# Depth-First Iterative-Deepening: An Optimal Admissible Tree Search* 

Richard E. Korf**<br>Department of Computer Science, Columbia University, New York, NY 10027, U.S.A.


#### Abstract

The complexities of various search algorithms are considered in terms of time. space. and cost of solution path. It is known that breadth-furst search requires 100 much space and depth-first search can use 100 much time and doesn't always find a cheapest path. A depth-first iterative-deepening algorithm is shown to be asymptorically oprimal along all three dimensions for exponential tree searches. The algorithm has been used successfully in chess programs, has been effectively combined with bi-directional search, and has been applied to best-first heuristic search as well. This heuristic depth-first iterative-deepening algorithm is the only known algorithm that is capable of finding oprimal solutions to randomly generated instances of the Fifteen Puzzle within practical resource limits.


## 1. Introduction

Search is ubiquitous in artificial intelligence. The performance of most AI systems is dominated by the complexity of a search algorithm in their inner loops. The standard algorithms, breadth-first and depth-first search, both have serious limitations, which are overcome by an algorithm called depth-first iterative-deepening. Unfortunately, current AI texts either fail to mention this algorithm [ $10,11,14$ ], or refer to it only in the context of two-person game searches $[1,16]$. The iterative-deepening algorithm, however, is completely general and can also be applied to uni-directional search. bi-directional search, and heuristic searches such as $A^{*}$. The purposes of this article are to demonstrate the generality of depth-first iterative-deepening, to prove its optimality for exponential tree searches, and to remind practitioners in the field that it is the search technique of choice for many applications.

Depth-first iterative-deepening has no doubt been rediscovered many times

[^0]* Present address: Department of Computer Science, University of California, Los Angeles, CA 90024, U.S.A.
independently. The first use of the algorithm that is documented in the literature is in Slate and Atkin's Chess 4.5 program [15]. Berliner [2] has observed that breadth-first search is inferior to the iterative-deepening algorithm. Winston [16] shows that for two-person game searches where only terminal-node static evaluations are counted in the cost, the extra computation required by iterative-deepening is insignificant. Pearl [12] initially suggested the iterative-deepening extension of $\mathrm{A}^{*}$, and Berliner and Goetsch [3] have implemented such an algorithm concurrently with this work.

We will analyze several search algorithms along three dimensions: the amount of time they take, the amount of space they use, and the cost of the solution paths they find. The standard breadth-first and depth-first algorithms will be shown to be inferior to the depth-first iterative-deepening algorithm. We will prove that this algorithm is asymptotically optimal along all three dimensions for exponential tree searches. Since almost all heuristic tree searches have exponential complexity, this is a fairly general result.

We begin with the problem-space model of Newell and Simon [9]. A problem space consists of a set of states and a set of operators that are partial functions that map states into states. A problem is a problem space together with a particular initial state and a set of goal states. The task is to find a sequence of operators that will map the initial state to a goal state.

The complexity of a problem will be expressed in terms of two parameters: the branching factor of the problem space, and the depth of solution of the problem. The node branching factor ( $b$ ) of a problem is defined as the number of new states that are generated by the application of a single operator to a given state, averaged over all states in the problem space. We will assume that the branching factor is constant throughout the problem space. The depth (d) of solution of a problem is the length of the shortest sequence of operators that map the initial state into a goal state. The time cost of a search algorithm in this model of computation is simply the number of states that are expanded. The reason for this choice is that we are interested in asymptotic complexity and we assume that the amount of time is proportional to the number of states expanded. Similarly, since we assume that the amount of space required is proportional to the number of states that are stored, the asymptotic space cost of an algorithm in this model will be the number of states that must be stored.

This work is focused on searches which produce optimal solutions. We recognize that for most applications, optimal solutions are not required and that their price is often prohibitive. There are occasions, however, when optimal solutions are needed. For example, in assessing the quality of nonoptimal solutions, it is often enlightening to compare them to optimal solutions for the same problem instances.

