

Complete Solution of the Eight-Puzzle and the Benefit of Node Ordering in IDA*

Alexander Reinefeld

Paderborn Center for Parallel Computing
Warburger Str. 100, D-33095 Paderborn, Germany

Abstract

The 8-puzzle is the largest puzzle of its type that can be completely solved. It is simple, and yet obeys a combinatorially large problem space of $9!/2$ states. The $N \times N$ extension of the 8-puzzle is NP-hard.

In the first part of this paper, we present complete statistical data based on an exhaustive evaluation of all possible tile configurations. Our results include data on the expected solution lengths, the ‘easiest’ and ‘worst’ configurations, and the density and distribution of solution nodes in the search tree.

In our second set of experiments, we used the 8-puzzle as a workbench model to evaluate the benefit of node ordering schemes in Iterative-Deepening A* (IDA*). One highlight of our results is that almost all IDA* implementations perform worse than would be possible with a simple random ordering of the operators.

1 Introduction

The 8-puzzle is a prominent workbench model for measuring the performance of heuristic search algorithms [Gaschnig, 1979; Nilsson, 1980; Pearl, 1985; Russell, 1992], learning methods [Laird *et al.*, 1987] and the use of macro operators [Korf, 1985a]. It is simple, but has a combinatorially large problem space of $9!/2$ states. The 8-puzzle is the largest possible N -puzzle that can be completely solved. There exist larger variants, e.g. the 15-puzzle [Johnson and Storey, 1879], which can also be solved [Korf, 1985b], but not to completion. The general $N \times N$ extension of the 8-puzzle is NP-hard [Ratner and Warmuth, 1986].

Complete Solution. We enumerated all $9!/2$ tile configurations and computed all optimal (shortest) solution paths for all problem instances with a fast iterative-deepening search algorithm. Aside from Schofield’s [1967] analysis of a weaker 8-puzzle variant, our work is the first that gives complete statistical data on this application. Our results (for a quick overview see the table in the Conclusions) includes statistical data on

- the lengths of optimal solutions,

- the solution density in the search tree,
- the “hardest” and “easiest” configurations,
- the average heuristic branching factor,
- the distribution of goal nodes in the search frontier.

Benefit of Node Ordering in IDA*. Our experiments were also motivated by the need for a better understanding of the effectiveness of node ordering schemes in the IDA* search algorithm [Korf, 1985b]. Like any other iterative-deepening search, IDA* benefits by a good node expansion order, which reduces the time spent in the last (goal) iteration. Common ordering techniques include basic hill-climbing methods, the longest-path (principal variation) heuristic, the history heuristic and transposition tables. While some of these methods seem to work in practice (e.g. for the traveling salesman problem) it remained unclear how effective they are and which combination of the methods performs best. Our results give evidence that

- chances to find a solution early in the last iteration are extremely poor for common IDA* implementations; simple random node selection is better,
- steepest ascent hill-climbing does not help,
- a “history” table that holds a success score for each operator is better than random ordering,
- exploring the longest path first yields good performance while requiring only $O(\text{depth})$ storage space,
- the benefits of transposition tables are mainly due to caching previously expanded nodes rather than to an improved operator ordering.

2 The 8-Puzzle

The objective of the 8-puzzle is to rearrange a given initial configuration of eight squared tiles on a 3×3 board into a specified goal configuration by successively sliding tiles into the orthogonally adjacent empty square (the *blank* square). While it would seem easy to find any solution to this problem, we are only interested in obtaining *optimal* solutions with the fewest moves. We take, by convention, the following configuration to be the goal state:




Figure 1: Distribution of optimal solution lengths f^*

	1	2
3	4	5
6	7	8

There exist $9!$ possible tile permutations on a 3×3 board, and every second permutation is solvable [Johnson and Storey, 1879]. Hence, there is a total of $9!/2 = 181\,440$ solvable problem instances.

A weaker variant of the 8-puzzle has been investigated by Schofield [1967] with a mixed analytical/empirical approach. His 8-puzzle variant includes only mappings from one ‘standard position’ into another, where a ‘standard position’ is one with the blank located in the middle square. The search space of this weaker puzzle has only $8!/2 = 20\,160$ states. Our definition of the 8-puzzle, in contrast, conforms to the larger 15- and 25-puzzles, which have the blank in the upper left square and which do not restrict the search space to ‘standard positions’.

