SOLVING PROBABILISTIC COMBINATORIAL GAMES

Ling Zhao* and Martin Müller**

Department of Computing Science
University of Alberta
Edmonton, Alberta, Canada, T6E 2E8

Abstract. Probabilistic combinatorial games (PCG) are a model for Go-like games recently introduced by Ken Chen. They differ from normal combinatorial games since terminal positions in each subgame are evaluated by a probability distribution. The distribution expresses the uncertainty in the local evaluation. This paper focuses on the analysis and solution methods for a special case, 1-level binary PCG. Monte-Carlo analysis is used for move ordering in an exact solver that can compute the winning probability of a PCG efficiently. Monte-Carlo interior evaluation is used in a heuristic player. Experimental results show that both types of Monte-Carlo methods work very well in this problem.

1 Introduction

Heuristic position evaluation in game playing programs should be a measure of the probability of winning. However, in point-scoring games such as Go or Amazons, evaluation functions usually approximate the score of the game instead of the winning probability. This can lead to serious blunders when programs make a risky moves to gain even more territory, instead of playing conservatively to preserve a comfortable lead [8,7].

One approach to dealing with uncertainty of evaluation is to replace scalar evaluation by partially ordered values such as probability distributions. For the case of minimax search such approaches were developed in [2,9]. [7] extends this approach to the case of combinatorial sums of independent games [3], and defines the framework of *probabilistic combinatorial games* (*PCG*). Figure 1 shows an idealized example of how PCG might be applied to Go. A board is split into three independent areas, and the result of local search in each area is represented by a PCG. Terminal nodes in the local search are evaluated by a probability distribution.

[7] defined the PCG model and gave a high-level solution algorithm. This paper presents first computational results. The main contributions are:

– Efficient exact and heuristic solvers for the special case of sums of 1-level binary PCG's, called *SPCG*.
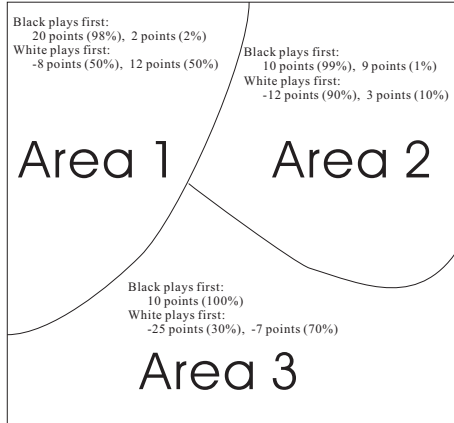
---

  * zhao@cs.ualberta.ca
 ** mmueller@cs.ualberta.ca

Black plays first:
   20 points (98%), 2 points (2%)
White plays first:
   -8 points (50%), 12 points (50%)

Black plays first:
   10 points (99%), 9 points (1%)
White plays first:
   -12 points (90%), 3 points (10%)

Area 1

Area 2

Black plays first:
   10 points (100%)
White plays first:
   -25 points (30%), -7 points (70%)

Area 3

Fig. 1. PCG example: sum of three subgames

- An analysis of the search space of SPCG.
- A Monte-Carlo move ordering technique used in alpha-beta search for the exact SPCG solver.
- A Monte-Carlo based heuristic evaluation technique for SPCG with performance close to that of the optimal player.
- An extensive experimental evaluation of the two solvers.

The remaining part of the paper is organized as follows. Section 2 gives definitions and notation for SPCG, and introduces the main problems addressed in this paper. Section 3 analyzes game trees in SPCG and describes methods to evaluate terminal and interior nodes. Section 4 develops the design of an exact solver and a strong heuristic player. Section 5 presents detailed experimental results and analysis, and future work is discussed in Section 6.

2 Definitions and Problem Identification

The following definition of probabilistic combinatorial games (PCG) is adapted from [7].

1. A terminal position, represented by a probability distribution $d = [(p_1, v_1), (p_2, v_2), \cdots, (p_n, v_n)]$, is a PCG. In d, the outcome is v_i with a probability of p_i, with $1 \leq i \leq n$, $0 \leq p_i \leq 1$ and $\sum_{i=1}^{n} p_i = 1$.
2. If $A_1, A_2, \cdots, A_n$ and $B_1, B_2, \cdots, B_n$ are all PCGs, then $\{A_1, A_2, \cdots, A_n \mid B_1, B_2, \cdots, B_n\}$ is a PCG. As in combinatorial games, the players are called Left and Right, $A_1, A_2, \cdots, A_n$ are the Left options and $B_1, B_2, \cdots, B_n$ are the Right options.
3. A sum of PCG is a PCG. Summation is understood in the sense of combinatorial game theory [3]: a move in a sum game consists of a move to an option in exactly one subgame and leaves all other subgames unchanged.

