK (421), this is a tremendous increase in coverage of over 50%.

At per-domain level, depicted in Table 2, as one might expect, RK^* with iPDB does not absolutely dominate other approaches. There are 17 domains where the configuration does not achieve maximal coverage, including AIRPORT and LOGISTICS-IPC2, where FI excels and BLOCKSWORLDand VISIT-ALL where SymK achieves the largest coverage. Other heuristics also exhibit superior performance: LMcut in DRIVERLOG, ORGANIC-SYNTHESIS-SPLIT, PATHWAYS-NN, PIPESWORLD-NOTANKAGE, SATELLITE, TRUCKS, and WOODWORKING, M&S in PSR-SMALL and WOODWORKING, and CEGAR in FREECELL, MPRIME, MYSTERY, PIPESWORLD-NOTANKAGE, and TRUCKS. Interestingly, the highest coverage in ORGANIC-SYNTHESIS-SPLIT is achieved by K^* with LMcut, while RK^* gets a slightly lower coverage due to the overhead of the partial order reduction used.

			LMcut		M&S		CEGAR		iPDB		blind	
	SymK	FI	K^*	RK^*	K^*	RK^*	K^*	RK^*	K^*	RK^*	K^*	RK^*
airport (50)	7	**18**	7	8	7	8	7	8	7	8	7	8
barman11 (20)	**4**	0	**4**	**4**	**4**	**4**	**4**	**4**	**4**	**4**	**4**	**4**
blocks (35)	**30**	12	28	28	23	23	18	18	26	26	18	18
data-network18 (20)	5	1	**7**	**7**	**7**	**7**	**7**	**7**	**7**	**7**	5	5
depot (22)	2	2	**4**	**4**	**4**	**4**	**4**	**4**	**4**	**4**	3	3
driverlog (20)	8	3	**11**	**11**	10	10	9	9	10	10	5	5
floortile11 (20)	**2**	0	**2**	**2**	**2**	**2**	**2**	**2**	**2**	**2**	0	0
freecell (80)	9	2	15	15	15	15	**22**	**22**	16	16	14	14
ged14 (20)	13	7	**15**	**15**	**15**	**15**	**15**	**15**	**15**	**15**	13	13
grid (5)	2	1	2	2	2	2	2	2	**3**	**3**	1	1
gripper (20)	2	2	**3**	**3**	**3**	**3**	**3**	**3**	**3**	**3**	**3**	**3**
hiking14 (20)	2	0	**4**	**4**	**4**	**4**	**4**	**4**	**4**	**4**	**4**	**4**
logistics00 (28)	6	**10**	7	9	7	9	7	8	7	9	6	7
logistics98 (35)	2	1	3	**4**	3	3	3	**4**	3	**4**	2	2
miconic (150)	35	16	**43**	**43**	**43**	**43**	**43**	**43**	**43**	**43**	40	40
movie (30)	3	1	4	**27**	4	**27**	4	**27**	4	**27**	4	**27**
mprime (35)	16	13	23	23	20	20	**25**	**25**	20	20	14	14
mystery (30)	15	18	20	22	19	20	**25**	**25**	21	21	16	17
nomystery11 (20)	13	7	14	14	16	16	12	12	**20**	**20**	8	8
openstacks08 (30)	1	**2**	**2**	**2**	**2**	**2**	**2**	**2**	**2**	**2**	**2**	**2**
openstacks (30)	4	0	**5**	**5**	**5**	**5**	**5**	**5**	**5**	**5**	**5**	**5**
organic-synthesis-split18 (20)	8	10	**15**	14	2	2	9	9	9	9	10	10
parcprinter08 (30)	6	25	8	**30**	8	27	8	**30**	8	**30**	8	**30**
parcprinter11 (20)	3	18	4	**20**	4	**20**	4	**20**	4	**20**	4	**20**
parking11 (20)	0	0	1	1	2	2	0	0	**4**	**4**	0	0
parking14 (20)	0	0	3	3	**4**	**4**	0	0	**4**	**4**	0	0
pathways (30)	2	1	4	**5**	4	4	4	4	4	4	4	4
pegsol08 (30)	28	23	27	27	27	27	27	27	**29**	**29**	27	27
pegsol11 (20)	18	13	17	17	17	17	17	17	**19**	**19**	17	17
pipes-notank (50)	9	9	**16**	**16**	15	15	**16**	**16**	15	15	13	13
pipes-tank (50)	6	4	8	8	**9**	**9**	8	8	**9**	**9**	6	6
psr-small (50)	44	45	45	47	46	**49**	45	47	46	48	46	48
satellite (36)	5	2	5	**9**	5	7	5	7	5	6	4	5
scanalyzer08 (30)	8	4	**10**	**10**	**10**	**10**	8	8	**10**	**10**	6	6
scanalyzer11 (20)	5	1	**7**	**7**	**7**	**7**	5	5	**7**	**7**	3	3
snake18 (20)	3	2	7	7	8	7	**11**	**11**	**11**	**11**	**11**	**11**
spider18 (20)	0	3	11	11	6	6	11	11	**13**	**13**	11	11
storage (30)	14	11	**15**	**15**	**15**	**15**	**15**	**15**	**15**	**15**	13	13
termes18 (20)	5	0	6	6	**7**	**7**	5	5	**7**	**7**	3	3
tetris14 (17)	2	2	**5**	**5**	2	2	**5**	**5**	**5**	**5**	**5**	**5**
tidybot11 (20)	1	3	**14**	**14**	4	4	**14**	**14**	**14**	**14**	12	12
tidybot14 (20)	0	0	**8**	**8**	0	0	**8**	7	**8**	**8**	4	4
transport08 (30)	11	6	11	11	11	11	11	11	**14**	**14**	11	11
transport11 (20)	6	1	6	6	6	6	6	6	**11**	**11**	6	6
transport14 (20)	5	0	6	6	4	4	5	5	**7**	**7**	4	4
trucks (30)	2	2	**5**	**5**	4	4	**5**	**5**	4	4	3	3
visitall11 (20)	**12**	6	10	10	9	9	9	9	11	11	8	8
visitall14 (20)	**6**	0	5	5	4	4	3	3	5	5	2	2
woodworking08 (30)	6	9	7	**21**	7	**21**	7	18	7	18	7	13
woodworking11 (20)	1	4	2	**15**	2	**15**	2	12	2	12	2	7
zenotravel (20)	8	5	**9**	**9**	**9**	**9**	**9**	**9**	**9**	**9**	7	7
Other (324)	**16**	**16**	**16**	**16**	**16**	**16**	**16**	**16**	**16**	**16**	**16**	**16**
Overall (1827)	421	341	536	636	489	582	521	609	558	**647**	447	525

