

Generalisation of Alpha-Beta Search for AND-OR Graphs With Partially Ordered Values*

Junkang Li^{1,2}, Bruno Zanuttini², Tristan Cazenave^{1,3} and Véronique Ventos¹

¹NukkAI, Paris, France

²Normandie Univ.; UNICAEN, ENSICAEN, CNRS, GREYC, 14 000 Caen, France

³LAMSADE, Université Paris-Dauphine, PSL, CNRS, France

junkang.li@nukk.ai, bruno.zanuttini@unicaen.fr, tristan.cazenave@lamsade.dauphine.fr, vventos@nukk.ai

Abstract

We define a new setting related to the evaluation of AND-OR directed acyclic graphs with partially ordered values. Such graphs arise naturally when solving games with incomplete information (e.g. most card games such as Bridge) or games with multiple criteria. In particular, this setting generalises standard AND-OR graph evaluation and computation of optimal strategies in games with complete information.

Under this setting, we propose a new algorithm which uses both alpha-beta pruning and cached values. In this paper, we present our algorithm, prove its correctness, and give experimental results on random games and on a card game with incomplete information.

1 Introduction

Search in graphs containing AND- and OR-nodes is used as a basis of many algorithmic solutions to Artificial Intelligence problems. In such graphs, OR-nodes typically model choice nodes where an agent can choose a successor, while AND-nodes model an opponent. For instance, in robust planning with nondeterministic actions, an AND-node models the outcome of an action: a strategy must be valid whatever the outcome [Kissmann and Edelkamp, 2009]. Similarly, when solving a zero-sum two-player (sequential) game for a player, OR-nodes are those at which it is her turn to play, while AND-nodes correspond to her opponent [Van Den Herik *et al.*, 2002]: the value of an OR-node is 1 if and only if at least one move leads to a node with value 1; dually, at AND-nodes, the values of the children are conjoined. More generally, in games that can have more than two outcomes, such as chess or checkers [Schaeffer *et al.*, 2007], AND-nodes (respectively OR-nodes) correspond to a minimum (respectively maximum) operator on the values of their children.

A fundamental question is that of evaluating rooted AND-OR directed acyclic graphs (DAGs), which means computing the value of their root given a value for each of their leaves. For instance, in games with complete information, selecting

the best move for the current turn amounts to evaluate each of the children of the root. This problem has been thoroughly studied in the literature [Marshall, 1986] under the setting of totally ordered values (Boolean or real) and the standard AND/OR or min/max operators. Following [Ginsberg and Jaffray, 2002], we investigate here a more general setting, where the values are taken from a distributive lattice (V, \wedge, \vee) (i.e. a partially ordered set with least upper bound and greatest lower bound for any two elements), and operators for AND- and OR nodes are taken to be the meet \wedge and join \vee , respectively. This setting arises naturally in many applications, in particular in games with incomplete information. Example of such games are Skat [Kupferschmid and Helmert, 2006; Rebstock *et al.*, 2019; Edelkamp, 2020], Bridge [Levy, 1989; Ginsberg, 2001; Cazenave and Ventos, 2020], Hearts and Spades [Sturtevant and White, 2006].

A well-known technique for evaluating AND-OR graphs is alpha-beta pruning, which maintains a lower bound α (respectively upper bound β) on the value of each OR-nodes (respectively AND-nodes), and uses them to prune some of their successors. This technique is currently used in strong chess programs [Haworth and Hernandez, 2021] combined with sophisticated evaluation functions such as NNUE neural networks that were first used in Shogi [Nasu, 2018]. However, the generalisation of alpha-beta pruning to AND-OR DAGs with partially ordered values is nontrivial since two values from a lattice are not always comparable. We build on the seminal work by [Ginsberg and Jaffray, 2002] and generalise it by proving the correctness of lattice-valued alpha-beta pruning with the consideration for heuristic functions.

Orthogonally, we investigate caching techniques for alpha-beta pruning in lattice-valued DAGs. The question is again nontrivial because nodes are in general revisited with different α and β than during previous visits. For this, we propose a new algorithm called ‘alpha-beta duo’. We state its correctness and experimentally evaluate its efficiency.

The paper is organised as follows. Preliminaries are given in Sections 2 and 3. We extend the work by [Ginsberg and Jaffray, 2002] in Section 4, and present alpha-beta duo in Section 5. We then report experimental results and conclude.

2 Preliminaries

The following definitions on posets and lattices are based on [Davey and Priestley, 2002].

*A long version with full proofs is available at the URL <https://hal.archives-ouvertes.fr/hal-03656224>.

Definition 1. Let V be a set and \preceq be a binary relation on V . Then (V, \preceq) is called a partially ordered set (poset) if \preceq is a partial order (i.e. reflexive, transitive, and antisymmetric).