PCG are played as follows:

– A move can be played in any nonterminal subgame.
– If all subgames are terminal, the game itself is terminal, and its probability distribution is the sum of the distributions of all subgames.
– A game is won for Left if the value is greater than 0, and won for Right otherwise. From a terminal game, the probability of winning is computed by adding the probabilities of all values greater than 0.

The PCG model contains combinatorial games as the special case where all probability distributions in terminal positions have a single point distribution of the form $d = [(1, v)]$.

This paper focuses on a different special case of PCG which models uncertainty but keeps the combinatorial games as simple as possible. Simple PCG (SPCG) are PCG which obey the following constraints:

1. A SPCG consists of $n \geq 1$ subgames.
2. Each subgame has exactly one option for Left and one option for Right.
3. Each option immediately leads to a terminal position represented by a probability distribution.
4. Each distribution d in a terminal position has only two values with associated probabilities: $d = [(p_1, v_1), (p_2, v_2)]$.

A SPCG of n subgames is of the form $\sum_{i=1}^{n} \{d_i^L \mid d_i^R\}$, where d_i^L and d_i^R are 2-valued probability distributions.

In SPCG, a move can be simply indicated by the number of the subgame. In the experiments and analysis in this paper, without loss of generality, Left always moves first. This paper focuses on two basic problems concerning SPCGs:

1. Efficient exact solution: What is the winning probability of Left in a given SPCG?
2. Approximate solution: When solving a SPCG is too slow, how to play as well as possible under time constraints?

3 Game Tree Analysis of SPCG

Game trees in SPCG contain many regularities that can be exploited in an efficient solver. Suppose a game G has n subgames. There are n choices of a subgames for the first move, $n-1$ choices for the next move, and the kth move has $n-k$ choices. If the root of a game tree is defined as a node at depth 0, then any node at depth k has exactly $n-k$ children, and there are $n!/(n-k)!$ nodes in total at depth k. The total number of nodes in the game tree is $\sum_{k=0}^{n} n!/(n-k)! = \sum_{k=0}^{n} n!/k!$, and there are $n!$ terminal nodes.

Since the subgames are independent, the order of play of the subgames chosen by the same players does not change the evaluation. Each position can be represented by three sets: The *Left set (Right set)* contains the indices of all subgames played by Left (Right), and the *open set* contains the indices of all

subgames that have not been played yet. For example, the sequence of subgames chosen (1, 2, 3) is equivalent to the sequence of (3, 2, 1). If $n = 5$, then the Left set after both these sequences would be $\{1, 3\}$, the Right set $\{2\}$, and the open set $\{4, 5\}$.

For a position after k moves (a node at depth k in the game tree), k subgames have been played by the two players, $\lceil k/2 \rceil$ by Left and $\lfloor k/2 \rfloor$ by Right. The number of distinct nodes at depth k is equal to the number of possible combinations of left and right sets, $\binom{n}{k}\binom{k}{\lfloor k/2 \rfloor}$. There are $\binom{n}{\lfloor n/2 \rfloor}$ distinct terminal nodes, and the total number of distinct nodes in the game tree is $\sum_{k=0}^{n}\binom{n}{k}\binom{k}{\lfloor k/2 \rfloor}$.

As an example, Table 1 lists the number of nodes for SPCG with 12 to 15 subgames. The total number of distinct nodes in a game increases roughly by a factor of 3 per subgame. The number of distinct terminal nodes increases by a factor close to 2 (exactly 2 if n is odd, $2 - \frac{2}{n+2}$ if n is even). There is a very large number of transpositions in the game tree, and there are many more distinct interior nodes than terminal nodes.

Table 1. Statistics for SPCG with 12 to 15 subgames.

# subgames	12	13	14	15
# all nodes	1,302,061,345	16,926,797,486	236,975,164,805	3,554,627,472,076
# terminal nodes	479,001,600	6,227,020,800	87,178,291,200	1,307,674,368,000
# all distinct nodes	143,365	414,584	1,201,917	3,492,117
# distinct terminal nodes	924	1,716	3,432	6,435

3.1 Terminal Node Evaluation

In a SPCG with n subgames, the value of each terminal node T in its game tree is a sum of n probability distributions.

$$T = \sum_{i=1}^{n}[(p_{1,i}, v_{1,i}), (p_{2,i}, v_{2,i})]$$

T itself can be expressed as a single, complex probability distribution over sums of values $v_{1,i}$ and $v_{2,i}$. The winning probability of T for Left is the sum of all probabilities of values greater than 0 in T. The following formula expresses the winning probability P_w of T. Let $q(i) \in \{1, 2\}$ ($1 \leq i \leq n$) be such that the result value $v_{q(i),i}$ is chosen at subgame i. Then

$$P_w = \sum\{p_{q(1),1}p_{q(2),2}\cdots p_{q(n),n} \mid \sum_{i=1}^{n} v_{q(i),i} > 0\}, \forall 2^n \text{ choices of } q(1), q(2), \cdots, q(n).$$
$$(1)$$

There are at least three methods to evaluate P_w.