3 Complete Solution

We generated all $9!/2$ solvable tile configurations and computed all optimal solutions for all problem instances. We used Korf’s IDA* algorithm (for a description see Section 4) with the Manhattan distance as a heuristic estimate function. The whole experiment took less than one hour CPU-time on a SUN-SparcStation.

Solution Lengths. The average length of all optimal solution paths is $f^* = 21.97$, that is, almost 22 moves are needed to solve a given random configuration. Figure 1 shows the distribution of the optimal solution path lengths for all tile configurations. The two cases with the




Figure 2: Number of solutions per problem

shortest (non-trivial) solution path $f^* = 1$ have either of tile 1 or 3 in the upper left corner:

1		2
3	4	5
6	7	8

 $h = 1$
 $f^* = 1$

3	1	2
	4	5
6	7	8

 $h = 1$
 $f^* = 1$

The two configurations with the longest optimal solution path $f^* = 31$ are:

8	7	6
	4	1
2	5	3

 $h = 21$
 $f^* = 31$

8		6
5	4	7
2	3	1

 $h = 21$
 $f^* = 31$

Interestingly, these configurations do not build the largest trees. Their (complete) search trees have both 15890 leaves, whereas the ‘hardest’ configurations spawn trees with three times as many nodes. Similar to other NP-complete problems [Cheeseman *et al.*, 1991], the 8-puzzle has a small number of very hard problem instances which require significantly more leaf node expansions than average.

Solutions per Problem. For the 181440 solvable configurations, we found a total of 500880 optimal solutions. This gives an average solution density of 2.76 per problem, with the minimum and maximum number lying at 1 and 64 solutions, respectively (see Figure 2). About 3/4 of the problem instances have only one, two or three solutions. The two configurations with the most (64) solutions are:

8	5	6
7	2	3
4	1	

 $h = 20$
 $f^* = 30$

8	5	4
7	6	3
2	1	

 $h = 20$
 $f^* = 30$

The tiles are almost reversely ordered in these configurations, resulting in long solution paths with many move transpositions. The search tree is relatively large (with

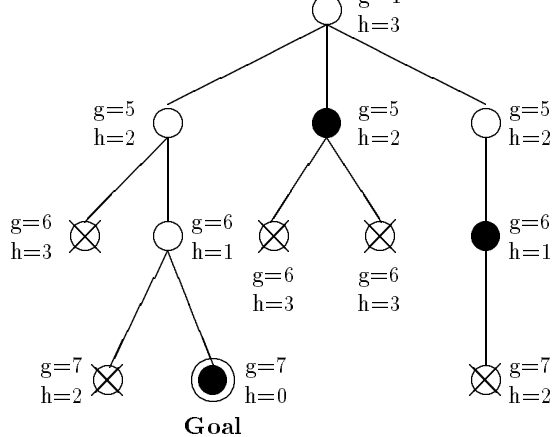


Figure 3: Sample tree searched with cost-bound 7. Filled circles are leaves.

11 333 total leaves), but it is not the largest (which has 41 794 leaves). There are 50 problems with more than 20 000 leaves on the optimal search frontier.

4 Benefit of Node Ordering in IDA*

Iterative-Deepening A*. IDA* [Korf, 1985b] is an efficient search method for applications with high heuristic branching factors and low solution densities [Rao *et al.*, 1991]. Typical applications of IDA* include single agent games like the N -puzzle, some variants of the traveling salesman problem, floorplan optimization and the cutting stock problem.

IDA* performs a series of depth-first searches with successively increased cost-bounds. The total cost $f(n)$ of a node n is made up of the cost already spent in reaching that node $g(n)$, plus the estimated cost of the path to a goal state $h(n)$. At the beginning, the cost bound is set to the heuristic estimate of the initial state, $h(\text{root})$. Then, for each iteration, the bound is increased to the minimum path value that exceeded the previous bound:

```

procedure IDA* ( $n$ );
 $bound := h(n)$ ;           { initial bound is  $h(\text{root})$  }
while not solved do     { iterate until solved ... }
     $bound := DFS(n, bound)$ ; { with increased cost bound }