## 2. Breadth-First Search

We begin our discussion with one of the simplest search algorithms, breadth-
first search. Breadth-first search expands all the states one step (or operator application) away from the initial state, then expands all states two steps from the initial state, then three steps, etc.. until a goal state is reached. Since it always expands all nodes at a given depth before expanding any nodes at a greater depth, the first solution path found by breadth-first search will be one of shortest length. In the worst case, breadth-first search must generate all nodes up to depth $d$, or $b+b^{2}+b^{3}+\cdots+b^{d}$ which is $\mathrm{O}\left(b^{d}\right)$. Note that on the average, half of the nodes at depth $d$ must be examined. and therefore the average-case time complexity is also $\mathrm{O}\left(b^{d}\right)$.

Since all the nodes at a given depth are stored in order to generate the nodes at the next depth, the minimum number of nodes that must be stored to search to depth $d$ is $b^{d-1}$, which is $\mathrm{O}\left(b^{d}\right)$. As with time, the average-case space complexity is roughly one-half of this, which is also $\mathrm{O}\left(b^{d}\right)$. This space requirement of breadth-first search is its most critical drawback. As a practical matter, a breadth-first search of most problem spaces will exhaust the available memory long before an appreciable amount of time is used. The reason for this is that the typical ratio of memory to speed in modern computers is a million words of memory for each million instructions per second (MIPS) of processor speed. For example, if we can generate a million states per minute and require a word to store each state, memory will be exhausted in one minute.

## 3. Depth-First Search

Depth-first search avoids this memory limitation. It works by always generating a descendant of the most recently expanded node, until some depth cutoff is reached, and then backtracking to the next most recently expanded node and generating one of its descendants. Therefore, only the path of nodes from the initial node to the current node must be stored in order to execute the algorithm. If the depth cutoff is $d$, the space required by depth-first search is only $\mathrm{O}(d)$.

Since depth-first search only stores the current path at any given point, it is bound to search all paths down to the cutoff depth. In order to analyze its time complexity, we must define a new parameter, called the edge branching factor (e), which is the average number of different operators which are applicable to a given state. For trees. the edge and node branching factors are equal, but for graphs in general the edge branching factor may exceed the node branching factor. For example, the graph in Fig. 1 has an edge branching factor of two, while its node branching factor is only one. Note that a breadth-first search of this graph takes only linear time while a depth-first search requires exponential time. In general, the time complexity of a depth-first search to depth $d$ is $O\left(e^{d}\right)$. Since the space used by depth-first search grows only as the log of the time required, the algorithm is time-bound rather than space-bound in practice.

Another drawback, however, to depth-first search is the requirement for an arbitrary cutoff depth. If branches are not cut off and duplicates are not


FIG. 1. Graph with linear number of nodes but exponential number of paths.
checked for, the algorithm may not terminate. In general, the depth at which the first goal state appears is not known in advance and must be estimated. If the estimate is too low, the algorithm terminates without finding a solution. If the depth estimate is too high, then a large price in running time is paid relative to an optimal search, and the first solution found may not be an optimal one.

## 4. Depth-First Iterative-Deepening

A search algorithm which suffers neither the drawbacks of breadth-first nor depth-first search on trees is depth-first iterative-deepening (DFID). The algorithm works as follows: First, perform a depth-first search to depth one. Then, discarding the nodes generated in the first search, start over and do a depth-first search to level two. Next, start over again and do a depth-first search to depth three, etc., continuing this process until a goal state is reached.

Since DFID expands all nodes at a given depth before expanding any nodes at a greater depth, it is guaranteed to find a shortest-length solution. Also, since at any given time it is performing a depth-first search, and never searches deeper than depth $d$, the space it uses is $\mathrm{O}(d)$.

The disadvantage of DFID is that it performs wasted computation prior to reaching the goal depth. In fact, at first glance it seems very inefficient. Below, however, we present an analysis of the running time of DFID that shows that this wasted computation does not affect the asymptotic growth of the run time for exponential tree searches. The intuitive reason is that almost all the work is done at the deepest level of the search. Unfortunately, DFID suffers the same drawback as depth-first search on arbitrary graphs, namely that it must explore all possible paths to a given depth.

Definition 4.1. A brute-force search is a search algorithm that uses no information other than the initial state, the operators of the space, and a test for a solution.

Theorem 4.2. Depth-first iterative-deepening is asymptotically optimal among brute-force tree searches in terms of time, space, and length of solurion.

