

SetA*: An Efficient BDD-Based Heuristic Search Algorithm

Rune M. Jensen, Randal E. Bryant, and Manuela M. Veloso

Computer Science Department, Carnegie Mellon University,
5000 Forbes Ave, Pittsburgh, PA 15213-3891, USA
{runej,bryant,mmv}@cs.cmu.edu

Abstract

In this paper we combine the goal directed search of A* with the ability of BDDs to traverse an exponential number of states in polynomial time. We introduce a new algorithm, SetA*, that generalizes A* to expand sets of states in each iteration. SetA* has substantial advantages over BDDA*, the only previous BDD-based A* implementation we are aware of. Our experimental evaluation proves SetA* to be a powerful search paradigm. For some of the studied problems it outperforms BDDA*, A*, and BDD-based breadth-first search by several orders of magnitude. We believe exploring sets of states to be essential when the heuristic function is weak. For problems with strong heuristics, SetA* efficiently specializes to single-state search and consequently challenges single-state heuristic search in general.

Introduction

During the last decade, powerful search techniques using an implicit state representation based on the *reduced ordered binary decision diagram* (BDD, Bryant 1986) have been developed in the area of *symbolic model checking* (McMillan 1993). Using blind exploration strategies these techniques have been successfully applied to verify systems with very large state spaces. Similar results have been obtained in well-structured AI search domains (Cimatti *et al.* 1997). However for hard combinatorial problems the search fringe often grows exponentially with the search depth.

A classical AI approach for avoiding the state explosion problem is to use heuristics to guide the search toward the goal states. The question is whether heuristics can be applied to BDD-based search such that their ability to efficiently expand a large set of states in each iteration is preserved. The answer is non-trivial since heuristic search algorithms require values to be associated with each state and manipulated during search a task for which BDDs often have proven less efficient.

In this paper, we present a new search algorithm called SetA*. The main idea is to avoid the above problem by generalizing A* (Hart, Nilsson, & Raphael 1968) from single states to sets of states in the search queue. Recall that A*

associates two values g and h with each state in the search queue. g is the cost of reaching the state and h is an estimate of the remaining cost of reaching the goal given by a *heuristic function*. In SetA* states with similar g and h values are merged such that we can represent them implicitly by a BDD without having to store any numerical information. In each iteration, SetA*: 1) pops the set with highest priority, 2) computes its next states, and 3) partitions the next states into child sets with unique g and h values, which are reinserted into the queue. A straightforward implementation of the three phases has disappointing performance (see PreSetA*, Table 2). A key idea of our work is therefore to combine phase 2 and 3. The technique fits nicely with the so called disjunctive partitioning of BDD-based search (Clarke, Grumberg, & Peled 1999). In addition it can be applied to any heuristic function. Our experimental evaluation of SetA* proves it a powerful search paradigm. For some problems it dominates both A* and BDD-based breadth-first search (see Table 2). In addition, it outperforms the only previous BDD-based implementation of A* (Edelkamp & Reffel 1998), we are aware of, by up to two orders of magnitude (see Table 3).

Directed BDD-based search has received little attention in symbolic model checking. The reason is that the main application of BDDs in this field is verification where all reachable states must be explored. For Computation Tree Logic (CTL) checking, guiding techniques have been proposed to avoid a blow-up of intermediate BDDs (Bloem, Ravi, & Somenzi 2000). However these techniques are not applicable to search since they are based on defining lower and upper bounds on the fixed-point. Directed search techniques are relevant for *falsification* where the goal is to find a state not satisfying an invariant. As far as we know, BDD-based directed search was first considered for this application (Yang & Dill 1998). The proposed algorithm is a simple best-first search where the search fringe is partitioned with a specialized BDD-operator according to the *Hamming distance* to the goal state (number of different bits). Even though this operation is fairly efficient for the Hamming distance, it is not obvious how to define it in general. The only previous BDD-based implementation of A* that we are aware of, is BDDA* (Edelkamp & Reffel 1998). BDDA* can use a general heuristic function and has been applied to planning as well as model checking. However, BDDA* re-




Figure 1: An example search problem consisting of four states and four transitions a, b, c , and d . The dashed lines indicate the two search fringes of a BDD-based breadth-first search from the initial state $i = (0,0)$ to the goal states $G = \{(1,1)\}$. The h -values is a heuristic function equal to the vertical goal distance.

quires arithmetic operations at the BDD level during search and includes no tools to control the growth of the search fringe or for cycle detection. In practice this often leads to substantial performance penalties (see Table 3). Another limitation of BDDA* is that the generalization to weighted A* is non-trivial.