```

```

function DFS ( $n, bound$ );
if  $f(n) > bound$ 
    then return  $f(n)$ ;    { path cost exceeds bound }
if  $h(n) = 0$ 
    then return solved;  { goal found }
return lowest value of  $DFS(n_i, bound)$  for all succ.  $n_i$  of  $n$ 

```

With an admissible (non-overestimating) heuristic function h , IDA* is guaranteed to find an optimal (shortest) solution path [Korf, 1985b]. Moreover, IDA* obeys the same asymptotic branching factor as A*, if the number of newly expanded nodes grows exponentially with the search depth [Korf, 1985b; Mahanti *et al.*, 1992]. This growth rate, the *heuristic branching factor* b_h , is defined as the node ratio of two consecutive iterations.

It depends on the number of applicable operators per node (the *edge branching factor* b_e) and the discrimination power of the heuristic estimate function h . Using the *Manhattan distance* (the sum of the minimum displacement of each tile from its goal position) as a heuristic estimate function, we computed $b_h = 3.81$ for the 8-puzzle. The corresponding value for the 15-puzzle, $b_h \approx 6.68$, is much higher because there are more interior squares with more move options [Reinefeld and Marsland, 1991].

The time to solve a problem is dominated by the time spent in the final (goal) iteration, which in turn depends on the expansion order of the leaves. A node is said to be *expanded*, when all its successors have been generated, or when it is a goal node. A *leaf* is an expanded node that has no expanded successors. Figure 3 illustrates a sample tree with three leaves. They all lie on a *search frontier* of an IDA* search with cost bound 7. We are primarily interested in *optimal* search frontiers that contain one or more goal nodes, as shown in the example.

Normalization. We want to find out how node ordering techniques improve the chances to find a solution early in the search frontier. For this purpose, we run IDA* on all 8-puzzle configurations and gathered statistical data on the location of goal nodes in the search front. In order to allow a direct comparison of the experimental data, we normalized the search frontiers (which vary in size between 1 and 41 794 leaves) to a common scale of 101 discrete data points. Here, 0 and 100 correspond to the location of the leftmost and the rightmost leaf node, respectively. When the leaves are numbered from left (= 1) to right (= n), the i -th leaf lies on position $\frac{i-1}{n-1}$ in the normalized search frontier.

For trees with more than 100 leaves, normalization is no problem. Care must be taken when there are only few leaves, because this causes the values to cluster at certain data points. As an example, the first leaf in Figure 3 lies at location $0/2 = 0$, the second at $1/2 = 50\%$ and the third at $2/2 = 100\%$ of the normalized search frontier.

Fixed Successor Ordering. For practical reasons, common IDA* implementations generate the node successors in an arbitrary, but fixed sequence (e.g. up, left, right, down). With such a scheme, one would expect the goal nodes to be evenly distributed over the whole search frontier. Surprisingly, this is not the case. Figure 4 depicts the distribution of all goal nodes in the last (deepest) search frontier ('all-solution-case'). The continuous plot shows the distribution of goal nodes with a fixed successor ordering. In the middle part, each goal location occurs ≈ 5000 times. This is as it should be, because the occurrences of all 101 goal locations must add up to the total solution number 500 880. The zigzag appearance with extreme points at specific locations ($\frac{1}{4}, \frac{1}{3}, \frac{1}{2}, \dots$) is caused by the normalization of small search frontiers with < 100 leaves, as discussed above.

More interesting are the 'shoulders' at both sides of the graph. They indicate a lower solution density in the left and right parts of the search frontier. This interesting phenomenon reveals a serious deficiency of our – and presumably most other – IDA* implementation: To