Table 2: Domain level coverage of (unordered) top-q planners for $q = 1$. Each entry in the table represents the number of tasks in the domain solved by the respective approach. The last two rows depict the aggregated coverage over domains where there is no difference between the tested configurations and the overall coverage.

Conclusions and Future Work

In this work, we exploit partial order reduction techniques, namely stubborn sets, to improve the performance of K^* search. For that, we show that no adaptation to these techniques is required if a single-goal planning task transformation is used. We point out the weaknesses of the transfor-

mation used in previous implementations of K^* and suggest using a recently proposed for the purpose of symmetry pruning variant, which preserves the pruning power of the original planning task. We perform an experimental evaluation that establishes our approach as the new state of the art for unordered top-quality planning.

In the future work, we intend to further improve top-k planning with partial order reduction. To do so, an efficient reconstruction of pruned reorderings must be derived. A previously suggested naive exhaustive DFS search (Katz et al. 2018) was found to be too slow to be competitive. Efficient methods can exploit the information on which actions were actually pruned at each state.

References

Aljazzar, H.; and Leue, S. 2011. K*: A Heuristic Search Algorithm for Finding the K Shortest Paths. *AIJ*, 175(18): 2129–2154.

Alkhazraji, Y.; Wehrle, M.; Mattmüller, R.; and Helmert, M. 2012. A Stubborn Set Algorithm for Optimal Planning. In *Proc. ECAI 2012*, 891–892.

Anonymous. 2023a. K* Search Over Orbit Space for Top-k Planning. Submitted to *IJCAI*. Under review.

Anonymous. 2023b. On K* Search for Top-k Planning. Submitted to *SoCS*. Under review.

Bäckström, C.; and Nebel, B. 1995. Complexity Results for SAS$^+$ Planning. *Computational Intelligence*, 11(4): 625–655.

Boddy, M.; Gohde, J.; Haigh, T.; and Harp, S. 2005. Course of Action Generation for Cyber Security Using Classical Planning. In *Proc. ICAPS 2005*, 12–21.

Chakraborti, T.; Isahagian, V.; Khalaf, R.; Khazaeni, Y.; Muthusamy, V.; Rizk, Y.; and Unuvar, M. 2020. From Robotic Process Automation to Intelligent Process Automation. In *Proc. BPM 2020: Blockchain and RPA Forum*, 215–228.