The remainder of the paper is organized as follows. First we briefly describe BDDs and BDD-based search techniques. We then define the SetA* algorithm and evaluate it experimentally in a range of search and planning domains. Finally we draw conclusions and discuss directions for future work.

BDD-based Search

A BDD is a directed acyclic graph representing a Boolean function on a set of ordered variables. Due to two reduction rules the BDD is a canonical and compact representation. Another advantage is that a large set of BDDs can share structure in a multi-rooted BDD and be efficiently manipulated by a general function for applying Boolean operators. Due to the space limitations of this paper we will treat BDDs as a black-box. Readers interested in a thorough introduction are referred to (Bryant 1986).

A search problem is a 4-tuple (S, T, i, G) . S is a set of states. $T : S \times S$ is a *transition relation* defining the search graph. $(s, s') \in T$ iff there exists a transition leading from s to s' . i is the initial state of the search while G is the set of goal states. A solution to a search problem is a path $\pi = s_0, \dots, s_n$ where $s_0 = i$ and $s_n \in G$ and $\bigwedge_{j=0}^{n-1} (s_j, s_{j+1}) \in T$. Assuming that states can be encoded as bit vectors, BDDs can be used to represent the *characteristic function* of a set of states and the transition relation. To make this clear, consider the simple search problem shown in Figure 1. A state s is represented by a bit vector with two elements $\vec{s} = (s_0, s_1)$. Thus the initial state is represented by a BDD for the expression $\neg s_0 \wedge \neg s_1$. Similarly we have $G = s_0 \wedge s_1$. To encode the transition relation, we need to refer to current state variables and next state variables. We adopt the usual notation in BDD literature of primed vari-

ables for the next state

$$\begin{aligned} T(s_0, s_1, s'_0, s'_1) &= \neg s_0 \wedge \neg s_1 \wedge s'_0 \wedge \neg s'_1 \\ &\vee \neg s_0 \wedge \neg s_1 \wedge \neg s'_0 \wedge s'_1 \\ &\vee \neg s_0 \wedge s_1 \wedge s'_0 \wedge s'_1 \\ &\vee s_0 \wedge s_1 \wedge s'_0 \wedge \neg s'_1. \end{aligned}$$

The main idea in BDD-based search is to stay at the BDD level when finding the next states of a set of states. This can be done by computing the image of a set of states V encoded in current state variables

$$\text{Img} = (\exists \vec{s}. V(\vec{s}) \wedge T(\vec{s}, \vec{s}'))[\vec{s}/\vec{s}'].$$

Consider the first step of the search from i in the example domain. We have $V(s_0, s_1) = \neg s_0 \wedge \neg s_1$. Thus,

$$\begin{aligned} \text{Img} &= (\exists \vec{s}. \neg s_0 \wedge \neg s_1 \wedge T(s_0, s_1, s'_0, s'_1))[\vec{s}/\vec{s}'] \\ &= (s'_0 \wedge \neg s'_1 \vee \neg s'_0 \wedge s'_1)[\vec{s}/\vec{s}'] \\ &= s_0 \wedge \neg s_1 \vee \neg s_0 \wedge s_1. \end{aligned}$$

A common problem in BDD-based search is that intermediate BDDs in the image computation tend to be large compared to the BDD representing the result. In symbolic model checking, a range of techniques has been proposed to avoid this problem. Among the most successful of these are *transition relation partitioning*. For search problems, where each transition normally only modifies a small subset of the state variables, the suitable partitioning technique is *disjunctive partitioning* (Clarke, Grumberg, & Peled 1999). To make a disjunctive partitioning, the part of the individual transition expressions keeping the unmodified variables unchanged is removed. The transition expressions are then partitioned according to what variables they modify. For our example we get two partitions

$$\begin{aligned} P_1 &= \neg s_0 \wedge \neg s_1 \wedge s'_0 \vee \neg s_0 \wedge s_1 \wedge s'_0 \\ m_1 &= (s_0) \\ P_2 &= \neg s_0 \wedge \neg s_1 \wedge s'_1 \vee s_0 \wedge s_1 \wedge \neg s'_1 \\ m_2 &= (s_1). \end{aligned}$$