Proof. As mentioned above, since DFID generates all nodes at a given depth before expanding any nodes at a greater depth. it always finds a shortest path to the goal, or any other state for that matter. Hence, it is optimal in terms of solution length.

Next, we examine the running time of DFID on a tree. The nodes at depth $d$ are generated once during the final iteration of the search. The nodes at depth $d-1$ are generated twice, once during the final iteration at depth $d$, and once during the penultimate iteration at depth $d-1$. Similarly, the nodes at depth $d-2$ are generated three times, during iterations at depths $d, d-1$, and $d-2$. etc. Thus the total number of nodes generated in a depth-first iterativedeepening search to depth $d$ is

$$
b^{d}+2 b^{d-1}+3 b^{d-2}+\cdots+d b
$$

Factoring out $b^{d}$ gives

$$
b^{d}\left(1+2 b^{-1}+3 b^{-2}+\cdots+d b^{1-d}\right)
$$

Letting $x=1 / b$ yields

$$
b^{d}\left(1+2 x^{1}+3 x^{2}+\cdots+d x^{d-1}\right)
$$

This is less than the infinite series

$$
b^{d}\left(1+2 x^{1}+3 x^{2}+4 x^{3}+\cdots\right)
$$

which converges to

$$
b^{d}(1-x)^{-2} \text { for } \operatorname{abs}(x)<1
$$

Since $(1-x)^{-2}$, or $(1-1 / b)^{-2}$, is a constant that is independent of $d$, if $b>1$ then the running time of depth-first iterative-deepening is $O\left(b^{d}\right)$.

To see that this is optimal, we present a simple adversary argument. The number of nodes at depth $d$ is $b^{d}$. Assume that there-exists an algorithm that examines less than $b^{d}$ nodes. Then, there must exist at least one node at depth $d$ which is not examined by this algorithm. Since we have no additional information, an adversary could place the only solution at this node and hence the proposed algorithm would fail. Hence, any brute-force algorithm must take at least $c b^{d}$ time, for some constant $c$.

Finally, we consider the space used by DFID. Since DFID at any point is engaged in a depth-first search, it need only store a stack of nodes which represents the branch of the tree it is expanding. Since it finds a solution of optimal length, the maximum depth of this stack is $d$, and hence the maximum amount of space is $\mathrm{O}(d)$.


Fic. 2. Graph of branching factor vs. constant coefficient as search depth goes to infinity.
To show that this is optimal, we note that any algorithm which uses $f(n)$ time must use at least $k \log f(n)$ space for some constant $k$ [7]. The reason is that the algorithm must proceed through $f(n)$ distinct states before looping or terminating, and hence must be able to store that many distinct states. Since storing $f(n)$ states requires $\log f(n)$ bits, and $\log b^{d}$ is $d \log b$, any brute-force algorithm must use $k d$ space, for some constant $k$.

The value of the constant $(1-1 / b)^{-2}$ gives an upper bound on how much computation is wasted in the lower levels of the search, since it is the limit of the constant coefficient as the search depth goes to infinity. Fig. 2 shows a graph of this constant versus the branching factor. As the branching factor increases, the constant quickly approaches one. For branching factors close to one, however, the value of the constant coefficient approaches infinity as the depth goes to infinity.

## 5. Bi-Directional Search

For those problems with a single goal state that is given explicitly and for which the operators have inverses, such as the Fifteen Puzzle, bi-directional search
[13] can be used. Bi-directional search trades space for time by searching forward from the initial state and backward from the goal state simultaneously, storing the states generated, until a common state is found on both search frontiers. Depth-first iterative-deepening can be applied to bi-directional search as follows: A single iteration consists of a depth-first search from one direction to depth $k$, storing only the states at depth $k$, and two depth-first searches from the other direction, one to depth $k$ and one to depth $k+1$, not storing states but simply matching against the stored states from the other direction. The search to depth $k+1$ is necessary to find odd-length solutions. This is repeated for $k$ from zero (to find solutions of length one) to $d / 2$. Assuming that a hashing scheme is used to perform the matching in constant time per node, this algorithm will find an optimal solution of length $d$ in time $O\left(b^{d / 2}\right)$ and space $O\left(b^{d / 2}\right)$. In experiments involving Rubik's Cube [8], which has an effective branching factor of 13.5, this algorithm was used to find solutions up to 11 moves long on a DEC VAX 11/780.