For a poset (V, \preceq) and $S \subseteq V$, an element $x \in V$ is called an upper bound (UB) of S if $s \preceq x$ holds for all $s \in S$, and x is called a least upper bound (LUB) if in addition $x \preceq y$ holds for any UB y of S . If S has an LUB, then it is unique. The greatest lower bound (GLB) of S is defined dually. For $x, y \in V$, we write $x \vee y$ (x join y) and $x \wedge y$ (x meet y) respectively for the LUB and GLB of $\{x, y\}$, when they exist.

Definition 2. A poset (V, \preceq) is called a distributive lattice if

- for all $x, y \in V$, $x \vee y$ and $x \wedge y$ exist,
- for all $x, y, z \in V$, $x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z)$,
- for all $x, y, z \in V$, $x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$.

It is moreover said to be bounded if there are elements $\perp, \top \in V$ satisfying $\perp \preceq x$ and $x \preceq \top$ for any $x \in V$.

In the remainder of this paper, we denote by $(V, \preceq, \wedge, \vee)$ an arbitrary bounded distributive lattice.

Example 1. Let S be a set and let 2^S denote its powerset. Then $(2^S, \subseteq, \cap, \cup)$ is a bounded distributive lattice with set inclusion \subseteq as partial order, set intersection \cap and set union \cup respectively as meet and join, \emptyset as \perp , and S as \top .

We denote any directed acyclic graph (DAG) by $G = (N, C)$, where N is the set of nodes, and $C : N \rightarrow 2^N$ is a function that yields the set of children of each node. A root r is a node without predecessor (i.e. for any $n \in N$, $r \notin C(n)$), and a leaf is a node without child. We denote the set of leaves of a DAG G by L_G . We only consider rooted DAGs, which contain a (necessarily unique) root r such that there exists a directed path to every vertex from r .

An AND-OR DAG is a rooted DAG (N, C, r) equipped with a labelling function $\ell : N \rightarrow \{A, O\}$. Nodes labelled by A and O are respectively called AND-nodes and OR-nodes. Note that we do not impose nodes to be alternating.

3 Problem Setting

We are interested in the problem of evaluating the root value of an AND-OR DAG, given values for all its leaves. Formally, given an AND-OR DAG $G = (N, C, r, \ell)$, a bounded distributive lattice $(V, \preceq, \wedge, \vee)$, and an evaluation function $e : L_G \rightarrow V$ assigning a value in V to each leaf of G , the goal is to compute $v(r)$, where the value $v(n)$ of $n \in N$ is defined recursively by:

- for a leaf node n , $v(n) := e(n)$;
- for an internal AND-node n , $v(n) := \bigwedge_{c \in C(n)} v(c)$;
- for an internal OR-node n , $v(n) := \bigvee_{c \in C(n)} v(c)$.

Since G is a DAG, the function $v : N \rightarrow V$ is well-defined.

Example 2. Consider the DAGs in Figure 1, where circle and square nodes represent AND-nodes and OR-nodes, respectively. On the left, the lattice is the set of Boolean vectors of length 4 (denoted as words), with bitwise AND and bitwise OR as meet and join, respectively. One can easily verify that $v(r) = 1100$. On the right, the lattice is the set of Boolean vectors of length 3, and $v(R) = 001$.




Figure 1: Two AND-OR DAGs with partially ordered values.

3.1 Example Applications

Many important problems are in fact AND-OR DAG evaluation in disguise. The simplest one is Boolean circuit evaluation. Here AND- and OR-nodes model AND and OR gates, the lattice is the Boolean algebra (i.e. $V = \{0, 1\}$ with $0 \preceq 1$ and logical conjunction and disjunction as meet and join), and the evaluation function encodes the inputs of the circuit.

Solving a game with complete information typically involves computing the minimax value of a game tree, which can be regarded as evaluating an AND-OR DAG: AND- and OR-nodes are respectively choice nodes of player MIN and of player MAX, the lattice is (V, \leq, \min, \max) with V a totally ordered set such as \mathbb{Z} or \mathbb{R} , and the evaluation function gives the value of terminal nodes, or a heuristic value if the search is cut at some depth. Then the root value is the minimax value of the game.

In games with incomplete information, nontrivial lattices come into play. For example, [Ginsberg, 2001] shows that computing the maxmin value of a player amounts to evaluate the game DAG with the lattice $(2^{2^S}, \preceq, \cap, \cup)$, where S is a finite set and $f \sqcap g = \{\alpha \cap \beta \mid \alpha \in f, \beta \in g\}$ for any $f, g \in 2^{2^S}$ (i.e. f and g are sets of subsets of S). We will discuss more about this in Section 6.