In addition to large space savings, disjunctive partitioning often lowers the complexity of the image computation which now can skip the quantifications on unchanged variables and operate on smaller expressions

$$\text{Img} = \bigvee_{j=1}^{|\mathbf{P}|} (\exists m_j. V(\vec{s}) \wedge P_j(\vec{s}, m'_j))[m_j/m'_j].$$

SetA*

SetA* is a generalization of weighted A* where the definition of f is changed from $f = g + h$ to $f = (1 - w)g + wh$, $w \in [0, 1]$ (Pohl 1970). Similar to BDDA*, SetA* assumes a finite search domain and unit-cost transitions. SetA* expands a set of states instead of just a single state. The main input is what we will call, an *improvement partitioning*. That is, a disjunctive partitioning where the transitions of a partition reduce the h -value by the same amount. The improvement partitioning is non-trivial to compute. The reason is that it may be intractable to calculate

```

function SetA*(IP, init, goal, u, w)
1  Q.initialize(u, w, goal)
2  g ← 0
3  h ← h(init)
4  Q.insert(init, g, h)
5  R.update(init, g)
6  while ¬Q.empty() and ¬Q.topAtGoal()
7    top ← Q.pop()
8    for j = 0 to |IP|
9      next ← image(top, IPj)
10     R.prune(next)
11     g ← top.g + 1
12     h ← top.h − impr(IPj)
13     Q.insert(next, g, h)
14     R.update(next, g)
15 if Q.empty() then NoPathExists
16 else R.extractPath()

```

Figure 2: The SetA* algorithm.

each transition expression in turn. Fortunately large sets of transitions are often described in more abstract terms (e.g., by *actions* or *guarded commands*) that can be directly translated into BDDs. This allows for an implicit way to partition a set of transitions according to their improvement. Assume that a set of transitions are represented by a BDD $T(\vec{s}, \vec{s}')$. Given a BDD $h(\vec{s}, \vec{v})$ encoding the heuristic function, such that \vec{v} is a bit vector representation of the h -value associated with state s , the set of transitions with improvement equal to k is

$$T(\vec{s}, \vec{s}') \wedge h(\vec{s}, \vec{v}) \wedge h(\vec{s}', \vec{v}') \wedge \vec{v} - \vec{v}' = \vec{k}.$$

The improvement partitioning is computed only once prior to the search, and in practice it turns out that it often can be produced directly from the description of transitions or by splitting the disjunctive partitioning. In fact, for the heuristics we have studied so far, no BDD encoding of the heuristic function has been necessary.

SetA* uses two main data structures: a prioritized queue Q and a reach structure R . Each node in Q contains a BDD representing a set of states with particular g and h values. The node with lowest f -value has highest priority. Ties are solved by giving highest priority to the entry with lowest h -value. An important parameter of Q is an upper bound u on the BDD sizes. When inserting a new node it is unioned with an existing node in Q with the same g and h value if the sum of the size of their two BDDs is less than u . Otherwise a new entry is created for the node. The reach structure is for loop detection. R keeps track of the lowest g -value of every reached state and is used to prune states from a set of next states already reached with a lower g -value. The algorithm is shown in Figure 2. All sets and set operations are carried out with BDDs. SetA* takes five arguments. **IP** is the improvement partitioning described above. *init* and *goal* are the initial and goal states of the search. u is the upper bound parameter of Q and w is the usual weight parameter of weighted A*. Initially the algorithm inserts the

initial state in Q . Observe that the h -value of the initial state has to be found. However since *init* is a single state this is trivial. Similar to the regular A* algorithm, SetA* continues popping the top node of the queue until the queue is either empty or the states of the top node overlaps with the goal. The top node is expanded by finding the image of it for each improvement partition in turn (1.9). Before being inserted in Q , the new nodes are pruned for states seen at a lower search depth, and the reach structure is updated (1.10-14). If the loop was aborted due to Q being empty no solution path exists. Otherwise the path is extracted by applying transitions backwards on the states in R from one of the reached goal states.

SetA* is *sound* due to the soundness of the image computation. Since no states reached by the search are pruned, SetA* is also *complete*. Given an *admissible heuristic* and $w = 0.5$, SetA* further finds *optimal length paths*. As for A*, the reason is that a state on the optimal path eventually will reach the top of Q because states on finalized but sub-optimal paths have higher f -value (Pearl 1984).

