

Optimizing $\alpha\mu$

Tristan Cazenave¹, Swann Legras², Véronique Ventos²

¹ LAMSADE, Université Paris-Dauphine, PSL, CNRS, France

² Nukkai, France

Tristan.Cazenave@dauphine.psl.eu slegras@nukk.ai vventos@nukk.ai

Abstract

$\alpha\mu$ is a search algorithm which repairs two defaults of Perfect Information Monte Carlo search: strategy fusion and non-locality. In this paper we optimize $\alpha\mu$ for the game of Bridge, avoiding useless computations. The proposed optimizations are general and apply to other imperfect information turn-based games. We define multiple optimizations involving Pareto fronts, and show that these optimizations speed up the search. Some of these optimizations are cuts that stop the search at a node, while others keep track of which possible worlds have become redundant, avoiding unnecessary, costly evaluations. We also measure the benefits of parallelizing the double dummy searches at the leaves of the $\alpha\mu$ search tree.

Introduction

The state of the art for imperfect information card games is Perfect Information Monte Carlo sampling (PIMC). It was first proposed by Levy (Levy 1989) for Bridge, and used in the popular program GIB (Ginsberg 2001). PIMC can be used in other trick-taking card games such as Skat (Buro et al. 2009; Kupferschmid and Helmert 2006), Spades and Hearts (Sturtevant and White 2006). Long analyzed the reasons why PIMC is successful in these games (Long et al. 2010). The principle of PIMC is to use determinization to generate possible worlds. For each possible move, PIMC plays the move, samples from a set of possible worlds, and then solves each world exactly as if it was a perfect information game. It then computes the average of the results among the possible worlds to evaluate each move.

Other approaches to imperfect information games are Information Set Monte Carlo Tree Search (Cowling, Powley, and Whitehouse 2012), counterfactual regret minimization that solved Poker (Zinkevich et al. 2008; Brown and Sandholm 2019), and Exploitability Descent (Lockhart et al. 2019). There are also various Reinforcement Learning algorithms for incomplete information games are available in the OpenSpiel framework (Lanctot et al. 2019).

Recursive Monte Carlo Search (Furtak and Buro 2013) is the adaptation of Nested Monte Carlo Search (Cazenave 2009) to multi-player and incomplete information games. The basic principle is the same: it uses Monte Carlo Search

to improve Monte Carlo Search made at a higher level of recursion. Recursive Monte Carlo Search improves on PIMC for Skat and Bridge (Bouzy, Rimbaud, and Ventos 2020), however, it takes much more time to complete than PIMC.

PIMC plays sub-optimally due to two main problems: strategy fusion and non-locality (Frank and Basin 2001). $\alpha\mu$ (Cazenave and Ventos 2020) is an anytime heuristic search algorithm for incomplete information games that assumes perfect information for the opponents. $\alpha\mu$ addresses the strategy fusion and non-locality problems encountered by PIMC. Other programs also address the strategy fusion problem in endgame play, for example, GIB (Ginsberg 2001) in Bridge (using a single dummy solver) and the Skat endgame solver of Stefan Edelkamp (Edelkamp 2020).

In this paper we propose improvements to speedup the $\alpha\mu$ algorithm without affecting its output.

The paper starts with a section on previous work in Bridge and on the $\alpha\mu$ algorithm. The next section describes the optimizations and the modified algorithm. The last section gives experimental results.

Prerequisites and Previous Work

In this section, we introduce some definitions used in the paper, we then briefly explain the game of Bridge, Computer Bridge, previous work on dealing with strategy fusion, non-locality, the early cut and the root cut that were defined previously in the $\alpha\mu$ paper (Cazenave and Ventos 2020).

Definitions

We define general naming conventions used throughout the article.

Vectors Given n different possible worlds, a boolean vector of size n keeps the status of the game for each possible world: A zero at index w means that the game is lost for world number w . A one means the game is won. Associated to the vector there is another vector of booleans indicating whether the world is possible in the current state. At the root of the search all worlds are possible but when an opponent makes a move, the move is usually only valid in some of the worlds and the set of valid worlds is reduced.

Pareto Fronts A set of vectors that are not currently dominated by other sets of vectors.

We say a Pareto front f_1 dominate f_2 iff:

$$\forall x_2 \in f_2, \exists x_1 \in f_1, x_1 \text{ dominates } x_2 \text{ or } x_1 = x_2.$$

We say a vector $x_1 \in \{0, 1\}^n$ dominates x_2 iff:

$$\forall i \in [1, n]: x_1[i] \geq x_2[i] \text{ and}$$

$$\exists i \in [1, n] \text{ such that } x_1[i] > x_2[i].$$

The score of a vector is the average among all possible worlds of the values contained in the vector.

