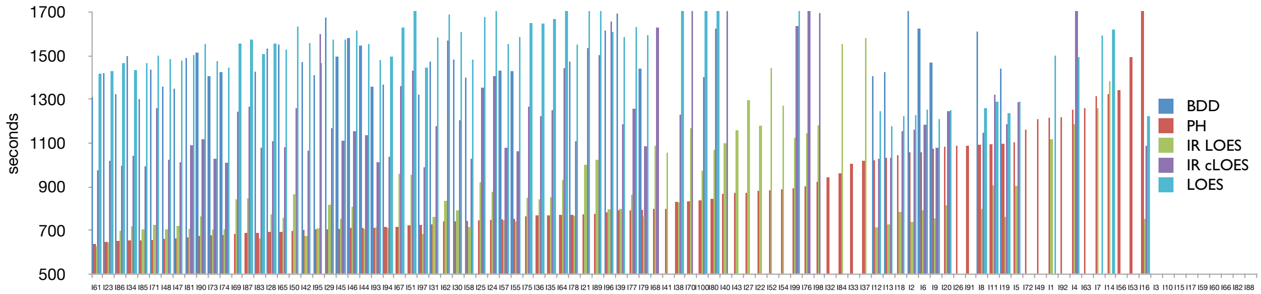


Figure 15: Planning time (parsing, analysis, PDB creation and search) over Korf's 100 instances of the 15-puzzle. $I23$ stands for instance 23. If a representation has no bar for an instance, it failed to terminate successfully in 30 minutes.

BDDs. We ran all instances with a hard 30 minute cut-off timer. Table 2 gives the results. Here PH fared the best, thanks to its very quick PDB lookups. While the LOES variants offered a noticeable relative reduction in PDB size, the absolute differences were relatively small. The results would probably change if the analysis component allowed larger PDBs (to the detriment of the BDD based representation). Figure 15 gives an overview of the total planning time over all hundred instances.

Table 2: Number of instances solved, PDB size and relative search speed over Korf's hundred 15-puzzle instances.

| Rep. | $\#_{solved}$ | size (MB) | $\frac{t_x}{t_{PH}}$ |
|---|---|---|---|
| PH | 91 | 20.0 | 1.00 |
| BDD | 40 | 117.6 | 0.11 |
| IR LOES | 82 | 11.4 | 0.36 |
| IR cLOES | 70 | 9.6 | 0.23 |
| LOES | 68 | 11.2 | 0.49 |

## Conclusion and Future Work

We believe techniques such as LOES offer exciting opportunities to better exploit dynamic programming and other memoization techniques in domain-independent planning. The approach allows for quite efficient representation of strong, precomputed heuristics. The very basic domain analysis we employed in our evaluation can only give a hint of the potential for ad-hoc abstraction in heuristic search. Of interest is also LOES' impedance match with BDDs. The inverse relation representation straightforwardly allows to adaptively interchange BDD and LOES based representations of state-sets. In this way, a domain-independent planner can leverage the superior efficiency of BDDs in appropriate domains while mitigating their lack of robustness by falling back to a LOES-based representation.

## References

Ball, M., and Holte, R. 2008. The compression power of symbolic pattern databases. In *International Conference on Automated Planning and Scheduling (ICAPS)*, 2–11.

Culberson, J., and Schaeffer, J. 1998. Pattern databases. *Computational Intelligence* 14(3):318–334.

Drager, K.; Finkbeiner, B.; and Podelski, A. 2006. Directed model checking with distance-preserving abstractions. In *Model checking software: 13th International SPIN Workshop, Vienna, Austria, March 30-April 1, 2006: proceedings*, 19. Springer-Verlag New York Inc.

Edelkamp, S. 2002. Symbolic pattern databases in heuristic search planning. In *Proceedings of the 6th International Conference on AI Planning and Scheduling (AIPS-02)*, 274–283.

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 National Conference on Artificial Intelligence*, volume 22, 1007. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999.

Haslum, P.; Bonet, B.; and Geffner, H. 2005. New admissible heuristics for domain-independent planning. In *Proceedings of the National Conference on Artificial Intelligence*, volume 20, 1163. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999.

Helmert, M.; Haslum, P.; and Hoffmann, J. 2007. Flexible abstraction heuristics for optimal sequential planning. In *Proc. ICAPS*, volume 2007, 176–183.

Helmert, M. 2006. The fast downward planning system. *Journal of Artificial Intelligence Research* 26(2006):191–246.

Jensen, R.; Bryant, R.; and Veloso, M. 2002. SetA*: An efficient bdd-based heuristic search algorithm. In *Proceeding of the 18th National Conference on Artificial Intelligence (AAAI-02)*, 668–673.

Korf, R. 1985. Depth-first iterative-deepening: An optimal admissible tree search. *Artificial intelligence*.

Schmidt, T., and Zhou, R. 2011. Succinct set-encoding for state-space search. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence*.