The upper bound u can be used to adjust how many states SetA* expands. If each partition in **IP** contains a single transition and $u = 0$ then SetA* specializes to A*. For problems with many shortest length solution paths like the DVM and Logistics described in the next section, it may be an advantage to focus on a subset of them by choosing a low u -value. A similar approach is used by A_ϵ^* described in (Pearl 1984).

The weight w has the usual effect. For $w = 0.5$ Set A* behaves like A*. For $w = 1.0$ it performs best-first search, and for $w = 0.0$ it carries out a regular breadth-first search. The fact that w can take any value in the range $[0, 1]$ is important in practice, since it can be used to increase an underestimating heuristic or decrease an overestimating heuristic.

We end this section by demonstrating SetA* on our example problem. For this demonstration we assume $w = 0.5$ and $u = \infty$. The heuristic function is the vertical distance to the goal state. In Figure 1 the states have been labeled with h -values. We see that **IP** must contain at least three partitions: one containing transition d that improves by minus one, one containing a and c that improve by zero, and one containing b that improves by one. Initially we have

$$\begin{aligned} Q_0 &= < (f = 0.5, g = 0, h = 1, \{(0, 0)\}) > \\ R_0 &= < (g = 0, \{(0, 0)\}) > . \end{aligned}$$

In the first iteration, state $(0, 0)$ is expanded to one child containing state $(1, 0)$ and one child containing $(0, 1)$. According to the improvements of the partitions, we get

$$\begin{aligned} Q_1 &= < (f = 0.5, g = 1, h = 0, \{(0, 1)\}), \\ &\quad (f = 1.0, g = 1, h = 1, \{(1, 0)\}) > \\ R_1 &= < (g = 0, \{(0, 0)\}), (g = 1, \{(0, 1), (1, 0)\}) > . \end{aligned}$$

In the second iteration, only the c transition can fire resulting in

$$\begin{aligned} Q_2 &= < (f = 1.0, g = 2, h = 0, \{(0, 1)\}), \\ &\quad (f = 1.0, g = 1, h = 1, \{(1, 0)\}) > \\ R_2 &= < (g = 0, \{(0, 0)\}), (g = 1, \{(0, 1), (1, 0)\}), \\ &\quad (g = 2, \{(1, 1)\}) > . \end{aligned}$$

The tie breaking rule causes the goal state to be at the top of Q at the beginning of the third iteration. Thus the while loop is aborted and the solution path $(0, 0), (0, 1), (1, 1)$ is extracted from R_2 .

Experimental Evaluation

SetA* has been implemented in the UMOP multi-agent planning framework (Jensen & Veloso 2000) to study its performance characteristics relative to blind bidirectional BDD-based breadth-first search (also implemented in UMOP) and an A* implementation with explicit state representation and cycle detection. In a second evaluation round we developed a domain independent STRIPS planning system called DOP. The state encoding and heuristic function used by the MIPS planner (Edelkamp & Helmert 2001) was reproduced in order to conduct a fair comparison with BDDA* implemented in MIPS. In addition to SetA*, two blind BDD-based breadth-first search algorithms were implemented in DOP, one searching forward and one searching backward.

All experiments were carried out on a Linux 5.2 PC with a 500 MHz Pentium 3 CPU, 512 KB L2 cache and 512 MB RAM. The time limit (TIME) was 600 seconds and the memory limit (MEM) was 450 MB. For UMOP and DOP the number allocated BDD nodes of the BDD-package and the number of partitions in the disjunctive partitioning were hand-tuned for best performance.

Artificial Problems

Two problems IG^k and $D^xV^yM^z$ were defined and studied using the minimum Hamming distance to a goal state as heuristic function. In these experiments the improvement partitioning was computed by splitting a disjunctive partitioning using a specialized BDD-function. Given an improvement k , this function traverses the BDD of an action and picks transitions of the action improving k . The complexity of the function is linear in the size of the action BDD when the goal is a conjunction and the variable ordering interleaves current and next state variables.