4 Alpha-Beta Pruning Under Partial Order

Most of the literature on alpha-beta pruning concerns only totally ordered values, such as real numbers. Since AND-OR DAGs with partially ordered values are useful to model richer problems, [Dasgupta *et al.*, 1996] proposed alpha-beta pruning in this new setting for multi-criteria game. [Ginsberg and Jaffray, 2002] gave the first thorough study on this subject, and proved in particular that deep pruning is sound for rational players if and only if the set of values is a distributive lattice. [Loddo and Saiu, 2010] showed that deep pruning is sound for tropical algebras if rationality is relaxed.

The form of deep α pruning considered by [Ginsberg and Jaffray, 2002] is given in Figure 2 (left). If $v \preceq \alpha$, then the subtree T can be deeply pruned. To show why this definition of deep pruning does not capture every cut an alpha-beta search should perform when values are partially ordered, consider Figure 2 (right). The lattice is again the set of Boolean vectors of length 3, with bitwise AND and bitwise OR as meet and join, respectively. When an alpha-beta search algorithm descends to the bottommost AND-node, the value of α would be 110, which is the join of the value of an already explored child of two ancestor OR-nodes. We would like the algorithm




Figure 2: Deep pruning vs expected pruning.

Algorithm 1: Alpha-beta search

```

1 def AlphaBeta (node n, α, β) :
2   I = h(n, α, β)
3   v ← I
4   determine the successor nodes n1, . . . , nb of n
5   for i in {1, . . . , b}:
6     if n is an OR-node:
7       | α ← α ∨ v
8     else:
9       | β ← β ∧ v
10    if α ≥ β:
11      | break
12    vchild ← AlphaBeta (ni, α, β)
13    if n is an OR-node:
14      | v ← v ∨ vchild
15    else:
16      | v ← v ∧ vchild
17  return v
    
```

to prune subtree T since the value of its parent node cannot be better than the current value of α (due to the sibling of T). However, deep pruning, as it is defined in the literature such as [Ginsberg and Jaffray, 2002], does not apply since the value of α does not come from a child of one single ancestor node. Note that this phenomenon is specific to lattices that are *not* totally ordered, since otherwise the meet or join (i.e. min or max) of two values is always one of them, hence deep pruning captures any pruning in a standard alpha-beta search.

Another question not formally addressed in the literature is the initialisation of node values. In standard alpha-beta search, one typically initialises the value of an OR-node (respectively AND-node) to be α (respectively β) [Knuth and Moore, 1975] or $-\infty$ (respectively $+\infty$) [Marsland, 1986] (note that $-\infty$ and $+\infty$ translate to \perp and \top in our context). However, one may have access to a heuristic evaluation of nodes, typically by evaluating a relaxed version of the problem which is easier to solve. For instance, a player can do no better in a game with incomplete information than in the same game but with complete information. The latter being much easier to solve, the value obtained can be used as a heuristic in the original game with incomplete information. Ideally, initialising values with an accurate heuristic should accelerate the search by finding cuts earlier.

In order to fill these two gaps in the literature, we first formalise alpha-beta search under partial order with initialisation function in Algorithm 1. We denote by h the initialisation

function. In general, its value depends on the current α and β , so we define it to yield a value $h(n, \alpha, \beta)$ for any node n and bounds α and β . Note that since non-trivial initial value can be used for a node n (Line 3), a cut may happen even before the first child of n is explored, hence we update α and β (Lines 7 and 9) and determine whether there is a cut (Line 10) at the beginning of the main loop. This is otherwise the same algorithm as in the literature (for instance, [Marsland, 1986]).

It can be seen that Algorithm 1 will perform the wished pruning in the example in Figure 2 (right). By mimicking the proof by [Knuth and Moore, 1975], we prove that such pruning is indeed sound, thereby extending the result by [Ginsberg and Jaffray, 2002] to its full form, provided that the initialisation function h satisfies a certain admissibility condition:

Definition 3. A heuristic function h is said to be admissible for G and V if for any node n in G and any $\alpha, \beta \in V$,

$$h(n, \alpha, \beta) \begin{cases} = v(n) & \text{if } n \text{ is a leaf node;} \\ \succeq v(n) \wedge \beta & \text{if } n \text{ is an AND-node;} \\ \preceq v(n) \vee \alpha & \text{if } n \text{ is an OR-node.} \end{cases}$$

Note that this condition is satisfied for the initialisation values usually used in the literature, such as α or $-\infty$ for OR-nodes (and β or $+\infty$ for AND-nodes). It is also satisfied when $h(n, \alpha, \beta)$ overestimates $v(n)$ for internal AND-nodes and underestimates it for internal OR-nodes.

We denote the value returned by Algorithm 1 with input n, α, β by $f(n, \alpha, \beta)$. The correctness of Algorithm 1 is a consequence of the following central result.

Proposition 1. If h is an admissible heuristic function for G and V , then for any node n of G and any $\alpha, \beta \in V$, we have $\beta \wedge (\alpha \vee f(n, \alpha, \beta)) = \beta \wedge (\alpha \vee v(n))$.