Impossible worlds When the search reaches a node where some worlds have disappeared because the moves that have been played are not possible in these worlds. These worlds are represented as "x" in figures in this article. New impossible worlds appear at min nodes when we consider a move from the defender possible in only a subset of all worlds.

Useless worlds Worlds for which we know for sure that nothing is lost by labelling them 0. These worlds are noted "-" on figures.

Comparison of Pareto Fronts When comparing two Pareto fronts from different depths in the tree or at min node between two actions that contain different impossible worlds, to test for the dominance of one over the other, one must carefully handle impossible worlds as they can be a win higher in the tree. Such impossible worlds must be evaluated as 1 unless proven useless higher in the tree when checking for dominance relation. As such, [1 x 0] is not dominated by [1 0 0] but is dominated by [1 1 0] or even by [1 x -].

Bridge in short

The interested reader can refer for instance to (Mahmood, Grant, and Sharif 2014) for a more complete presentation of the game of Bridge.

Bridge is a trick-taking card game with four players (denoted by West, North, East and South or W,N,E,S) divided in two partnerships (East-West and North-South). A standard 52 card deck is shuffled and each player receives a hand of 13 cards that is only visible to them. A Bridge deal (or *board*) is divided into two major phases: the bidding (out of the scope of the paper) and the card play.

The goal of the bidding is to reach a contract which determines the minimum number of tricks the pair commits to win during the card play, either with no trump (NT) or with a determined suit as trump.

In the following, we assume that the North-South pair reached the contract of 3NT (resp. 7NT) and that South is the agent who plays the board (i.e the *declarer*).

During the card play, the goal is to fulfill (for the declarer) or to defeat (for the *defenders*) the contract reached during the bidding phase. For the contract of 3NT (resp. 7NT), the minimum number of tricks required to win the board is 9 (resp. 13).

The player on the left of the declarer exposes the first card of the game. The declarer's partner (called the *Dummy*) then lays their cards face up on the table.

When playing in a NT contract, there is only one simple rule : each player is required to follow suit if possible and can play any card of the suit.

When the four players have played a card, the player who played the highest-ranked card in the suit (2<3<...<10,J,Q,K,A) wins the trick and they will be on lead at the following trick.

The board is over when all the cards have been played.

Computer Bridge

Since 1996, the best bridge programs can annually participate in the World Computer-Bridge Championship (WCBC). The selected bridge AI used in our experiments was developed by Yves Costel for the bridge program Wbridge5¹. The boosted version of Wbridge5 (Ventos et al. 2017) won the WCBC in 2016, 2017 and 2018.

The best Bridge bots use PIMC, with a double-dummy solver (DDS) to evaluate each simulated world. They also use some additional heuristics, for example to prefer the less-revealing of equivalent actions.

DDS gives the number of tricks won by each side for Spade, Heart, Diamond, Club or No Trump contracts when all four players know the placement of the 52 cards, and each player plays optimally.

A tree-search can then be used in this simplified game, as well as standard alpha-beta pruning methods. Values of the leaves are computed using the double-dummy solver. The growing number of tricks won during the play is used as upper and lower bounds in the alpha-beta process.

A very efficient DDS has been written (and made public) by Bo Haglund (Haglund 2010).

Currently, the average level of the best current Bridge AIs are still far from the level of professional players, and closer to the level of good amateurs. This is partly due to the aforementioned problems of PIMC: namely strategy fusion and non-locality.

The strategy fusion problem of PIMC comes from the fact that DDS plays different Max moves in the different worlds of an information set. In the real game, Max has to play the same move in all the worlds of an information set. $\alpha\mu$ addresses this problem by evaluating every Max move jointly for all the possible worlds in the information set during its search.

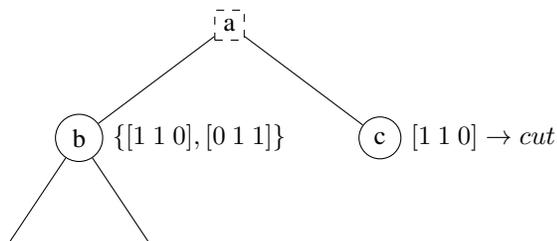
The second problem of PIMC is non-locality. It happens that a move which is optimal at a node is not the best move to backup from a global perspective. In some cases it is better to keep a locally suboptimal move that gives a better result at the root of the search tree than the locally optimal one. $\alpha\mu$ addresses this problem by backing up Pareto fronts instead of vectors.