IG^k This problem is simplest to define using the STRIPS language (Fikes & Nilsson 1971). A state is a set of facts and an action is a triple of sets of facts. In a given state S , an action (pre, add, del) is applicable if $pre \subseteq S$, and the resulting state is $S' = (S \cup add) \setminus del$. The actions are

$$\begin{array}{lll} \mathbf{A}_1^1 & \mathbf{A}_j^1 \ j = 2, \dots, n & \mathbf{A}_j^2 \ j = 1, \dots, n \\ pre : \{I^*\} & pre : \{I^*, G_{j-1}\} & pre : \{\} \\ add : \{G_1\} & add : \{G_j\} & add : \{I_j\} \\ del : \{\} & del : \{\} & del : \{I^*\}. \end{array}$$

The initial state is $\{I^*\}$ and the goal state is $\{G_j | k < j \leq n\}$. Only \mathbf{A}_j^1 actions should be applied to reach the goal. Applying an \mathbf{A}_j^2 action in any state leads to a wild path since I^* is deleted. The only solution is $\mathbf{A}_1^1, \dots, \mathbf{A}_n^1$ which is non-trivial to find, since the heuristic gives no information to guide the search on the first k steps. The purpose of the experiment is to investigate how well SetA* copes with this situation compared to A*. For SetA* $w = 0.5$ and $u = \infty$. n of IG^k is 16. The results are shown in Table 1. The experiment shows a

k	SetA*		A*	
	Time	#it	Time	#it
0	0.2	16	0.1	16
2	0.2	16	0.4	145
4	0.2	16	7.4	2861
6	0.2	16	80.1	24931
8	0.2	16	344.0	90080
10	0.2	16	TIME	-
12	0.2	16	TIME	-
14	0.2	16	TIME	-

Table 1: Results for the IG^k problem. Time is in seconds. #it is the number of iterations.




Figure 3: The initial state of $D^5V^3M^7$.

fast degradation of A*'s performance with the number of unguided steps. A* gets lost expanding an exponentially growing set of states on wild paths. SetA* is hardly affected by the lack of guidance. The reason is that all transitions on the unguided part improve by zero. Thus on this part, SetA* performs a regular BDD-based breadth-first search, which scales well due to the structure of the problem.

$D^xV^yM^z$ In this domain a set of sliders are moved between the corner positions of hypercubes. In any state, a corner position can be occupied by at most one slider. The dimension of the hypercubes is y . There are z sliders of which x are moving on the same cube. The remaining $z - x$ sliders are moving on individual cubes. Figure 3 shows the initial state of $D^5V^3M^7$. The purpose of $D^xV^yM^z$ is to investigate the performance of SetA* for hard combinatorial problems relative to A* and BDD-based breadth-first search. In this experiment we study the $D^xV^4M^{15}$ problem. We also show the results of PreSetA*, a earlier version of SetA* finding the next states and splitting them in two separate phases. Both versions of SetA* were run with $w = 0.5$ and $u = 200$. The results are shown in Table 2. For this problem disjunctive partitioning of the transition relation is crucial for large values of x . Despite applying this technique, BDD-based bidirectional search does not scale due to a blow-up of the search fringe in both directions. A* works well when x is small since f is a perfect or near perfect discriminator. However when the quality of the heuristic degrades A* gets lost tracking equally promising paths. The good performance of SetA* is due to the upper bound setting that focuses the search on a reasonable subset of the paths. Interestingly the search time is very low even for the hardest problems. Time and memory are spent on building and splitting the transition relation. Separating the next state

x	SetA*	PreSetA*	A*	BiDir
1	0.6	0.8	1.1	0.7
2	0.7	0.9	1.1	0.7
3	0.6	1.4	1.1	1.6
4	0.6	1.5	1.1	8.1
5	0.6	3.5	1.0	334.0
6	0.8	14.4	TIME	TIME
7	1.3	39.8	TIME	TIME
8	2.1	50.7	TIME	TIME
9	6.8	202.6	TIME	TIME
10	16.3	297.2	TIME	TIME
11	39.3	TIME	TIME	TIME
12	MEM	TIME	TIME	TIME

Table 2: Results for the $D^xV^4M^{15}$ problem.

computation and the splitting as done by the earlier version of SetA* seems to come with a large performance penalty.