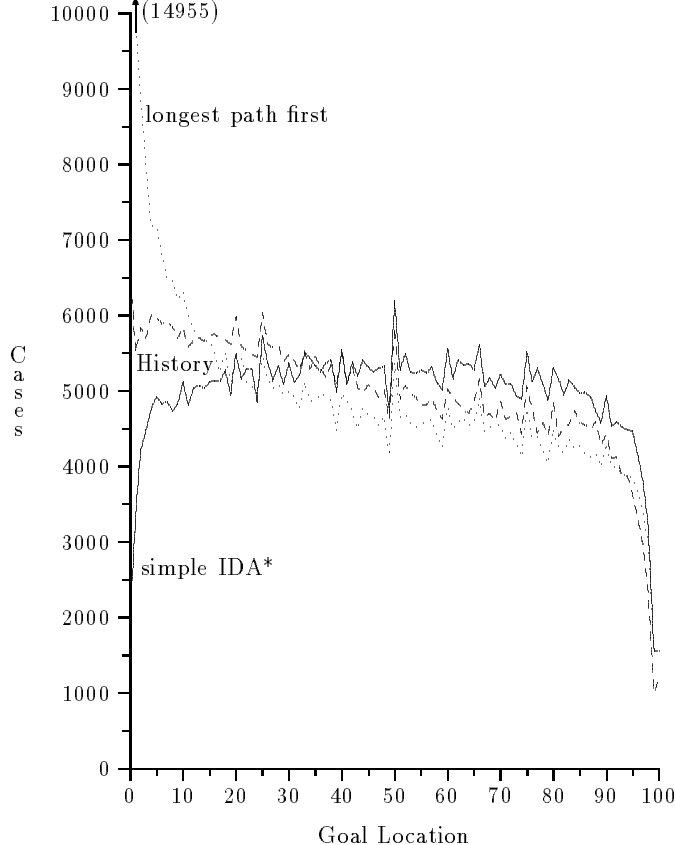


Figure 4: Location of *All* Solutions in the Search Front

be time efficient, common implementations use a table driven operator selection scheme, which applies the operators always in the same order. Our program (and also Korf's [1985b]) first tries to move the blank up, then to the left, to the right and finally down. If done in a row, this results in a circular move path which finally runs into a dead end situation. Hence, chances to find a solution in the first path of the search tree are low.

Random Node Selection is Better. Figure 4 clearly indicates, that a fixed operator sequence is worse than average. The chances to hit a solution right at the beginning of the search tree can be greatly improved by randomly perturbing the operator sequence in each node. The resulting solution density graph (not shown here) is almost flat in the left part, as it should be.

Expanding the Longest Path First. While simple random node ordering outperforms any kind of fixed operator sequence, we can do much better by applying domain dependent heuristics. One possibility is to investigate the deepest path of the previous iteration first. In multi-agent games, this path is called *principal variation*. It is the preferred move sequence, if both parties adhere to the minimax principle. In single-agent (IDA*) search, the longest path is the one that came closest to the goal. Since all nodes in the search frontier have the same f -value, the leaves on the longest paths have the highest g - and consequently lowest h -values.