Previous Work: the $\alpha\mu$ Algorithm

The $\alpha\mu$ algorithm repairs the strategy fusion and non-locality problems of PIMC by searching multiple moves ahead and by manipulating Pareto fronts at Max and Min

¹<http://www.wbridge5.com>

Figure 1: Example of an early cut at node c.



nodes. The $\alpha\mu$ algorithm assumes that the defence has perfect information whereas the declarer has incomplete information.

At Max nodes, each possible move returns a Pareto front. The final Pareto front for a node is the union of all the Pareto fronts returned by the search beginning with each of the different possible moves. The idea is to keep all the possible options for Max, i.e. Max has the choice between all the vectors of the overall Pareto front. In order to optimize computations and memory, vectors that are dominated by another vector in the same Pareto front are removed.

The Min players can choose different moves in different possible worlds, so they take the minimum outcome over all possible moves for a possible world. When they can choose between two vectors they take for each index the minimum between the two values at this index of the two vectors.

The $\alpha\mu$ algorithm is related to GIB single dummy solver (Ginsberg 2001). However the single dummy solver of Ginsberg is limited to the endgame and its complexity explodes rapidly for earlier states that $\alpha\mu$ handles easily at the price of incompleteness. Given enough time and using all possible worlds, $\alpha\mu$ converges to the single dummy solver results.

Generation of Possible Worlds

The possible worlds are generated using random generation followed by verification of the constraints. We currently use three types of constraints: the constraints coming from the bidding phase, the constraints on the West hand due to following rules for the opening lead, the constraints due to sluff.

Early Cut

If a Pareto front at a Min node is dominated by the Pareto front of the upper Max node it can safely be cut since the evaluation is optimistic for the Max player. A deeper search will always return a worse (or the same) result for the Max player due to strategy fusion.

Figure 1 gives an example of an early cut at a Min node. The root node a is a Max node, the first move played at a returned $\{[1\ 1\ 0], [0\ 1\ 1]\}$. The second move is then tried leading to node c and the initial Pareto front calculated with double dummy searches at node c is $[1\ 1\ 0]$. It is dominated by the Pareto front of node a so node c can be cut.

Root Cut

If a move at the root of $\alpha\mu$ for M Max moves gives the same probability of winning than the best move of the previous iteration of iterative deepening for $M - 1$ Max moves, the search can be safely be stopped since it is not possible to find a better move.

Optimizations

In this section we present different improvements of the $\alpha\mu$ algorithm that make it faster but do not change the result of a search.

Maintaining Useful Worlds

As we descend the tree it is possible to remove the worlds that are useless to evaluate, allowing us to run fewer DDS evaluations at the leaves, and enabling cuts.

A transposition table perfectly recalls the results of the previous search at nodes. For example if the DDS result stored in the transposition table for a world at a Min node is zero, it will also be zero in the whole subtree. It is therefore useless to calculate it again and the world can be marked as useless.

If at a Min node the maximum value of a world for the current Pareto front is zero, the world can be marked as useless as it will always have a zero value in the Pareto front returned by the node.

At the leaves only useful worlds are solved by DDS.

At Min nodes it is useless to unify the set of possible moves with the possible moves of the useless worlds. We thus only take the union of the possible moves in useful worlds.

World Cuts

Cuts due to having only zero or one useful world are called world cuts.

If the search is at a node without useful worlds it can be safely cut.

If there are no useful worlds left, search can be cut as the Pareto front is known to be the vector with only zeros.

Search can also be cut if there is only one useful world left. The reason for this is that all of the useless worlds will eventually evaluate to zero at the root, and so we only need to compute the DDS result associated to the single useful world and we can return a single vector containing the result for the useful world.

Figure 2 gives an example of a world cut with no useful worlds and figure 3 gives an example of a world cut with one useful world.

Calculating Vectors at Interior Nodes

It happens that some moves are not explored due to root cut for example. However when the evaluation of the best move at the root gets worse with more search, these moves are then explored to a depth greater than one. In order to verify if an early cut is possible at a depth smaller than the search depth for these moves, the algorithm has to calculate the Pareto fronts at smaller depths. If it can cut with the resulting front it has gained much search efficiency.

Figure 2: Example of a world cut with no useful worlds at node c .