Planning Problems

Like MIPS, the DOP planning system uses an approximation to the HSPr heuristic (Bonet & Geffner 1999) for STRIPS domains. In addition, it performs similar analysis to minimize the state encoding length. HSPr is an efficient but non-admissible heuristic. We approximate it by summing the depth $d(f)$ of each fact in a state given by a relaxed forward breadth-first search. The heuristic is applied in a backward search from the goal states to the initial state. For any action (*pre*, *add*, *del*) leading from S to S' (when applied in forward direction), we assume

$$del \subseteq pre \text{ and } add \not\subseteq pre.$$

Since the search is backward the improvement of the action is

$$\begin{aligned}
impr &= h(S') - h(S) \\
&= h(S' \cap (pre \cup add)) - h(S \cap (pre \cup add)) \\
&= \sum_{f \in add \setminus S} d(f) - \sum_{f \in del} d(f).
\end{aligned}$$

Thus the improvement of an action can be computed without any BDD-based encoding of the heuristic function. Each action is partitioned in up to $2^{|add|}$ sets of transitions with different improvement.

The problems, we consider, are *Gripper* from the STRIPS track of the AIPS-98 planning competition (Long 2000) and *Logistics* from the first round of the STRIPS track of the AIPS-00 planning competition (Bacchus 2001). The purpose of these experiments is to compare the performance of SetA* and BDDA*, not to solve the problems particularly fast. In that case, a more informative heuristic like the FF heuristic (Hoffmann 2001) should be applied.

Gripper This domain considers a robot with two grippers moving an increasing number of balls between two connected rooms. The first experiment compares forward BDD-based breadth-first search, SetA* with $w = 1.0$ and $u = \infty$, backward BDD-based breadth-first search, pure

#p	For-ward	SetA*	Back-ward	Pure BDDA*	BDDA*
2	0.1	0.1	0.1	3.97	3.89
4	0.2	0.3	0.6	8.01	7.13
6	0.4	0.6	1.4	22.37	15.36
8	0.9	1.0	2.9	72.54	47.08
10	1.2	1.4	5.3	317.15	TIME
12	1.7	2.0	9.1	TIME	TIME
14	2.2	2.7	17.2	TIME	TIME
16	3.5	3.5	19.7	TIME	TIME
18	3.9	4.9	27.5	TIME	TIME
20	5.0	5.8	37.2	TIME	TIME

Table 3: Results of the first gripper experiment. #p is the problem number and time is in seconds.

w	p	#it	Time
0.0	125	360	7.6
0.2	125	354	7.9
0.4	125	338	8.0
0.6	125	204	6.1
0.8	125	204	6.2
1.0	125	204	5.9

Table 4: Results of the second gripper experiment for problem 20. w is the weight, $|p|$ is the solution length and #it is the number of iterations. Time is in seconds.

BDDA* and BDDA*. Pure BDDA* performs best-first search. The results are shown in Table 3. All algorithms find optimal length solutions. BDDA* spends up to 10% of the time computing BDD formulas for the arithmetic operations. During the search even the moderate growth of the search fringe impairs the efficiency considerably. The problem turns out to be most efficiently solved with BDD-based breadth-first search. However it is challenged by SetA* even though this algorithm relies on the slower backward expansion.

The second experiment shows the impact of the weight setting in problem 20. The results are shown in Table 4. Even though SetA* can solve the problem performing breadth-first search the heuristic improves its performance.

Logistics This domain considers moving packages with trucks between sub-cities and with airplanes between cities. In the first experiment SetA* was run with $w = 1.0$ and $u = 200$. The results are shown in Table 5. The HSPr heuristic is very efficient in this domain. The solutions of BDDA* are a couple of steps shorter than the solutions of SetA*. SetA* and pure BDDA* produce plans of similar quality. Again the construction of arithmetic formulas takes up considerable time for BDDA*. The upper bound of 200 speeds up SetA* on the last five problems where the fringe BDDs otherwise grow considerably.

