SymK – A Versatile Symbolic Search Planner

David Speck

Linköping University, Linköping, Sweden
david.speck@liu.se

Abstract

SymK is a planner that performs symbolic search using Binary Decision Diagrams to find a single optimal or the best k plans. It is designed to be a versatile planner by supporting several important and expressive extensions to the classical planning formalism. Our planner, SymK, therefore natively supports features relevant for compact modeling of planning tasks, such as conditional effects and derived predicates with axioms.

Introduction

Symbolic search is a state-space exploration technique that originated in model checking (McMillan 1993). Symbolic search algorithms are similar to their explicit counterparts. However, they expand and generate entire sets of states rather than individual states. Over the years, symbolic search has proven to be a highly competitive approach to optimal planning, yielding impressive results at previous International Planning Competitions (Edelkamp and Helmert 2001; Torralba et al. 2014; Kissmann, Edelkamp, and Hoffmann 2014; Edelkamp, Kissmann, and Torralba 2015; Torralba et al. 2017; Speck, Geißer, and Mattmüller 2018b; Franco et al. 2018).

One strength of symbolic search is that it does not require strong heuristics to be competitive, since symbolic bidirectional blind search is among the strongest search strategies (Torralba et al. 2017; Speck, Geißer, and Mattmüller 2020; Fišer, Torralba, and Hoffmann 2022). For this reason, symbolic search does not necessarily suffer from the restriction of strong state-of-the-art heuristics to the planning formalism. With this in mind, the SymK planner was developed with the goal of being a versatile symbolic search planner that supports several expressive extensions of traditional classical planning while retaining the core of the formalism (Speck 2022). Among other things, SymK can find multiple optimal solutions or even all solutions for a given planning task (Speck, Mattmüller, and Nebel 2020; von Tschammer, Mattmüller, and Speck 2022), and supports conditional effects, oversubscribed goal descriptions (Speck and Katz 2021), state-dependent action costs (Speck, Geißer, and Mattmüller 2018a), and complex state descriptions with derived predicates and axioms (Speck et al. 2019). Despite the broad feature support, SymK also implements a symbolic search that is tailored to efficiently search for a single optimal solution for a given classical planning task.

The following describes the details of the SymK configuration submitted to the *optimal track* of the 2023 International Planning Competition to find a single optimal solution, and which and how SymK supports the PDDL language features of the competition.

Implementation

SymK is based on Fast Downward 22.06 (Helmert 2006) and SymBA* (Torralba et al. 2014). For preprocessing, we use the h^2 preprocessor for invariant computation and spurious action pruning (Alcázar and Torralba 2015). For the competition, we chose to perform a bidirectional symbolic blind search, which is known to be one of the dominant search strategies for symbolic search (Torralba et al. 2017; Speck, Geißer, and Mattmüller 2020). At each search iteration, either a forward or a backward search step is performed. To decide which direction is more promising, the runtime of the last forward step is compared to the runtime of the last backward step. To represent formulas, sets of states, and transition relations, we use Binary Decision Diagrams (BDDs) (Bryant 1986) of the CUDD library (Somenzi 2015). We also use a fixed variable ordering based on an analysis of the causal graph known as the Gamer variable ordering described in Kissmann and Hoffmann (2013, 2014). Finally, we combine as many actions as possible into a transition relation until the BDD representation exceeds 100k nodes, perform state set partitioning based on the resulting transition relations, and use mutexes to prune spurious states during search (Torralba and Alcázar 2013; Torralba, Edelkamp, and Kissmann 2013; Torralba 2015; Torralba et al. 2017).

Language Support

SymK supports *PDDL 2.2 Level 1* (Fox and Long 2003) plus the *action cost requirement* from PDDL 3.1 and *all ADL features* such as quantified and conditional effects and negation, disjunction, and quantification in conditions. In particular, SymK natively supports conditional effects and derived predicates with axioms, which are rarely supported by optimal planners. SymK supports *conditional effects* by encoding them directly in the transition relations as described in Kissmann, Edelkamp, and Hoffmann (2014). *Derived pred-*

icates and axioms are supported by SymK using the symbolic translation approach of Speck et al. (2019), where all occurrences of derived predicates in the planning task are replaced by their corresponding primary representation using BDDs as the underlying data structure. Finally, beyond the IPC requirements, SymK implements top-k planning to generate *many* or even *all plans* with increasing costs as an output stream and supports *state-dependent action costs*.