## 6. Heuristic Search

Depth-first iterative-deepening can also be combined with a best-first heuristic search such as $A^{*}[6]$. The idea is that successive iterations correspond not to increasing depth of search, but rather to increasing values of the total cost of a path. For $A^{*}$, this total cost is composed of the cost so far in reaching the node ( $g$ ) plus the estimated cost of the path from the node to a goal state ( $h$ ). Iterative-deepening- $A^{*}$ (IDA*) works as follows: At each iteration, perform a depth-first search, cutting off a branch when its total cost $(g+h)$ exceeds a given threshold. This threshold starts at the estimate of the cost of the initial state, and increases for each iteration of the algorithm. At each iteration, the threshold used for the next iteration is the minimum cost of all values that exceeded the current threshold.

A well-known property of $A^{*}$ is that it always finds a cheapest solution path if the heuristic is admissible, or in other words never overestimates the actual cost to the goal [6]. This property also holds for iterative-deepening-A*. Furthermore, IDA* expands the same number of nodes, asymptotically, as $\mathrm{A}^{*}$ in an exponential tree search.

The proofs of these results are much simpler and more intuitive if we restrict our attention to cost functions which are monotonically non-decreasing along any path in the problem space. Such a heuristic is called monotone or consistent [11]. Formally,

Definition 6.1. A cost function $f(n)$ is monotone if for all nodes $n$ and $s(n)$, where $s(n)$ is a successor of $n, f(n) \leqslant f(s(n)$ ).

This restriction is not essential, and slightly more complex proofs will establish the same results without it. As a practical matter, however, almost all
reasonable cost functions are monotone [11]. In fact. using an idea proposed by Mérö [17], we can formally make this assumption without loss of generality, as shown in the following lemma.

Lemma 6.2. For any admissible cost function f, we can construct a monotone admissible function $f^{\prime}$ which is at least as informed as $f$.

Proof. We construct $f^{\prime}$ recursively from $f$ as follows: if $n$ is the initial state, then $f^{\prime}(n)=f(n)$; otherwise, $f^{\prime}(s(n))=\max \left[f(s(n)), f^{\prime}(n)\right]$. Clearly, $f^{\prime}$ is monotone since $f^{\prime}(n) \leqslant f^{\prime}(s(n))$. In order to show that $f^{\prime}$ is admissible, note that $f^{\prime}(n)$ is equal to the maximum value of $f$ applied to all the predecessors of $n$ along the path back to the initial state. Since $f$ is admissible, the maximum value of $f$ along a path is a lower bound on the cost of that path, and hence a lower bound on the cost of $n$. Thus, $f^{\prime}$ does not violate admissibility. Furthermore, $f^{\prime}$ is at least as informed as $f$ since for all $n, f^{\prime}(n) \geqslant f(n)$ and hence $f^{\prime}(n)$ is at least as accurate an estimate as $f(n)$.

Note that this lemma provides a simple and intuitive proof of the admissibility of $A^{*}$. If we restrict our attention to cost functions which are monotone non-decreasing, and $A^{*}$ always expands the open node of least cost, it is clear that the first solution it finds will be one of least cost. Similarly, the result below follows just as easily.

Lemma 6.3. Given an admissible monotone cost function, iterative-deepening$A^{*}$ will find a solution of least cost if one exists.

Proof. Since the initial cost cutoff of IDA* is the heuristic estimate of the cost of the initial state, and the heuristic never overestimates cost, the length of the shortest solution cannot be less than the initial cost cutoff. Furthermore, since the cost cutoff for each succeeding iteration is the minimum value which exceeded the previous cutoff, no paths can have a cost which lies in a gap between two successive cutoffs. Therefore, since IDA* always expands all nodes at a given cost before expanding any nodes at a greater cost, the first solution it finds will be a solution of least cost.