The second experiment was carried out on problem 7 of the Logistics domain. In this experiment SetA* was run with $u = \infty$. The results are shown in Table 6. As depicted HSPr

#p	SetA*	Pure BDDA*	For- ward	BDDA*	Back- ward
4	0.2	6.5	0.3	7.7	0.38
5	0.3	6.7	0.5	9.5	0.81
6	0.3	6.7	0.4	8.4	0.94
7	0.9	13.9	99.0	TIME	396.93
8	1.0	14.1	59.5	138.5	TIME
9	0.9	14.0	100.0	132.6	TIME
10	2.5	25.1	MEM	TIME	TIME
11	2.2	25.2	MEM	TIME	TIME
12	2.0	24.9	MEM	TIME	TIME
13	8.5	57.5	MEM	TIME	TIME
14	7.7	56.7	MEM	TIME	TIME
15	7.3	53.9	MEM	TIME	TIME

Table 5: Results of the first logistics experiment. #p is the problem number and time is in seconds.

w	p	#it	Time
0.0	25	279	8.6
0.2	25	203	8.9
0.4	25	102	4.7
0.6	29	49	0.9
0.8	31	31	0.8
1.0	31	31	0.9

Table 6: Results of the second logistics experiment. w is the weight, $|p|$ is the solution length and #it is the number of iterations. Time is in seconds.

is a good heuristic for this domain increasing the speed significantly while preserving a relative high solution quality.

Conclusion and Outlook

In this paper, we have successfully combined BDD-based search and heuristic search into a new search paradigm. The experimental evaluation of SetA* proves it a powerful algorithm often several orders of magnitude faster than BDD-based breadth-first search and A*. Today planning problems are efficiently solved by heuristic single-state search algorithms. However as recently noticed, the success may be due to an inherent simplicity of the benchmark domains when using the right heuristics (Hoffmann 2001). For less domain-tuned heuristics, we believe that the ability of SetA* to explore an exponentially growing set of paths in polynomial time is essential. Our ongoing research includes identifying such problems and comparing the performance of SetA* and single-state search algorithms.

Acknowledgments

This research is sponsored in part by the Danish Research Agency and the United States Air Force under Grants Nos F30602-00-2-0549 and F30602-98-2-0135. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency, the Air Force, or the US Government.

References

- Bacchus, F. 2001. AIPS'00 planning competition : The fifth international conference on artificial intelligence planning and scheduling systems. *AI Magazine* 22(3):47–56.
- Bloem, R.; Ravi, K.; and Somenzi, F. 2000. Symbolic guided search for CTL model checking. In *Proceedings of the 37th Design Automation Conference (DAC'00)*, 29–34. ACM.
- Bonet, B., and Geffner, H. 1999. Planning as heuristic search: New results. In *Proceedings of the European Conference on Planning (ECP-99)*. Springer.
- Bryant, R. E. 1986. Graph-based algorithms for boolean function manipulation. *IEEE Transactions on Computers* 35:677–691.
- Cimatti, A.; Giunchiglia, E.; Giunchiglia, F.; and Traverso, P. 1997. Planning via model checking: A decision procedure for AR. In *Proceedings of the 4th European Conference on Planning (ECP'97)*, 130–142. Springer.
- Clarke, E.; Grumberg, O.; and Peled, D. 1999. *Model Checking*. MIT Press.
- Edelkamp, S., and Helmert, M. 2001. MIPS the model-checking integrated planning system. *AI Magazine* 22(3):67–71.
- Edelkamp, S., and Reffel, F. 1998. OBDDs in heuristic search. In *Proceedings of the 22nd Annual German Conference on Advances in Artificial Intelligence (KI-98)*, 81–92. Springer.
- Fikes, R. E., and Nilsson, N. J. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence* 2:189–208.
- Hart, P. E.; Nilsson, N. J.; and Raphael, B. 1968. A formal basis for heuristic determination of minimum path cost. *IEEE Transactions on SSC* 100(4).
- Hoffmann, J. 2001. Local search topology in planning benchmarks: An empirical analysis. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, 453–458. Morgan Kaufmann.
- Jensen, R., and Veloso, M. M. 2000. OBDD-based universal planning for synchronized agents in non-deterministic domains. *Journal of Artificial Intelligence Research* 13:189–226.
- Long, D. 2000. The AIPS-98 planning competition. *AI Magazine* 21(2):13–34.
- McMillan, K. L. 1993. *Symbolic Model Checking*. Kluwer Academic Publ.
- Pearl, J. 1984. *Heuristics : Intelligent Search Strategies for Computer Problem Solving*. Addison-Wesley.
- Pohl, I. 1970. First results on the effect of error in heuristic search. *Machine Intelligence* 5:127–140.
- Yang, C. H., and Dill, D. L. 1998. Validation with guided search of the state space. In *Proceedings of the 35th Design Automation Conference (DAC'98)*, 599–604. ACM.