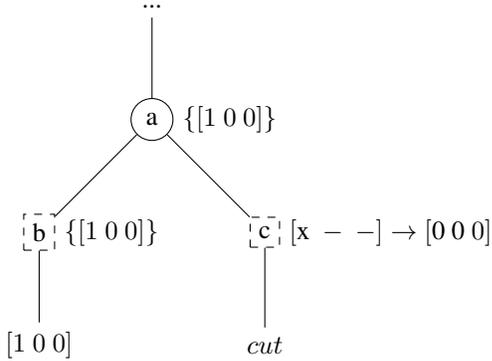
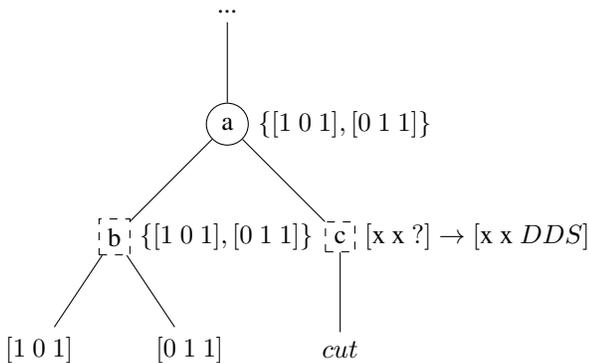


Figure 3: Example of a world cut with one useful world at node c .



We call this optimization the *Empty Entry* optimization.

Cut on Win

If at a Max node a move is found for which all worlds are won for the current search depth, the search can be cut, as more search cannot improve on this result.

α Cut

Figure 4 gives an example of an α cut at a Min node. The root node a is a Max node, the first move played at a returned $\{[1\ 1\ 0], [0\ 1\ 1]\}$. The second move is then tried leading to node c and the initial Pareto front calculated with double dummy searches at node c is $[1\ 1\ 1]$. The first Min move at c returns the $[x\ 1\ 0]$ Pareto front. The x means that the Min move is not possible in the first world and thus there is no evaluation associated to this world for the first Min move. The current Pareto front at node c is updated by taking the minimum in all the possible and evaluated worlds of the outcomes, leading to a $[1\ 1\ 0]$ Pareto front. This Pareto front is dominated by the Pareto front of the left move at node a and therefore the search can be cut.

Figure 5 gives an example of a deep α cut at a Min node. For each max node earlier in the path, if there exists one node whose Pareto front dominates currently evaluated front, further search can be avoided.

Algorithm 1 shows how to go up in the tree from a candidate node to check if the candidate node can be cut. The principle of the algorithm is to check all upper Max nodes to see if one of them dominates the current candidate node. We first save the current node to be evaluated against upper node (line 2), we then crawl the tree upward until reaching a max node (line 4 to 7). If we do not find any Max node it means that we reached the tree root and can return false, otherwise we test for Pareto dominance (line 11). If the upper max node does not dominate the candidate node, we search for the next Max node in the tree (line 3).

Algorithm 1 The deep α cut

```

1:  $\alpha$ cut(node)
2: candidate  $\leftarrow$  node
3: while node exist do
4:   node  $\leftarrow$  parent(node)
5:   while node exist and Min parent(node) do
6:     node  $\leftarrow$  parent(node)
7:   end while
8:   if node doesn't exist then
9:     return false
10:  end if
11:  if candidate.front  $\leq$  node.front then
12:    return true
13:  end if
14: end while
15: return false

```

The $\alpha\mu$ Algorithm with Cuts

Algorithm 2 gives the $\alpha\mu$ algorithm with cuts. M is the number of Max moves to search. If this number is equal to zero

Figure 4: Example of an α cut at node c.

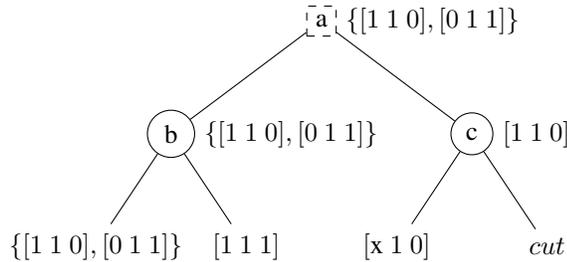
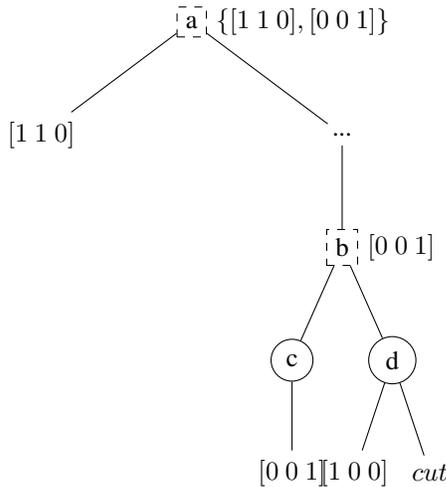


Figure 5: Example of a deep α cut at node d.



or if the state is terminal or if there is at most one useful world left the search is stopped (lines 2-5).