Intuitively, Proposition 1 states that the value returned by Algorithm 1 is the exact value of n up to a factor of α and β . So even for partially ordered values, alpha-beta search can be interpreted as search with a pruning window.

Now it follows that Algorithm 1 is correct in the sense that if we use a lower and an upper bound as the initial search window at a node, we can recover its exact value using the returned value of the algorithm.

Corollary 1. If h is admissible for G and V , then for any node n of G and any $\underline{v}, \bar{v} \in V$ satisfying $\underline{v} \preceq v(n) \preceq \bar{v}$, we have $v(n) = \bar{v} \wedge (\underline{v} \vee f(n, \underline{v}, \bar{v}))$. In particular, $v(n) = f(n, \perp, \top)$.

Proof. Plug $\alpha = \underline{v}$ and $\beta = \bar{v}$ into the equality in Proposition 1, and use the fact that $\underline{v} \vee v(n) = v(n) = \bar{v} \wedge v(n)$. \square

Importantly, contrary to the case of totally ordered values, it is *not* true in general that the stronger equality $v(n) = f(n, \underline{v}, \bar{v})$ holds, as the example in Figure 1 (left) shows. Let $h(r, \cdot, \cdot) = \perp = 0000$ and $h(n, \cdot, \cdot) = \top = 1111$, so that h is admissible. Recall that $v(r) = 1100$. For $\underline{v} = 1000$ and $\bar{v} = 1110$, indeed $\underline{v} \preceq v(r) \preceq \bar{v}$. However, $f(r, \underline{v}, \bar{v}) = 1101 \neq v(r)$ since an α -cut happens in the call on n (Line 10) after examining its first child, since at this point $\alpha = \beta = 1100$. However, since $\alpha = \underline{v}$ and $\beta = \bar{v}$ yield no constraint on the fourth component of the values, we still have $v(r) = \bar{v} \wedge (\underline{v} \vee f(r, \underline{v}, \bar{v}))$, as stated in Corollary 1.

5 Alpha-Beta Duo Algorithm

We now come to the main contribution of our work, namely a caching scheme for reusing previously computed values in an alpha-beta search for partially ordered values. The trouble of standard alpha-beta search is that the returned value is equal to the exact value only up to a factor of α and β . It is therefore a nontrivial question how to reuse a previously computed value, since subsequent revisits of a node may come with different search windows.

In alpha-beta search with cache for totally ordered values (for instance [Marsland, 1986]), one can exploit the fact that with usual value initialisation, the value $f(n, \alpha, \beta)$ satisfies $f(n, \alpha, \beta) < \beta \Rightarrow v(n) \leq f(n, \alpha, \beta)$ and $f(n, \alpha, \beta) > \alpha \Rightarrow v(n) \geq f(n, \alpha, \beta)$. In particular, if $\alpha < f(n, \alpha, \beta) < \beta$, then $f(n, \alpha, \beta) = v(n)$. Hence by comparing $f(n, \alpha, \beta)$ to α and β , one can determine whether it is exact, a lower, or an upper bound, and store it in the cache with an appropriate flag.

However, this does not hold in general for partially ordered values, as shown on Figure 1 (right). For $\alpha = 010$ and $\beta = 110$, a β -cut happens after evaluating the first child of C , and an α -cut after evaluating the first child of D . Hence, the algorithm returns 010 for R , which is neither a lower nor an upper bound of the exact value 001. In fact, these two values are incomparable in the lattice. If R is an internal node in a DAG, then caching this returned value 010 may cause an evaluation error when the algorithm revisits R .

To tackle this difficulty, we propose a new algorithm named ‘alpha-beta duo’, which computes a pair of values for all nodes instead of one single value. The algorithm is presented in Algorithm 2, where `Cache` refers to a transposition table the entries of which are pairs of values indexed by nodes of G , and (\underline{h}, \bar{h}) refers to a pair of initialisation functions for which we assume the following property.

Definition 4. A pair (\underline{h}, \bar{h}) of initialisation functions is said to be admissible for G and V if for any node n in G , $\underline{h}(n) \leq v(n) \leq \bar{h}(n)$ holds, and in addition, if n is a leaf node, $\underline{h}(n) = \bar{h}(n) = v(n)$ holds.

In other words, admissible \underline{h} and \bar{h} respectively underestimates and overestimates the value of a node. Note that \underline{h} and \bar{h} that assign respectively \perp and \top to all internal nodes form an admissible pair, which can always be used if one does not have better heuristic functions.