Post-Competition Performance Analysis

After the great successes of planners, based heavily or even exclusively on symbolic search and symbolic data structures in the optimal track of the last two International Planning Competitions, namely with the A* symbolic bidirectional planner SymBA* (Torralba et al. 2014) which won in 2014 and the Complementary planner based on symbolic pattern databases (Franco et al. 2018) which came second in 2018, this year the results for symbolic search planners are mixed, including the result for SymK.

The results are mixed, mainly because planners that relied heavily on symbolic search and symbolic data structures, such as the 2023 versions of SymK, ComplementaryPDB (Franco, Edelkamp, and Moraru 2023), and SymBD (Torralba 2023), found themselves in the middle of the scoreboard rather than in the winning range. However, it's worth noting that similar to 2018, when symbolic search in the form of SymBA* played a critical role in the Delfi portfolio planner that topped the leaderboard, SymK was a critical component of the winning Ragnarok portfolio planner in this edition of the competition.

Several factors underlie the comparatively modest performance of planners relying solely on symbolic search and symbolic data structures in this edition of the IPC. One factor is the significant challenges posed by some of this year's competition domains, which were characterized by exceptionally large task characteristics, such as the vast number of conditional effects in the Rubik's Cube domain. Another factor is the difficulty of competing with portfolio planners that include stand-alone planners such as SymK as components, which then simply outperform the included stand-alone planners.

In summary, symbolic search played a role in the optimal track of the 2023 IPC, but less in the form of a stand-alone planner and more as an integral part of portfolio planning systems. In addition, the introduction of the new domains has opened up exciting research directions using symbolic search and data structures.

Conclusions

SymK is a planner based on symbolic search. It focuses on optimal planning while supporting extensions to the classical planning formalism. It is a new planner in the sense that SymK has never participated in a previous International Planning Competition (IPC), although SymBA* and Fast Downward, on which SymK is based, have. For the competition, we chose to use symbolic bidirectional blind search, for which several optimizations have been published over the years, which we have summarized in this planner abstract.

While SymK as a stand-alone planner did not place at the top of the IPC 2023 optimal track leaderboard, it was an essential component of the Ragnarok planner (Drexler et al. 2023) that won the optimal track. The latest version of SymK is available online:

`https://github.com/speckdavid/symk`

Acknowledgments

We would like to thank all the contributors to Fast Downward and Álvaro Torralba and Vidal Alcázar for the h^2 preprocessor. Special thanks to Daniel Fišer and Florian Pommerening for organizing the competition.

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. This research was partially supported by TAILOR, a project funded by the EU Horizon 2020 research and innovation programme under grant agreement no. 952215. The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS) partially funded by the Swedish Research Council through grant agreement no. 2022-06725.

References

Alcázar, V.; and Torralba, Á. 2015. A Reminder about the Importance of Computing and Exploiting Invariants in Planning. In Brafman, R.; Domshlak, C.; Haslum, P.; and Zilberstein, S., eds., *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling (ICAPS 2015)*, 2–6. AAAI Press.

Bryant, R. E. 1986. Graph-Based Algorithms for Boolean Function Manipulation. *IEEE Transactions on Computers*, 35(8): 677–691.

Drexler, D.; Gnad, D.; Höft, P.; Seipp, J.; Speck, D.; and Ståhlberg, S. 2023. Ragnarok. In *Tenth International Planning Competition (IPC-10): Planner Abstracts*.

Edelkamp, S.; and Helmert, M. 2001. The Model Checking Integrated Planning System (MIPS). *AI Magazine*, 22(3): 67–71.

Edelkamp, S.; Kissmann, P.; and Torralba, Á. 2015. BDDs Strike Back (in AI Planning). In Bonet, B.; and Koenig, S., eds., *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI 2015)*, 4320–4321. AAAI Press.

Fišer, D.; Torralba, Á.; and Hoffmann, J. 2022. Operator-Potential Heuristics for Symbolic Search. In Honavar, V.; and Spaan, M., eds., *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2022)*, 9750–9757. AAAI Press.

Fox, M.; and Long, D. 2003. PDDL2.1: An Extension to PDDL for Expressing Temporal Planning Domains. *Journal of Artificial Intelligence Research*, 20: 61–124.

Franco, S.; Edelkamp, S.; and Moraru, I. 2023. ComplementaryPDB Planner. In *Tenth International Planning Competition (IPC-10): Planner Abstracts*.

Franco, S.; Lelis, L. H. S.; Barley, M.; Edelkamp, S.; Martines, M.; and Moraru, I. 2018. The Complementary1 Planner in the IPC 2018. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 28–31.

Helmert, M. 2006. The Fast Downward Planning System. *Journal of Artificial Intelligence Research*, 26: 191–246.