At a Min node if there is an early cut the search stops (lines 9-11). Otherwise the union of the sets of legal moves of the useful worlds is calculated (lines 12-16). All moves are tried, maintaining the possible worlds (line 21), making a recursive call (line 22), updating the Pareto front (line 23), updating the useful worlds (line 24) and making a cut if possible (lines 25-27).

At a Max node similar operations are performed, the root cut is also tested (lines 43-47).

Experimental Results

In the experiments we have $\alpha\mu$ play either the 3NT contract or the 7NT contract against PIMC with 20 simulated worlds at each decision point, or against WBridge5.

Search Times with and without Optimizations

We fix a number of cards for the initial state and we play a game from this state using $\alpha\mu$, recording the average time per move. The initial number of cards is either 32 or 52. We play 100 games, and thus 3200 searches for 32 cards and 5200 searches for 52 cards. We make experiments both for 20 and 40 simulated worlds at each decision point. In all experiments the early cut and the root cut are enabled.

Table 1 gives the average times used by the program with and without optimizations for deals with 52 cards, 20 worlds and three Max moves. Cards is the number of Cards for the starting hand of each game. M is the number of Max moves allowed during the search. Worlds is the number of worlds used by the search. U is for maintaining useful worlds. E is the use of the Empty Entry optimization. α is the use of the α cut. W is the use of the World cuts. Win is the use of Cut on Win.

Table 2 gives the times for deals with 32 cards, 20 worlds and three Max moves.

Table 3 gives the times for deals with 32 cards, 40 worlds and three Max moves. We see that the speedup factor carries through to the case with twice as many simulated worlds.

Games Against WBridge 5

In order to compare two algorithms we use duplicate games. It means that the two algorithms play the same 10 000 deals against the same algorithm for defense. Table 4 shows the results of duplicate games against different algorithms. The P1 column shows the Max player associated to the winrate. The P2 column shows the other player to which it is compared. The D column indicates the algorithm used for playing the defense. P1 and P2 are playing the declarer for a 7 No Trump contract. The \neq column gives the number of games where P1 has a different outcome than P2. The Winrate column shows the percentage of these different games that were won by P1 and lost by P2. The standard deviation is given in column σ . The players are $\alpha\mu$, WBridge5 (WB5) the winner of the 2016, 2017 and 2018 World Computer-Bridge Championships, Perfect Information Monte Carlo (PIMC), and the Double Dummy Solver (DDS) which has complete information and plays knowing the hands of all players.

Algorithm 2 The $\alpha\mu$ search algorithm with cuts.

```

1: Function  $\alpha\mu$  (state, M, Worlds,  $\alpha$ )
2:   if stop(state, M, Worlds, result) then
3:     update the transposition table
4:     return result
5:   end if
6:   t  $\leftarrow$  entry in the transposition table
7:   if Min node then
8:     mini  $\leftarrow$   $\emptyset$ 
9:     if t.front  $\leq$   $\alpha$  then
10:      return mini
11:    end if
12:    Worlds = updateUsefulWorlds(t.front, Worlds)
13:    allMoves  $\leftarrow$   $\emptyset$ 
14:    for w  $\in$  Worlds do
15:      l  $\leftarrow$  legalMoves (w)
16:      allMoves = allMoves  $\cup$  l
17:    end for
18:    move t.move in front of allMoves
19:    for move  $\in$  allMoves do
20:      s  $\leftarrow$  play (move, state)
21:      W1  $\leftarrow$  {w  $\in$  Worlds : move  $\in$  w}
22:      f  $\leftarrow$   $\alpha\mu$  (s, M, W1,  $\emptyset$ )
23:      mini  $\leftarrow$  min(mini, f)
24:      Worlds = updateUsefulWorlds(mini, Worlds)
25:      if mini  $\leq$  to an upper front then
26:        break
27:      end if
28:    end for
29:    update the transposition table
30:    return mini
31:  else
32:    front  $\leftarrow$   $\emptyset$ 
33:    for w  $\in$  Worlds do
34:      l  $\leftarrow$  legalMoves (w)
35:      allMoves = allMoves  $\cup$  l
36:    end for
37:    move t.move in front of allMoves
38:    for move  $\in$  allMoves do
39:      s  $\leftarrow$  play (move, state)
40:      W1  $\leftarrow$  {w  $\in$  Worlds : move  $\in$  w}
41:      f  $\leftarrow$   $\alpha\mu$  (s, M - 1, W1, front)
42:      front  $\leftarrow$  max(front, f)
43:      if root node then
44:        if  $\mu$ (front) =  $\mu$  of previous search then
45:          break
46:        end if
47:      end if
48:    end for
49:    update the transposition table
50:    return front
51:  end if

```

Table 1: Comparison of the average time per move of different configurations of $\alpha\mu$ with three Max moves and 20 worlds on deals with 52 cards.