Eppstein, D. 1998. Finding the k shortest paths. *SICOMP*, 28(2): 652–673.

Gnad, D.; Hoffmann, J.; and Wehrle, M. 2019. Strong Stubborn Set Pruning for Star-Topology Decoupled State Space Search. *JAIR*, 65: 343–392.

Haslum, P.; Botea, A.; Helmert, M.; Bonet, B.; and Koenig, S. 2007. Domain-Independent Construction of Pattern Database Heuristics for Cost-Optimal Planning. In *Proc. AAAI 2007*, 1007–1012.

Helmert, M. 2006. The Fast Downward Planning System. *JAIR*, 26: 191–246.

Helmert, M.; and Domshlak, C. 2009. Landmarks, Critical Paths and Abstractions: What's the Difference Anyway? In *Proc. ICAPS 2009*, 162–169.

Helmert, M.; Haslum, P.; and Hoffmann, J. 2007. Flexible Abstraction Heuristics for Optimal Sequential Planning. In *Proc. ICAPS 2007*, 176–183.

Katz, M.; Ram, P.; Sohrabi, S.; and Udrea, O. 2020. Exploring Context-Free Languages via Planning: The Case for Automating Machine Learning. In *Proc. ICAPS 2020*, 403–411.

Katz, M.; and Sohrabi, S. 2022. Who Needs These Operators Anyway: Top Quality Planning with Operator Subset Criteria. In *Proc. ICAPS 2022*.

Katz, M.; Sohrabi, S.; and Udrea, O. 2020. Top-Quality Planning: Finding Practically Useful Sets of Best Plans. In *Proc. AAAI 2020*, 9900–9907.

Katz, M.; Sohrabi, S.; and Udrea, O. 2022. Bounding Quality in Diverse Planning. In *Proc. AAAI 2022*.

Katz, M.; Sohrabi, S.; Udrea, O.; and Winterer, D. 2018. A Novel Iterative Approach to Top-k Planning. In *Proc. ICAPS 2018*, 132–140.

Keren, S.; Gal, A.; and Karpas, E. 2018. Strong Stubborn Sets for Efficient Goal Recognition Design. In *Proc. ICAPS 2018*, 141–149.

Nguyen, T. A.; Do, M. B.; Gerevini, A.; Serina, I.; Srivastava, B.; and Kambhampati, S. 2012. Generating diverse plans to handle unknown and partially known user preferences. *AIJ*, 190: 1–31.

Röger, G.; Helmert, M.; Seipp, J.; and Sievers, S. 2020. An Atom-Centric Perspective on Stubborn Sets. In *Proc. SoCS 2020*, 57–65.

Seipp, J.; and Helmert, M. 2018. Counterexample-Guided Cartesian Abstraction Refinement for Classical Planning. *JAIR*, 62: 535–577.

Sievers, S.; and Wehrle, M. 2021. On Weak Stubborn Sets in Classical Planning. In *Proc. IJCAI 2021*, 4167–4174.

Sohrabi, S.; Riabov, A. V.; and Udrea, O. 2016. Plan Recognition as Planning Revisited. In *Proc. IJCAI 2016*, 3258–3264.

Speck, D.; Mattmüller, R.; and Nebel, B. 2020. Symbolic Top-k Planning. In *Proc. AAAI 2020*, 9967–9974.

Vadlamudi, S. G.; and Kambhampati, S. 2016. A Combinatorial Search Perspective on Diverse Solution Generation. In *Proc. AAAI 2016*, 776–783.

Valmari, A. 1989. Stubborn sets for reduced state space generation. In *Proc. APN 1989*, 491–515.

Wehrle, M.; and Helmert, M. 2012. About Partial Order Reduction in Planning and Computer Aided Verification. In *Proc. ICAPS 2012*, 297–305.

Wehrle, M.; and Helmert, M. 2014. Efficient Stubborn Sets: Generalized Algorithms and Selection Strategies. In *Proc. ICAPS 2014*, 323–331.

Wilhelm, A.; Steinmetz, M.; and Hoffmann, J. 2018. On Stubborn Sets and Planning with Resources. In *Proc. ICAPS 2018*, 288–297.

Winterer, D.; Alkhazraji, Y.; Katz, M.; and Wehrle, M. 2017. Stubborn Sets for Fully Observable Nondeterministic Planning. In *Proc. ICAPS 2017*, 330–338.