Not only does IDA* find a cheapest path to a solution and use far less space than $A^{*}$, but it expands approximately the same number of nodes as $A^{*}$ in a tree search. Combining this fact with several recent results on the complexity and optimality of $A^{*}$ allows us to state and prove the following general result:

Theorem 6.4. Given an admissible monotone heuristic with constant relative error, then iterative-deepening. $A^{*}$ is optimal in terms of solution cost. time, and space, over the class of admissible best-first searches on a tree.

Proof. From Lemma 6.3, we know that IDA* produces a solution of optimal cost.

To determine the time used by IDA*, consider the final iteration, in other words the one which finds a solution. It must expand all descendents of the initial state with values greater than or equal to the initial cost estimate and less than the optimal solution cost, plus some number of nodes whose cost equals the optimal solution cost. If $A^{*}$ employs the tie-breaking rule of 'most recently generated', it must also expand these same nodes. Thus, the final iteration of IDA* expands the same set of nodes as $A^{*}$ under this tie-breaking rule. Furthermore, if the graph is a tree, each of these nodes will be expanded exactly once. IDA* must also expand nodes during the previous iterations as well. However, Pearl has shown that if the heuristic used by $A^{*}$ exhibits constant relative error, then the number of nodes generated by the algorithm increases exponentially with depth [11]. Thus, we can use an argument similar to the proof of Theorem 4.2 to show that the previous iterations of IDA* do not affect the asymptotic order of the total number of nodes [18]. Thus, IDA * expands the same number of nodes, asymptotically, as $A^{*}$. Furthermore, a recent result of Dechter and Pearl [5] shows that $A^{*}$ is optimal, in terms of number of nodes expanded, over the class of admissible best-first searches with monotone heuristics. Therefore, IDA* is asymptotically optimal in terms of time for tree searches.

Since the number of nodes grows exponentially, we can again appeal to the argument in the proof of Theorem 4.2 to show that the space used by IDA* is also asymptotically optimal.

Is the assumption of constant relative error, i.e. that the error in the estimate grows at the same rate as the magnitude of the actual cost, valid for heuristics? Pearl observes that heuristics with better accuracy almost never occur in practice. For example, most physical measurements are subject to constant relative error [11]. Thus, we can conclude that heuristic depth-first iterativedeepening is asymptotically optimal for most best-first tree searches which occur in practice.

An additional benefit of IDA* over $A^{*}$ is that it is simpler to implement since there are no open or closed lists to be managed. A simple recursion performs the depth-first search inside an outer loop to handle the iterations.

As an empirical test of the practicality of the algorithm, both IDA* and A* were implemented for the Fifteen Puzzle. The implementations were in PASCAL and were run on a DEC 2060. The heuristic function used for both was the Manhattan distance heuristic: for each movable tile, the number of grid units between the current position of the tile and its goal position are computed, and these values are summed for all tiles. The two algorithms were tested against 100 randomly generated, solvable initial states. IDA* solved all instances with a median time of 30 CPU minutes, generating over 1.5 million nodes per minute. The average solution length was 53 moves and the maximum was 66 moves. $A^{*}$
solved none of the instances since it ran out of space after about 30000 nodes were stored. An additional observation is that even though IDA* generated more nodes than $A^{*}$, it actually ran faster than $A^{*}$ on the same problem instances, due to less overhead per node. The data from this experiment are summarized in Table 1. These are the first published optimal solution lengths to randomly generated instances of the Fifteen Puzzle. Although the Fifteen Puzzle graph is not strictly a tree, the edge branching factor is only slightly greater than the node branching factor, and hence the iterative-deepening algorithm is still effective.

Table 1. Optimal solution lengths for 100 randomly generated Fifteen Puzzle instances using iterative-deepening-A* with Manhattan distance heuristic function