The dotted graph in Figure 4 shows a much increased

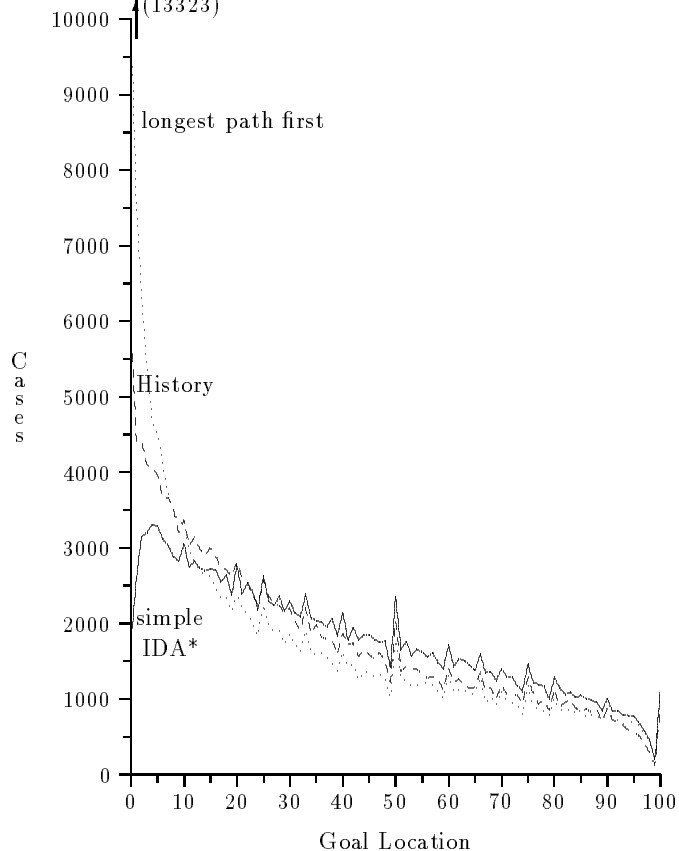


Figure 5: Location of *First* Solution in the Search Front

goal density in the left hand part of the search frontier. The chances to find a goal in the first leaf node are more than three times higher than with a random node selection scheme. Consequently, the right part of the graph lies at a lower level, because all cases again add up to 500 880. Note that the longest-path heuristic affects only the first few leaf expansions, because only one path is stored. However, additional experiments with a refined version that stores *all* longest paths were not successful, because the branching factor is too low in the 8-puzzle. Sophisticated variants only pay off in domains with a large branching factor, like the traveling salesman problem [Reinefeld and Marsland, 1991].

In practice, it often suffices to determine one optimal solution. When the search is stopped after a first solution, the benefit of the longest-path heuristic is even more pronounced, as can be seen in Figure 5. Other favorite aspects of this heuristic are its low space complexity of $O(f^*)$ and its negligible CPU time overhead.

Steepest Ascent Hill-Climbing does not Help. Hill-climbing methods have been found useful in many AI applications. They sort the node successors in increasing order of their heuristic estimates. Successors with low h -values are expanded first, with the expectation that they are closer to a goal. Our empirical results, however, did not reveal any significant advantage of hill-climbing. This observation confirms an earlier assumption of Powley and Korf [1991] and our own experiments

with the larger 15-puzzle. In the 15-puzzle, we spotted a slight advantage of hill climbing in the left part of the tree, but this does not pay for the increased overhead of move pre-sorting. A quick investigation [Reinefeld and Marsland, 1991] reveals that hill-climbing favors tile configurations with the blank located in either an edge or border square, because these configurations enjoy (statistically) lower h -values. Such configurations, however, have a lower ‘mobility’ and are thus less desirable.

History Heuristic. Dynamically acquired knowledge usually outperforms static sorting schemes (such as hill-climbing). One dynamic method, that proved useful in adversary games, is Schaeffer’s [1989] *history heuristic*. It maintains a score table, the *history table*, for every move seen in the search graph. In a given position, all applicable moves are examined in order of their previous success.

The history heuristic does not depend on domain specific knowledge (like heuristic estimate functions). It simply learns from the success in previously expanded subtrees. For the 8-puzzle, one needs a three dimensional array that holds a measure of the goodness of a move for each possible tile, each source position and each move direction. This gives 9 (tiles) \times 9 (positions) \times 4 (move directions) = 324 scores. As a measure of the goodness of a move, we counted the number of occurrences the specific move led to the deepest subtree.

Compared to the longest-path heuristic, the history heuristic provides ordering information for all nodes in the tree, and not only for the nodes on the first path. Even so, our empirical results (see the dashed lines in Figures 4 and 5) indicate, that the history heuristic is dominated by the longest-path heuristic. This might be attributed to the low branching factor of the 8-puzzle, and the inexact methods to measure the goodness of a move. We expect the history heuristic to be more suitable in domains with a larger branching factor and a fine grained evaluation function.

Transposition Table. Transposition tables are used in two-player games [Zobrist, 1970] to avoid unnecessary re-expansions of duplicate nodes (due to move transpositions and due to node re-expansions in successive iterations). When there is enough memory space available, transposition tables are also used to store node information for guiding the search process into the most promising direction.

We implemented a transposition table that holds a representation of every tile configuration seen in the search, the cost-bound to which the configuration has been searched and a best move (the move leading to the deepest subtree). As expected, the transposition table greatly reduces the number of node expansions, see Figure 6. We achieved an average 35% node count reduction per problem. However, most of the savings are due to eliminated transpositions. Node ordering plays only a minor role, and so the solution density graph looks very similar to that of the longest-path heuristic in Figures 4 and 5 (hence we did not plot it). This is again caused by the difficulties to distinguish between ‘good’ and ‘bad’ moves. We took the depth of the emanating subtree as