Kissmann, P.; Edelkamp, S.; and Hoffmann, J. 2014. Gamer and Dynamic-Gamer – Symbolic Search at IPC 2014. In *Eighth International Planning Competition (IPC-8): Planner Abstracts*, 77–84.

Kissmann, P.; and Hoffmann, J. 2013. What's in It for My BDD? On Causal Graphs and Variable Orders in Planning. In Borrajo, D.; Kambhampati, S.; Oddi, A.; and Fratini, S., eds., *Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling (ICAPS 2013)*, 327–331. AAAI Press.

Kissmann, P.; and Hoffmann, J. 2014. BDD Ordering Heuristics for Classical Planning. *Journal of Artificial Intelligence Research*, 51: 779–804.

McMillan, K. L. 1993. *Symbolic Model Checking*. Kluwer Academic Publishers.

Somenzi, F. 2015. CUDD: CU Decision Diagram Package – Release 3.0.0. https://github.com/ivmai/cudd. Accessed: 2023-02-20.

Speck, D. 2022. *Symbolic Search for Optimal Planning with Expressive Extensions*. Ph.D. thesis, University of Freiburg.

Speck, D.; Geißer, F.; and Mattmüller, R. 2018a. Symbolic Planning with Edge-Valued Multi-Valued Decision Diagrams. In de Weerdt, M.; Koenig, S.; Röger, G.; and Spaan, M., eds., *Proceedings of the Twenty-Eighth International Conference on Automated Planning and Scheduling (ICAPS 2018)*, 250–258. AAAI Press.

Speck, D.; Geißer, F.; and Mattmüller, R. 2018b. SYMPLE: Symbolic Planning based on EVMDDs. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 91–94.

Speck, D.; Geißer, F.; and Mattmüller, R. 2020. When Perfect Is Not Good Enough: On the Search Behaviour of Symbolic Heuristic Search. In Beck, J. C.; Karpas, E.; and Sohrabi, S., eds., *Proceedings of the Thirtieth International Conference on Automated Planning and Scheduling (ICAPS 2020)*, 263–271. AAAI Press.

Speck, D.; Geißer, F.; Mattmüller, R.; and Torralba, Á. 2019. Symbolic Planning with Axioms. In Lipovetzky, N.; Onaindia, E.; and Smith, D. E., eds., *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling (ICAPS 2019)*, 464–472. AAAI Press.

Speck, D.; and Katz, M. 2021. Symbolic Search for Oversubscription Planning. In Leyton-Brown, K.; and Mausam, eds., *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021)*, 11972–11980. AAAI Press.

Speck, D.; Mattmüller, R.; and Nebel, B. 2020. Symbolic Top-k Planning. In Conitzer, V.; and Sha, F., eds., *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*, 9967–9974. AAAI Press.

Torralba, Á. 2015. *Symbolic Search and Abstraction Heuristics for Cost-Optimal Planning*. Ph.D. thesis, Universidad Carlos III de Madrid.

Torralba, Á. 2023. SymBD: A Symbolic Bidirectional Search Baseline. In *Tenth International Planning Competition (IPC-10): Planner Abstracts*.

Torralba, Á.; and Alcázar, V. 2013. Constrained Symbolic Search: On Mutexes, BDD Minimization and More. In Helmert, M.; and Röger, G., eds., *Proceedings of the Sixth Annual Symposium on Combinatorial Search (SoCS 2013)*, 175–183. AAAI Press.

Torralba, Á.; Alcázar, V.; Borrajo, D.; Kissmann, P.; and Edelkamp, S. 2014. SymBA*: A Symbolic Bidirectional A* Planner. In *Eighth International Planning Competition (IPC-8): Planner Abstracts*, 105–109.

Torralba, Á.; Alcázar, V.; Kissmann, P.; and Edelkamp, S. 2017. Efficient Symbolic Search for Cost-optimal Planning. *Artificial Intelligence*, 242: 52–79.

Torralba, Á.; Edelkamp, S.; and Kissmann, P. 2013. Transition Trees for Cost-Optimal Symbolic Planning. In Borrajo, D.; Kambhampati, S.; Oddi, A.; and Fratini, S., eds., *Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling (ICAPS 2013)*, 206–214. AAAI Press.

von Tschammer, J.; Mattmüller, R.; and Speck, D. 2022. Loopless Top-K Planning. In Thiébaux, S.; and Yeoh, W., eds., *Proceedings of the Thirty-Second International Conference on Automated Planning and Scheduling (ICAPS 2022)*, 380–384. AAAI Press.