| NUM | 2 INITIAL STATE | ESTIMATE | ACTUAL | TOTAL NODES |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 1 |  | 41 | 57 | 276.361.933 |  |
| 2 |  | 43 | 55 | 15.300.442 |  |
| 3 |  | 41 | 59 | 565.994.203 |  |
| 4 |  | 42 | 56 | 62,643.179 |  |
| 5 |  | 42 | 56 | 11.020,325 |  |
| 6 |  | 36 | 52 | 32.201.660 |  |
| 7 | 2111555134667128101993140 | 30 | 52 | 387,138,094 |  |
| 8 |  | 32 | 50 | 39.118.937 |  |
| 9 |  | 32 | 46 | 1.650 .696 |  |
| 10 | 13118890015710443614551221 | 43 | 59 | 198.758.703 |  |
| 11 |  | 43 | 57 | 150.346.072 |  |
| 12 |  | 35 | 45 | 546,344 |  |
| 13 |  | 36 | 46 | 11.861.705 |  |
| 14 |  | 41 | 59 | 1.369 .596 .778 |  |
| 15 |  | 44 | 62 | 543.598,067 |  |
| 16 |  | 24 | 42 | 17.984.051 |  |
| 17 |  | 46 | 66 | 607.399.560 |  |
| 18 |  | 43 | 55 | 23.711 .067 |  |
| 19 |  | 36 | 46 | 1.280.495 |  |
| 20 |  | 36 | 52 | 17.954.870 |  |
| 21 |  | 34 | 54 | 257.064.810 |  |
| 22 |  | 41 | 59 | 750.746,755 |  |
| 23 |  | 33 | 49 | 15.971.319 |  |
| 24 |  | 34 | 54 | 42,693.209 |  |
| 25 |  | 32 | 52 | 100.734.844 |  |
| 26 |  | 40 | 58 | 226.668.645 |  |
| 27 | 141981512600391210134471511 | 33 | 53 | 306.123.421 |  |
| 28 |  | 36 | 52 | 5.934.442 |  |
| 29 |  | 38 | 54 | 117.076.111 |  |
| 30 |  | 35 | 47 | 2.196.593 |  |
| 31 |  | 38 | 50 | 2.351.811 |  |
| 32 |  | 43 | 59 | 661.041 .936 |  |
| 33 |  | 42 | 60 | 480,637.867 |  |
| 34 |  | 36 | 52 | 20.671.552 |  |
| 35 |  | 39 | 55 | 47.506.056 |  |
| 36 |  | 36 | 52 | 59,802.602 |  |
| 37 |  | 40 | 58 | 280,078,791 |  |
| 38 |  | 41 | 53 | 24.492.852 |  |
| 39 | 9041011415312657111388 | 35 | 49 | 19.355.806 |  |
| 40 |  | 36 | 54 | 63.276 .188 |  |
| 41 |  | 36 | 54 | 51.501.544 |  |
| 42 |  | 30 | 42 | 877.823 |  |
| 43 |  | 48 | 64 | 41.124,767 |  |
| 44 |  | 32 | 50 | 95.733 .125 |  |
| 45 |  | 39 | 51 | 6.158.733 |  |
| 46 |  | 35 | 49 | 22,119.320 |  |
| 47 |  | 35 | 47 | 1.411.294 |  |
| 48 |  | 39 | 49 | 1.905.023 |  |
| 49 |  | 33 | 59 | 1.809 .933 .698 |  |
| 50 |  | 39 | 53 | 63.036 .422 |  |
| 51 | 102844150124111336971512 | 44 | 56 | 26.622.863 |  |

Table 1. Continued

$\begin{array}{lllllllllllll}0 & 1 & 2 & 3 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 & 13 \\ 1\end{array} 415$

## 7. Two-Person Games

In the discussion so far, we have assumed a single-agent search to find a solution to a problem, and have been concerned with minimizing time and space subject to a fixed solution depth and branching factor. However. a
two-person game such as chess with static evaluation and mini-max search is a somewhat different situation. In this case, we assume that accuracy of the static evaluation increases with increasing search depth, and hence we want to maximize search depth subject to fixed time and space constraints. Since depth-first iterative-deepening minimizes, at least asymptotically, time and space for any given search depth, it follows that it maximizes the depth of search possible for any fixed time and space restrictions as well.

Another reason that DFID is used in game programs is that the amount of time required to search the next deeper level in the tree is not known when the ply begins, and the search ply may have to be aborted due to time constraints. In this case, the complete search at the next shallower depth can be used to make the move.

Finally, the information from previous iterations of a DFID search can be used to order the nodes in the search tree so that alpha-beta cutoff is more efficient. In fact, the best move at a given iteration has been shown experimentally to terminate the next iteration in about $70 \%$ of cases. This improvement in ordering, which is critical to alpha-beta efficiency, is only possible with the use of iterative-deepening [4].