Alpha-beta duo search works in the following manner:

- First, variables \underline{c} and \bar{c} denote respectively the best lower and upper bound of $v(n)$ available to the algorithm before this call. If n has already been visited, then \underline{c} and \bar{c} are retrieved from the cache. Otherwise, they are given by the initialisation functions \underline{h} and \bar{h} . If $\underline{c} = \bar{c}$, (\underline{c}, \bar{c}) is returned immediately.
- Otherwise, by symmetry consider the case when n is an OR-node. During the main loop, \underline{v} and \bar{v} are respectively the cumulative lower and upper bound of $v(n)$ (notice that they are both initialised to \perp for an OR-node). If a cut ever happens, it means not all children of n have been evaluated, hence \bar{v} is not a valid upper bound of $v(n)$. Then we take \bar{v} to be \bar{c} , the best upper bound previously known. On the other hand, \underline{v} , which is the join of lower

Algorithm 2: Alpha-beta duo search

```

1 def AlphaBetaDuo (node  $n$ ,  $\alpha$ ,  $\beta$ ) :
2   if there is an entry for  $n$  in the cache:
3     |  $(\underline{c}, \bar{c}) \leftarrow \text{Cache}(n)$ 
4   else:
5     |  $(\underline{c}, \bar{c}) \leftarrow (\underline{h}(n), \bar{h}(n))$ 
6   if  $\underline{c} = \bar{c}$ :
7     | return  $(\underline{c}, \bar{c})$ 
8    $\alpha \leftarrow \alpha \vee \underline{c}$ 
9    $\beta \leftarrow \beta \wedge \bar{c}$ 
10  if  $n$  is an OR-node:
11    |  $(\underline{v}, \bar{v}) \leftarrow (\perp, \perp)$ 
12  else:
13    |  $(\underline{v}, \bar{v}) \leftarrow (\top, \top)$ 
14  determine the children  $n_1, \dots, n_b$  of  $n$ 
15  for  $i$  in  $\{1, \dots, b\}$ :
16    | if  $\alpha \succeq \beta$ :
17      | if  $n$  is an OR-node:
18        |  $\bar{v} = \bar{c}$ 
19      | else:
20        |  $\underline{v} = \underline{c}$ 
21      | break
22    |  $\underline{v}', \bar{v}' \leftarrow \text{AlphaBetaDuo}(n_i, \alpha, \beta)$ 
23    | if  $n$  is an OR-node:
24      |  $\underline{v} \leftarrow \underline{v} \vee \underline{v}'$ 
25      |  $\bar{v} \leftarrow \bar{v} \vee \bar{v}'$ 
26      |  $\alpha \leftarrow \alpha \vee \underline{v}'$ 
27    | else:
28      |  $\underline{v} \leftarrow \underline{v} \wedge \underline{v}'$ 
29      |  $\bar{v} \leftarrow \bar{v} \wedge \bar{v}'$ 
30      |  $\beta \leftarrow \beta \wedge \bar{v}'$ 
31  |  $\underline{v} \leftarrow \underline{v} \vee \underline{c}$ 
32  |  $\bar{v} \leftarrow \bar{v} \wedge \bar{c}$ 
33  store  $(\underline{v}, \bar{v})$  in the cache under an entry for  $n$ 
34  return  $(\underline{v}, \bar{v})$ 

```

bounds of children of n that have been evaluated, is a valid lower bound so we keep it.

- Finally, after the main loop, \underline{v} and \bar{v} are the lower and upper bounds of $v(n)$ computed by the current call. Hence they are combined with the previously known bounds \underline{c} and \bar{c} to yield to best currently known bounds on $v(n)$ and they are cached.

We now prove that alpha-beta duo is correct. For this, we first need the following notion.

Definition 5. A cache `Cache` is said to be coherent for G and V if for any node n in G , if there is an entry for n in the cache, then $\text{Cache}(n) = (\underline{c}, \bar{c})$ where $\underline{c} \preceq v(n) \preceq \bar{c}$, and in addition, if n is a leaf node, then $\underline{c}(n) = \bar{c}(n) = v(n)$.

Obviously, an empty cache is coherent for any G and V .

In the following, we denote the pair of values returned by Algorithm 2 with input n , α , β by $(f(n, \alpha, \beta), \bar{f}(n, \alpha, \beta))$. We first show that if the cache is initially coherent, then it remains coherent after the execution, and that any interval stored in it cannot become looser.

Proposition 2. *If (\underline{h}, \bar{h}) is admissible and Cache is initially coherent for G and V , then for any node n in G and any $\alpha, \beta \in V$, we have $\underline{f}(n, \alpha, \beta) \preceq v(n) \preceq \bar{f}(n, \alpha, \beta)$. Moreover, if there is an entry (\underline{c}, \bar{c}) in the cache for n before the call, then $\underline{c} \preceq \underline{f}(n, \alpha, \beta)$ and $\bar{f}(n, \alpha, \beta) \preceq \bar{c}$.*

We can now prove results parallel to those in Section 4.