Cards	M	Worlds	U	E	α	W	Win	Time
52	3	20	n	n	n	n	n	3.020
52	3	20	y	n	n	n	n	2.916
52	3	20	n	y	n	n	n	3.321
52	3	20	n	n	y	n	n	3.342
52	3	20	n	n	n	y	n	3.648
52	3	20	n	n	n	n	y	2.837
52	3	20	y	y	y	y	y	1.032
52	3	20	n	y	y	y	y	3.230
52	3	20	y	n	y	y	y	1.438
52	3	20	y	y	n	y	y	1.469
52	3	20	y	y	y	n	y	1.365
52	3	20	y	y	y	y	n	3.566

Table 2: Comparison of the average time per move of different configurations of $\alpha\mu$ with three Max moves and 20 worlds on deals with 32 cards.

Cards	M	Worlds	U	E	α	W	Win	Time
32	3	20	n	n	n	n	n	2.016
32	3	20	y	n	n	n	n	1.465
32	3	20	n	y	n	n	n	1.739
32	3	20	n	n	y	n	n	1.704
32	3	20	n	n	n	y	n	1.972
32	3	20	n	n	n	n	y	1.259
32	3	20	y	y	y	y	y	0.605
32	3	20	n	y	y	y	y	0.989
32	3	20	y	n	y	y	y	0.648
32	3	20	y	y	n	y	y	0.647
32	3	20	y	y	y	n	y	0.629
32	3	20	y	y	y	y	n	1.207

Table 3: Comparison of the average time per move of different configurations of $\alpha\mu$ with three Max moves and 40 worlds on deals with 32 cards.

Cards	M	Worlds	U	E	α	W	Win	Time
32	3	40	n	n	n	n	n	4.381
32	3	40	y	n	n	n	n	2.891
32	3	40	n	y	n	n	n	3.939
32	3	40	n	n	y	n	n	3.852
32	3	40	n	n	n	y	n	4.338
32	3	40	n	n	n	n	y	2.989
32	3	40	y	y	y	y	y	1.297
32	3	40	n	y	y	y	y	2.420
32	3	40	y	n	y	y	y	1.382
32	3	40	y	y	n	y	y	1.410
32	3	40	y	y	y	n	y	1.329
32	3	40	y	y	y	y	n	2.443

Table 4: Results of duplicate games against WBridge5 and PIMC on 10 000 deals at 7 No Trump.

P1	P2	D	\neq	Winrate	σ
$\alpha\mu$	WB5	PIMC	812	0.567	0.0174
$\alpha\mu$	WB5	WB5	755	0.428	0.0180
$\alpha\mu$	WB5	DDS	567	0.697	0.0193
$\alpha\mu$	PIMC	PIMC	757	0.551	0.0181
$\alpha\mu$	PIMC	WB5	687	0.569	0.0189
$\alpha\mu$	PIMC	DDS	467	0.647	0.0221

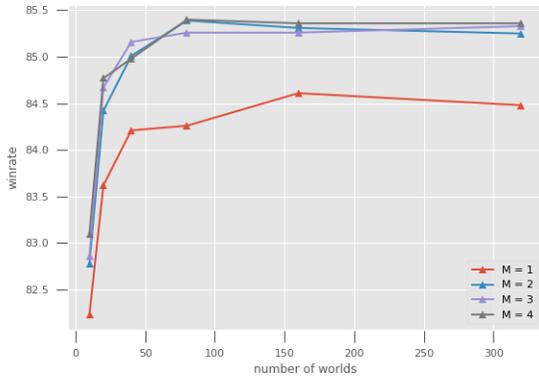


Figure 6: The evolution of the winrate with the number of worlds and the depth of the search.

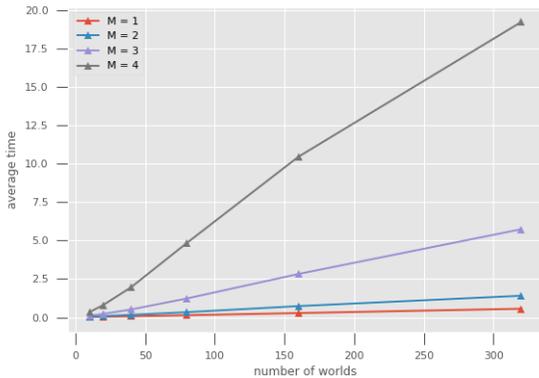


Figure 7: The evolution of the thinking time with the number of worlds and the depth of the search.