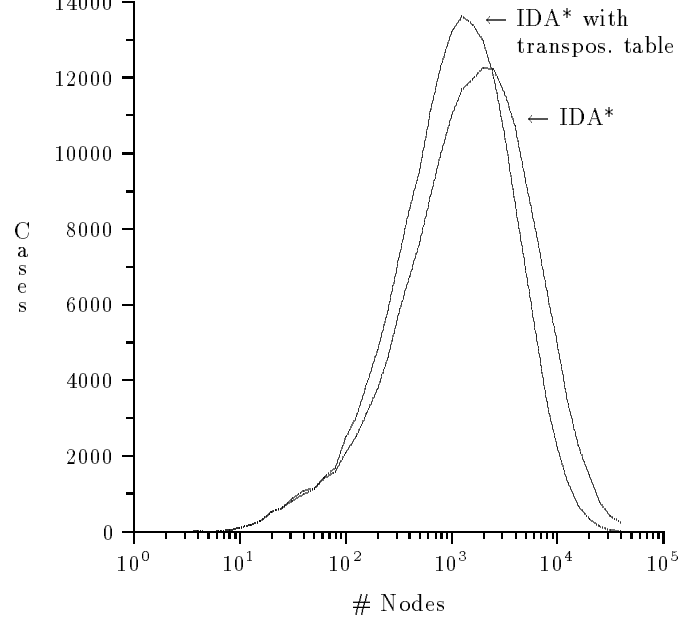


Figure 6: Node expansions to first solution

a utility measure. But many subtrees end at the same level and hence a finer grained scoring function is needed. Another problem is the low branching factor of this specific application, which leaves only few node successors to be sorted.

5 Conclusions

Table 1 gives a summary of our empirical results. The 8-puzzle data has been derived by exhaustively solving all possible board configurations. For the 15-puzzle, no exact data can be given, because the search space is too large to be completely enumerated. We used Korf’s [1985b] selection of 100 randomly generated problem instances as a test set¹.

In our second set of experiments, we used the 8-puzzle as a workbench to evaluate the benefit of node ordering techniques in iterative-deepening search. Surprisingly, we found that common IDA* implementations with a fixed operator sequence (e.g. up, left, right, down) perform worse than average. A simple random operator selection scheme is better!

The longest-path heuristic was found most effective. Consisting of a linear moves array, it is easy to implement and its space overhead of $O(\text{depth})$ is negligible. More sophisticated ordering techniques did not yield better performance, because the 8-puzzle has a low branching factor and it lacks a clear criterion for measuring the goodness of a move.

¹We identified three cases with differing node counts, which are probably due to typographical errors in the original paper [Korf, 1985b]:

<u>No</u>	<u>Korf</u>	<u>Our Version</u>	<u>Difference</u>
22	750,746,755	750,745,755	-1,000
88	6,009,130,748	6,320,047,980	+310,917,232
89	166,571,097	166,571,021	-76

	8-puzzle	15-puzzle																																		
Number of states in the search space	$9!/2 = 181\,440$	$16!/2 \approx 10^{13}$																																		
Edge branching factor b_e	1.67	2.00																																		
Heuristic branching factor b_h	3.81	≈ 6.68																																		
Avg. IDA* node expansions to first solution	3 443	$\approx 363 * 10^6$																																		
Max. IDA* node expansions to first solution	122 417	$\geq 6 * 10^9$																																		
Average initial heuristic distance, h	14.00	≈ 37																																		
Average length of solution	21.97	≈ 53																																		
Maximal length of solution	64	≥ 66																																		
Average number of solutions per problem	2.76	≈ 17																																		
Maximal number of solutions per problem	64	≈ 153																																		
Configurations with longest solution path	<table border="1" style="display: inline-table; margin-right: 20px;"> <tr><td>8</td><td>7</td><td>6</td></tr> <tr><td></td><td>4</td><td>1</td></tr> <tr><td>2</td><td>5</td><td>3</td></tr> </table> $f^* = 31$ <table border="1" style="display: inline-table; margin-right: 20px;"> <tr><td>8</td><td></td><td>6</td></tr> <tr><td>5</td><td>4</td><td>7</td></tr> <tr><td>2</td><td>3</td><td>1</td></tr> </table> $f^* = 31$	8	7	6		4	1	2	5	3	8		6	5	4	7	2	3	1	<table border="1" style="display: inline-table;"> <tr><td>15</td><td>14</td><td></td><td>4</td></tr> <tr><td>11</td><td>1</td><td>6</td><td>13</td></tr> <tr><td>7</td><td>5</td><td>8</td><td>9</td></tr> <tr><td>3</td><td>2</td><td>10</td><td>12</td></tr> </table> $f^* = 66$	15	14		4	11	1	6	13	7	5	8	9	3	2	10	12
8	7	6																																		
	4	1																																		
2	5	3																																		
8		6																																		
5	4	7																																		
2	3	1																																		
15	14		4																																	
11	1	6	13																																	
7	5	8	9																																	
3	2	10	12																																	