## 8. Conclusions

The standard algorithms for brute-force search have serious drawbacks. Breadth-first search uses too much space, and depth-first search in general uses too much time and is not guaranteed to find a shortest path to a solution. The depth-first iterative-deepening algorithm, however, is asymptotically optimal in terms of cost of solution, running time, and space required for brute-force tree searches. DFID can also be applied to bi-directional search, heuristic best-first search, and two-person game searches. Since almost all heuristic searches have exponential complexity, iterative-deepening- $A^{*}$ is an optimal admissible tree search in practice. For example, IDA* is the only known algorithm that can find optimal paths for randomly generated instances of the Fifteen Puzzle within practical time and space constraints.

## ACKNOWLEDGEMENT

Judea Pearl originally suggested the application of iterative-deepening to $A^{*}$. Hans Berliner pointed out the use of iterative-deepening for ordering nodes to maximize alpha-beta cutoffs. Michael Lebowitz. Andy Mayer, and Mike Townsend read earlier drafts of this paper and suggested many improvements. Andy Mayer implemented the $A^{*}$ algorithm that was compared with IDA*. An anonymous referee suggested the shortcomings of depth-first search on a graph with cycles. Finally, Jodith Fried drew the figures.

## REFERENCES

1. Barr, A. and Feigenbaum, E.A. (Eds.), Handbook of Artificial Intelligence (Kaufmann, Los Altos. CA, 1981).
2. Berliner, H., Search, Artificial Intelligence Syllabus, Department of Computer Science, CarnegieMellon University, Pittsburgh, PA, 1983.
3. Berliner, H. and Goetsch. G.. A quantitative study of search methods and the effect of constraint satisfaction, Tech. Rept. CMU-CS-84-147. Department of Computer Science, Carnegie-Mellon University, Pittsburgh, PA, 1984.
4. Berliner. H., Personal communication, 1984.
5. Dechter, R. and Pearl, J., The optimality of $A^{*}$ revisited, in: Proceedings of the National Conference on Artificial Intelligence, Washington, DC (August, 1983) 95-99.
6. Hart, P.E., Nilsson, N.J. and Raphael, B., A formal basis for the heuristic determination of minimum cost paths, IEEE Trans. Systems Sci. Cybernet. 4(2) (1968) 100-107.
7. Hopcroft, J.E. and Ullman, J.D., Introducrion to Automata Theory, Languages, and Computation (Addison-Wesley, Reading, MA, 1979).
8. Korf, R.E., Leaming to Solve Problems by Searching for Macro-Operators (Pittman, London, 1985).
9. Newell, A. and Simon, H.A., Human Problem Solving (Prentice-Hall, Englewood Cliffs, NJ, 1972).
10. Nilsson, N.J., Principles of Artificial Intelligence (Tioga, Palo Alto, CA, 1980).
11. Pearl, J., Heuristics (Addison-Wesley, Reading, MA, 1984).
12. Pearl, J., Personal communication, 1984.
13. Pohl, I., Bi-directional search, in: B. Meltzer and D. Michie (Eds.), Machine Intelligence 6 (American Elsevier, New York, 1971) 127-140.
14. Rich, E., Artificial Intelligence (McGraw-Hill, New York, 1983).
15. Slate, D.J. and Atkin, L.R., CHESS 4.5 - The Northwestern University Chess Program (SpringerVerlag, New York, 1977).
16. Winston, P.H., Artificial Intelligence (Addison'Wesley, Reading, MA, 1984).
17. Mêró, L., A heuristic search algorithm with modifiable estimate, Artificial Intelligence 23 (1984) 13-27.
18. Korf, R.E., Iterative-deepening- $A^{*}$ : an optimal admissible tree search, in: Proceedings Ninth International Joint Conference on Arrificial Intelligence, Los Angeles, CA, 1985.

Received December 1984; revised version received March 1985


[^0]:    - This research was supported in part by the Defense Advanced Research Projects Agency under contract N00039-82-C-0427, and by the National Science Foundation Division of Information Science and Technology grant IST-84-18879.