We can see on the first line that $\alpha\mu$ outperforms WBridge 5 against PIMC as the defense, but we see on the second line that WBridge5 performs better against itself playing as the defense. After discussing with the author of WBridge5, this may be due to the fact that WBridge5 as a declarer partly models itself as the defense during its search. The third line shows that $\alpha\mu$ performs much better than WBridge5 against a perfect complete information defense (DDS). The next three lines deal with PIMC as the second player of the duplicate games. $\alpha\mu$ performs better than PIMC against the three defenders, and especially against DDS (as in the duplicate with WBridge5).

Figure 6 shows the evolution of the winning percentage of different versions of $\alpha\mu$ against WBridge5. The x-axis is the number of worlds used by $\alpha\mu$. It starts at 20 worlds and doubles the number of worlds until 320 worlds. There are curves for each number of Max moves between 1 and 4. Note that $\alpha\mu$ with 1 max Move is PIMC. We can observe that the curves are asymptotic and that the asymptote of PIMC is below the asymptote of $\alpha\mu$ with numbers of Max moves greater than 1. Going from 2 Max moves to 3 or 4 does not improve the results.

Figure 7 gives the average time per move of the different versions of $\alpha\mu$. We can observe that $\alpha\mu$ with two Max moves is close to PIMC. However for equivalent thinking times for $\alpha\mu$ with two Max moves and for PIMC with 320 worlds, if we refer to figure 6, $\alpha\mu$ with two Max moves has better results than PIMC for equivalent times.

Leaf Parallelization

Previous experiments did not use parallelization. We now present a simple parallelization of the algorithms with cuts.

In Monte Carlo Tree Search there are three kinds of parallelization: root, leaf and tree parallelization (Cazenave and Jouandeau 2007; Chaslot, Winands, and van den Herik 2008; Cazenave and Jouandeau 2008). Root parallelization runs independent searches in parallel and sums the result to choose the most simulated move. Tree parallelization has multiple threads sharing a common tree. Leaf parallelization parallelizes the playouts at the leaves of the tree.

In order to speedup $\alpha\mu$, we are interested in the simplest form of parallelization, namely leaf parallelization. When reaching a leaf, $\alpha\mu$ does an $\alpha\beta$ complete information double dummy search on each possible world compatible with the cards played so far. These $\alpha\beta$ searches are independent and can safely be run in parallel. This is what we call the leaf parallelization of $\alpha\mu$.

In order to optimize Leaf Parallelization each thread is assigned the next available world as soon as the thread becomes available. There is a mutual exclusion on the assignment of a world to a thread.

Table 5 gives the comparison of the times spent per move for $\alpha\mu$ with and without parallelization for different numbers of Max moves and different depth. Leaf parallelization was done with OpenMP and mutual exclusion to choose the next available thread. We can see that there is little to gain for PIMC with 20 worlds, going from 0.019 seconds for 1 thread to 0.015 seconds for 6 threads. For PIMC with 80 worlds going from 1 to 6 threads halves the average time per

Table 5: Comparison of the average time per move with Leaf Parallelization using 1 thread and 6 threads.

Cards	M	Worlds	1 thread	6 threads
52	1	20	0.019	0.015
52	1	40	0.041	0.024
52	1	80	0.083	0.041
52	1	160	0.160	0.078
52	2	20	0.067	0.042
52	2	40	0.145	0.081
52	2	80	0.333	0.152

move. Similar speedups occur for $\alpha\mu$ with two Max moves: the speedup factor is approximately two for 80 worlds. Leaf parallelized $\alpha\mu$ with two Max moves and 40 worlds has approximately the same thinking time as PIMC with 160 worlds, and we can see in figure 6 that it has a better winrate.

Conclusion

We presented five different optimizations that speed up search for the $\alpha\mu$ algorithm, and showed how each optimization contributes to the final speedup when combined. In our experiments, the algorithm with the optimizations was shown to be three times faster than without the optimizations. The optimized algorithm has also been parallelized with leaf parallelization which gave a speedup factor of 2. The evolution of the winrate and the thinking time with the number of worlds has shown an asymptotic behavior both for PIMC and $\alpha\mu$ with different numbers of Max moves. The asymptote of $\alpha\mu$ is higher than the asymptote of PIMC, meaning that for usual thinking times and number of worlds, optimized $\alpha\mu$ performs better than PIMC. Moreover, optimized $\alpha\mu$ outperforms the former computer Bridge world champion WBridge5 in the 7NT contract when the defense is played by PIMC or DDS.