Proposition 3. *If (\underline{h}, \bar{h}) is admissible and Cache is initially coherent for G and V , then for any node n in G and any $\alpha, \beta \in V$ we have $\beta \wedge (\alpha \vee \underline{f}(n, \alpha, \beta)) = \beta \wedge (\alpha \vee \bar{f}(n, \alpha, \beta))$.*

Corollary 2. *If (\underline{h}, \bar{h}) is admissible and Cache is initially coherent for G and V , then for any node n in G and any $\underline{v}, \bar{v} \in V$ satisfying $\underline{v} \preceq v(n) \preceq \bar{v}$, we have $v(n) = \bar{v} \wedge (\underline{v} \vee \underline{f}(n, \underline{v}, \bar{v})) = \bar{v} \wedge (\underline{v} \vee \bar{f}(n, \underline{v}, \bar{v}))$. In particular, $v(n) = \underline{f}(n, \perp, \top) = \bar{f}(n, \perp, \top)$.*

Proof. By Proposition 2, $\underline{f}(n, \underline{v}, \bar{v}) \preceq v(n) \preceq \bar{f}(n, \underline{v}, \bar{v})$. Using $\bar{v} \wedge (\underline{v} \vee v(n)) = v(n)$, we get

$$\bar{v} \wedge (\underline{v} \vee \underline{f}(n, \underline{v}, \bar{v})) \preceq v(n) \preceq \bar{v} \wedge (\underline{v} \vee \bar{f}(n, \underline{v}, \bar{v})).$$

These inequalities are in fact equalities since from Proposition 3 we have $\bar{v} \wedge (\underline{v} \vee \underline{f}(n, \underline{v}, \bar{v})) = \bar{v} \wedge (\underline{v} \vee \bar{f}(n, \underline{v}, \bar{v}))$. \square

6 Experiments

To assess the efficiency of alpha-beta duo (hereafter ‘ABD’), we ran experiments comparing it to three other algorithms:

- alpha-beta without cache (Algorithm 1, ‘AB’ for short);
- an alpha-beta search which only caches the value computed for a node if no cut is found during the search in the subtree rooted at this node (hereafter ‘ABC’);
- a minimax search algorithm without alpha-beta pruning, but with a cache (hereafter ‘MMC’).

The code of ABD was slightly optimized by refining the values computed on Lines 31 and 32 with the corresponding bounds of all fully explored children. It is easy to show that this preserves the correctness of the algorithm.

For all experiments, we measured the number of nodes of the DAG visited at least once, the total number of node visits (equivalently, the total number of recursive calls), and the time taken for solving the problem. Intuitively, we expect ABD to be better than ABC, ABC to be better than MMC (because MMC does not use alpha-beta pruning), and MMC to be better than AB (because the latter does not cache its results and hence, recomputes several times for the same node).

We used two synthetic sets of benchmarks. The first (hereafter ‘random’) consists of random DAGs of the same kind as the one in Figure 1. Random DAGs with parameters d (depth), b (branching factor), and v (number of variables) are generated in the following manner:

- d layers $0, 1, \dots, d - 1$ are built: layer i consists of $3^{\min(i, \frac{3d}{2} - i)}$ nodes (which yields diamond-shaped DAGs);
- from each node n , a set of b nodes is randomly chosen from nodes in the next layer to be $C(n)$;

- each internal node is labelled AND or OR at random;
- the value of each terminal node is a uniformly drawn subset of $\{1, \dots, v\}$, or equivalently a random Boolean vector of length v .

We also consider strictly alternating DAGs, in which all nodes in layer i are OR-nodes (respectively AND-nodes) if i is even (respectively odd). In particular, the root is an OR-node.

The second set of benchmarks consists of a simplified version of the card game Bridge that we call ‘racing’. There are two players, MIN and MAX. Each has a hand of h cards drawn uniformly from the deck $\{1, \dots, d\}$ where $d \geq 2h$. Players only see their own hand. MIN begins the game. During each trick, the player who begins plays a card from her hand, the other sees it and plays a card in turn. The player who played the highest card wins this trick and starts the next one. No new card is ever drawn from the deck. The game ends when the players have no more cards or when one has won in total g tricks. MAX wins if she is the first one to reach g tricks. For the benchmark, each instance with parameter h, d , and g consists of a randomly drawn hand with h cards for MAX and a randomly drawn card that is supposed to be played by MIN during the first trick. Notice that when $d > 2h$, each player has incomplete information. We use evaluation of AND-OR DAGs to compute optimal strategies for MAX against the best defence adversary model defined in [Frank and Basin, 1998].