Table 1: Summary: The 8-puzzle results are exact; the 15-puzzle data is based on Korf’s random problem set

Acknowledgements

Thanks to Sam Loyd for providing the AI community with such an entertaining research subject. As early as 1878, he surely must have recognized the beauty of simple problem structures that spawn combinatorially large search spaces. Thanks also to Martin Lehmann (Univ. Hamburg) for his continuous interest and many stimulating discussions.

References

- [Cheeseman *et al.*, 1991] P. Cheeseman, B. Kanefsky, W.M. Taylor. *Where the really hard problems are*. 12th IJCAI Conf., Sydney (1991), 331–337.
- [Gaschnig, 1979] J. Gaschnig. *Performance measurement and the analysis of certain search algorithms*. Ph.D. diss., Carnegie Mellon Univ., Pittsburgh (1979).
- [Hart *et al.*, 1968] P.E. Hart, N.J. Nilsson, B. Raphael. *A formal basis for the heuristic determination of minimum cost paths*. IEEE SSC-4,2(1968), 100–107.
- [Johnson and Storey, 1879] W.W. Johnson and W.E. Storey. *Notes on the ‘15’-Puzzle*. Amer. J. Math. 2(1879), 397–404.
- [Korf, 1985a] R.E. Korf. *Macro-operators: A weak method for learning*. Art. Intell. 26(1985), 35–77.
- [Korf, 1985b] R.E. Korf. *Depth-first iterative-deepening: An optimal admissible tree search*. Artificial Intelligence 27(1985), 97–109.
- [Laird *et al.*, 1987] J.E. Laird, A. Newell, P.S. Rosenbloom. *SOAR: An architecture for general intelligence*. Artificial Intelligence 33(1987), 1–64.
- [Mahanti *et al.*, 1992] A. Mahanti, S. Ghosh, D.S. Nau, A.K. Pal and L. Kanal. *Performance of IDA* on trees and graphs*. AAAI-92, San Jose, CA, (1992), 539–544.
- [Nilsson, 1980] N.J. Nilsson. *Principles of Artificial Intelligence*. Tioga Publishing, Palo Alto, CA, (1980).
- [Pearl, 1985] J. Pearl. *Heuristics. Intelligent Search Strategies for Computer Problem Solving*. Addison-Wesley, Reading, MA, (1984).
- [Powley and Korf, 1991] C. Powley and R.E. Korf. *Single-agent parallel window search*. IEEE PAMI-13,5 (1991), 466–477.
- [Rao *et al.*, 1991] V.N. Rao, V. Kumar, R.E. Korf. *Depth-first vs. best-first search*. AAAI-91, Anaheim, CA, (1991), 434–440.
- [Ratner and Warmuth, 1986] D. Ratner and M. Warmuth. *Finding a shortest solution for the $N \times N$ extension of the 15-puzzle is intractable*. AAAI-86, 168–172.
- [Reinefeld and Marsland, 1991] A. Reinefeld and T.A. Marsland. *Memory functions in iterative-deepening search*. Univ. Hamburg, FB Informatik, Tech. Rep. FBI-HH-M-198/91. Submitted for publication.
- [Russell, 1992] S. Russell. *Efficient memory-bounded search methods*. ECAI-92, Vienna, (1992), 1–5.
- [Schaeffer, 1989] J. Schaeffer. *The history heuristic and alpha-beta search enhancements in practice*. IEEE PAMI-11,11(1989), 1203–1212.
- [Schofield, 1967] P.D.A. Schofield. *Complete solution of the ‘Eight-Puzzle’*. N.L. Collins, D. Michie (eds.), Machine Intell. 1, Amer. Elsevier, NY (1967), 125–133.
- [Zobrist, 1970] A.L. Zobrist. *A new hashing method with applications for game playing*. Tech. Rep. 88, Comp. Sc. Dept., Univ. Wisconsin, Madison (Apr. 1970).