Acknowledgment

Thanks to Dan Braun for proofreading.

References

Bouzy, B.; Rimbaud, A.; and Ventos, V. 2020. Recursive Monte Carlo Search for Bridge Card Play. In *2020 IEEE Conference on Games (CoG)*, 1–8.

Brown, N.; and Sandholm, T. 2019. Superhuman AI for multiplayer poker. *Science* 365(6456): 885–890.

Buro, M.; Long, J. R.; Furtak, T.; and Sturtevant, N. 2009. Improving state evaluation, inference, and search in trick-based card games. In *Twenty-First International Joint Conference on Artificial Intelligence*.

Cazenave, T. 2009. Nested Monte-Carlo Search. In Boutilier, C., ed., *IJCAI*, 456–461.

Cazenave, T.; and Jouandeau, N. 2007. On the parallelization of UCT. In *proceedings of the Computer Games Workshop*, 93–101.

Cazenave, T.; and Jouandeau, N. 2008. A Parallel Monte-Carlo Tree Search Algorithm. In *Computers and Games, 6th International Conference, CG 2008, Beijing, China, September 29 - October 1, 2008. Proceedings*, 72–80.

Cazenave, T.; and Ventos, V. 2020. The $\alpha\mu$ Search Algorithm for the Game of Bridge. In *Monte Carlo Search at IJCAI*.

Chaslot, G. M.-B.; Winands, M. H.; and van den Herik, H. J. 2008. Parallel monte-carlo tree search. In *International Conference on Computers and Games*, 60–71. Springer.

Cowling, P. I.; Powley, E. J.; and Whitehouse, D. 2012. Information set monte carlo tree search. *IEEE Transactions on Computational Intelligence and AI in Games* 4(2): 120–143.

Edelkamp, S. 2020. Representing and Reducing Uncertainty for Enumerating the Belief Space to Improve Endgame Play in Skat. In *ECAI 2020*.

Frank, I.; and Basin, D. 2001. A theoretical and empirical investigation of search in imperfect information games. *Theoretical Computer Science* 252(1-2): 217–256.

Furtak, T.; and Buro, M. 2013. Recursive Monte Carlo search for imperfect information games. In *2013 IEEE Conference on Computational Intelligence in Games (CIG), Niagara Falls, ON, Canada, August 11-13, 2013*, 1–8.

Ginsberg, M. L. 2001. GIB: Imperfect Information in a Computationally Challenging Game. *J. Artif. Intell. Res.* 14: 303–358.

Haglund, B. 2010. Search algorithms for a bridge double dummy solver.

Kupferschmid, S.; and Helmert, M. 2006. A skat player based on Monte-Carlo simulation. In *International Conference on Computers and Games*, 135–147. Springer.

Lanctot, M.; Lockhart, E.; Lespiau, J.-B.; Zambaldi, V.; Upadhyay, S.; Pérolat, J.; Srinivasan, S.; Timbers, F.; Tuyls, K.; Omidshafiei, S.; et al. 2019. OpenSpiel: A framework for reinforcement learning in games. *arXiv preprint arXiv:1908.09453*.

Levy, D. N. 1989. The million pound bridge program. *Heuristic Programming in Artificial Intelligence The First Computer Olympiad* 95–103.

Lockhart, E.; Lanctot, M.; Pérolat, J.; Lespiau, J.-B.; Morrill, D.; Timbers, F.; and Tuyls, K. 2019. Computing Approximate Equilibria in Sequential Adversarial Games by Exploitability Descent. *arXiv preprint arXiv:1903.05614*.

Long, J. R.; Sturtevant, N. R.; Buro, M.; and Furtak, T. 2010. Understanding the success of perfect information monte carlo sampling in game tree search. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*.

Mahmood, Z.; Grant, A.; and Sharif, O. 2014. *Bridge for Beginners: A Complete Course*. Pavilion Books.

Sturtevant, N. R.; and White, A. M. 2006. Feature construction for reinforcement learning in hearts. In *International Conference on Computers and Games*, 122–134. Springer.

Ventos, V.; Costel, Y.; Teytaud, O.; and Thépaut Ventos, S. 2017. Boosting a Bridge Artificial Intelligence. In *Proc. International Conference on Tools with Artificial Intelligence (ICTAI)*, 1280–1287. IEEE.

Zinkevich, M.; Johanson, M.; Bowling, M.; and Piccione, C. 2008. Regret minimization in games with incomplete information. In *Advances in neural information processing systems*, 1729–1736.