In games with incomplete information where S is the set of all possible hidden configurations, [Frank and Basin, 1998; Frank and Basin, 2001] define the maxmin value of player MAX to be the set of all subsets S' of S such that there is a uniform strategy winning in any configuration of S' . [Ginsberg, 2001] shows that computing this value amounts to evaluate the game DAG with the lattice $(2^{2^S}, \preceq, \sqcap, \sqcup)$. Intuitively, using \sqcup at OR-nodes models the fact that player MAX can choose any child as her strategy, and \sqcap at AND-nodes models the fact that a strategy must be robust to all adversarial strategies, hence a strategy wins in $s \in S$ if and only if it wins in s whatever action her opponent chooses.

In the same setting of games with incomplete information, one can also be interested in non-uniform strategies which allow player MAX’s actions to depend on the hidden information. This can be useful for computing heuristic values of for the game with incomplete information. It can be seen that the set of all configurations for which there is a non-uniform winning strategy can be computed as the value of the game DAG with the lattice $(2^S, \subseteq, \cap, \cup)$.

Hence, for both benchmarks, we consider the two lattices $(2^{2^S}, \sqcap, \sqcup)$ and $(2^S, \cap, \cup)$. In ‘random’, $S = \{1, \dots, v\}$, while in ‘racing’ S is the set of all possible hands of player MIN.

For space reasons, we only give the most representative results, in terms of computation time. For each parameter setting and each algorithm, we averaged over 10 runs. Figure 3 shows two examples where, as can be expected, it is more efficient to cache bounds, even more to perform cuts, and still more to compute and store two bounds per node. On the top plot, the gain of using ABD is exponential: with the branching factor increasing, ABD gets a better advantage




Figure 3: Experimental results on random (top) and racing (bottom). Top: $d = 15$, $v = 10$ (varying b), alternating DAGs. Bottom: $d = 20$, $g = 5$ (varying h). The lattice is 2^5 in both cases.

of computing and caching two bounds. On the bottom plot, ABC and AB (not represented) are exponentially worse, and ABD is better than MMC when the branching factor is high.

Now Figure 4 shows two examples where it turns out that it is not always better to use alpha-beta pruning with cache.

On the top plot, not caching results at all turns out to be better: the overhead due to the additional operations from the lattice 2^{2^S} (which are necessary to maintain the cache) seems to compensate the advantage of ABC or ABD in terms of number of visited nodes and recursive calls (the curves are reversed for this metrics, not shown here). On the bottom plot, it turns out that sometimes alpha-beta pruning even degrades performance. Again, this is due to the overhead of manipulating values from a large lattice (for a fixed hand size, the lattice grows exponentially with the deck size).

To complete these results, let us mention that in most experiments, the numbers of nodes explored and visited by each method are ordered as expected, with ratios varying from linear to exponential. In particular, for these metrics, ABD is most of the time better, and often much better, than the other three algorithms.

Summarising, ABD seems to provide a real gain in (brute) performance for DAGs with high branching factors. Contrastingly, when \wedge and \vee from the lattice are too expensive to compute (as is the case in some large lattices), it may sometimes be better not to use cache and alpha-beta pruning together, due to the overhead to maintain the coherence of the cache.




Figure 4: Experimental results on random (top) and racing (bottom). Top: $d = 15$, $b = 4$ (varying v), alternating DAGs, lattice 2^{2^S} . Bottom: $h = 7$, $g = 5$ (varying d), lattice 2^S .

7 Conclusion

We investigated alpha-beta search for AND-OR DAGs with values from a lattice, which has direct applications such as solving games with incomplete information. We have extended previous formal analyses, in particular to the use of heuristic as initialisation functions. Then we have proposed a new algorithm named ‘alpha-beta duo’, which caches both a lower and an upper bound of the value of each visited node, and we have formally proved its correctness. Experiments show that it is more efficient than other algorithms in terms of number of visited nodes and recursive calls. As for time efficiency, alpha-beta duo turns out to be more efficient than other algorithms for DAGs with large branching factors and reasonably-sized lattices. As an interesting conclusion, our experiments also put forth that in other cases, it may be better not to use a cache with alpha-beta pruning.

Our future work includes algorithmic optimisations for alpha-beta search with cache applied to games with incomplete information. We will also investigate the use of efficient knowledge representations [Bienvenu *et al.*, 2010; Niveau and Zanuttini, 2016] to accelerate lattice operations in such context. Another perspective is to apply our work to games with multiple criteria instead of scalar outcomes.

References

- [Bienvenu *et al.*, 2010] Meghyn Bienvenu, H el ene Fargier, and Pierre Marquis. Knowledge Compilation in the Modal

- Logic S5. In *Proc. 24th AAAI Conference on Artificial Intelligence (AAAI 2010)*, pages 261–266, 2010.
- [Cazenave and Ventos, 2020] Tristan Cazenave and Véronique Ventos. The $\alpha\mu$ search algorithm for the game of bridge. In *Proc. Monte Carlo Search International Workshop*, pages 1–16. Springer, 2020.
- [Dasgupta *et al.*, 1996] Pallab Dasgupta, P. P. Chakrabarti, and S. C. De Sarkar. Searching game trees under a partial order. *Artif. Intell.*, 82(1-2):237–257, 1996.
- [Davey and Priestley, 2002] Brian A. Davey and Hilary A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, 2nd edition, 2002.
- [Edelkamp, 2020] Stefan Edelkamp. Representing and reducing uncertainty for enumerating the belief space to improve endgame play in skat. In *Proc. 24th European Conference on Artificial Intelligence (ECAI 2020)*, pages 395–402. IOS Press, 2020.
- [Frank and Basin, 1998] Ian Frank and David A. Basin. Search in games with incomplete information: A case study using bridge card play. *Artif. Intell.*, 100(1-2):87–123, 1998.
- [Frank and Basin, 2001] Ian Frank and David A. Basin. A theoretical and empirical investigation of search in imperfect information games. *Theor. Comput. Sci.*, 252(1-2):217–256, 2001.
- [Ginsberg and Jaffray, 2002] M Ginsberg and Alan Jaffray. Alpha-beta pruning under partial orders. In Richard Nowakowski, editor, *More Games of No Chance*, number 42 in Mathematical Sciences Research Institute Publications, pages 37–48. Cambridge University Press, 2002.
- [Ginsberg, 2001] Matthew L. Ginsberg. GIB: imperfect information in a computationally challenging game. *J. Artif. Intell. Res.*, 14:303–358, 2001.
- [Haworth and Hernandez, 2021] Guy Haworth and Nelson Hernandez. The 20th top chess engine championship, TCEC20. *ICGA Journal*, 43(1):62–73, 2021.
- [Kissmann and Edelkamp, 2009] Peter Kissmann and Stefan Edelkamp. Solving fully-observable non-deterministic planning problems via translation into a general game. In Bärbel Mertsching, Marcus Hund, and Muhammad Zaheer Aziz, editors, *Proc. 32nd Annual German Conference on Artificial Intelligence (KI 2009)*, pages 1–8. Springer, 2009.
- [Knuth and Moore, 1975] Donald E. Knuth and Ronald W. Moore. An analysis of alpha-beta pruning. *Artif. Intell.*, 6(4):293–326, 1975.
- [Kupferschmid and Helmert, 2006] Sebastian Kupferschmid and Malte Helmert. A skat player based on monte-carlo simulation. In *Proc. 5th International Conference on Computers and Games (CG 2006)*, pages 135–147. Springer, 2006.
- [Levy, 1989] David NL Levy. The million pound bridge program. *Heuristic Programming in Artificial Intelligence: The First Computer Olympiad*, pages 95–103, 1989.
- [Loddo and Saiu, 2010] Jean-Vincent Loddo and Luca Saiu. How to correctly prune tropical trees. In Serge Autexier, Jacques Calmet, David Delahaye, Patrick D. F. Ion, Laurence Rideau, Renaud Rioboo, and Alan P. Sexton, editors, *Proc. 10th International Conference on Intelligent Computer Mathematics*, volume 6167 of *Lecture Notes in Computer Science*, pages 101–115. Springer, 2010.
- [Marsland, 1986] T. A. Marsland. A review of game-tree pruning. *J. Int. Comput. Games Assoc.*, 9(1):3–19, 1986.
- [Nasu, 2018] Yu Nasu. Efficiently updatable neural-network-based evaluation functions for computer shogi. *The 28th World Computer Shogi Championship Appeal Document*, 2018.
- [Niveau and Zanuttini, 2016] Alexandre Niveau and Bruno Zanuttini. Efficient Representations for the Modal Logic S5. In *Proc. 25th International Joint Conference on Artificial Intelligence (IJCAI 2016)*, 2016.
- [Rebstock *et al.*, 2019] Douglas Rebstock, Christopher Solinas, Michael Buro, and Nathan R Sturtevant. Policy based inference in trick-taking card games. In *Proc. 2019 IEEE Conference on Games (CoG 2019)*, pages 1–8. IEEE, 2019.
- [Schaeffer *et al.*, 2007] Jonathan Schaeffer, Neil Burch, Yngvi Björnsson, Akihiro Kishimoto, Martin Müller, Robert Lake, Paul Lu, and Steve Sutphen. Checkers is solved. *Science*, 317(5844):1518–1522, 2007.
- [Sturtevant and White, 2006] Nathan R Sturtevant and Adam M White. Feature construction for reinforcement learning in hearts. In *Proc. 5th International Conference on Computers and Games (CG 2006)*, pages 122–134. Springer, 2006.
- [Van Den Herik *et al.*, 2002] H Jaap Van Den Herik, Jos WHM Uiterwijk, and Jack Van Rijswijk. Games solved: Now and in the future. *Artif. Intell.*, 134(1-2):277–311, 2002.