1. Direct evaluation of all 2^n combinations above.
2. A dynamic programming algorithm, representing the distribution as a list of bins and adding one distribution at a time. This is effective when the values v are all integers within a small range. However, in the worst case (real numbers or large integers), the final distribution can contain up to 2^n distinct non-zero entries, and this algorithm does not improve on the direct one.
3. Using the fact that the probability density of a sum is the convolution of the probability density functions of summands. Again, in the worst case the result will contain up to 2^n non-zero components.

In the experiments reported in this paper, for the range of n tested, methods 1 and 2 perform similarly well. Method 3 was not implemented. Method 2 can be further improved by noting that the final result required is only the value P_w, not the exact distribution. All intermediate results that are far enough from 0 to ensure a win or loss can be bagged and removed from further processing. This improves the efficiency. In experiments with $n = 14$ subgames, optimized dynamic programming was about twice as fast as the direct approach for a range of values $v \in [-1000, 1000]$ but twice as slow for a larger range $v \in [-5000, 5000]$. Since the dynamic programming version is efficient only when the range is limited, method 1 was used for all further experiments and discussion in this paper.

It is costly or even infeasible to compute the exact winning probability of a terminal node when n is large. In such a situation, it is desirable to have a good approximation method with controllable statistical error. In cases when a quick estimate is sufficient, approximation can improve the performance of a SPCG solver or player.

Monte-Carlo sampling is used to approximate winning probabilities. For each distribution d, a value is generated. For example, if $d = [(p_1, v_1), (p_2, v_2)]$, then v_1 is generated with probability p_1 and v_2 with probability $p_2 = 1 - p_1$. In a SPCG terminal node T, n values are generated from the n distributions of subgame outcomes, and the sum of these n values is the result of the sample. Left wins if and only if the sum is greater than 0. The fraction of wins P_w^k from k independent samples is used to approximate the winning probability in T.

For each sample, the winning probability is P_w, and $1 - P_w$ is the probability of a loss. The mean of the distribution is P_w, and the standard deviation is $\sqrt{P_w(1 - P_w)}$. According to the Central Limit Theorem, the mean of random samples drawn from a distribution tends to have a normal distribution. When k is sufficiently large, P_w^k has a normal distribution with mean P_w and standard deviation $\sqrt{\frac{P_w(1-P_w)}{k}} \leq \frac{1}{2\sqrt{k}}$. This inequality can be used to select the minimum value of k for a given required accuracy. For example, according to the normal distribution, almost for sure (99.7%) the difference between P_w and P_w^k is no more than $\frac{3}{2\sqrt{k}}$. With 10000 samples, the difference is 99.7% likely to be within 0.015.

Experiments shown in Table 2 list the difference between P_w and P_w^k in random terminal nodes, and compare it with the theoretical bounds. For games consisting of n subgames, 100 terminal nodes are randomly chosen as test cases. For each terminal node, probabilities were uniformly generated from between 0% and 100% with granularities of 1% and 0.001%. Values were uniformly generated between -1000 to 1000 with a granularity of 1.

For each terminal node, the error estimate, which is computed by $P_w - P_w^k$, is recorded for k from 3 to 33333. The standard deviation of the estimated error from the 100 terminal nodes is compared with its theoretical bound $(\frac{1}{2\sqrt{k}})$ in Table 2 for n is 14, 17, and 20 respectively.

Table 2. Standard deviation with probability granularity of 1% and 0.001%

# samples	3	10	33	100	333	1000	3333	10000	33333
$1/2\sqrt{k}$	0.2887	0.1581	0.0870	0.0500	0.0274	0.0158	0.0087	0.0050	0.0027
14 subgames	0.1768	0.1118	0.0488	0.0320	0.0189	0.0118	0.0070	0.0050	0.0042
17 subgames	0.2294	0.1176	0.0717	0.0348	0.0203	0.0121	0.0074	0.0059	0.0052
20 subgames	0.2427	0.1020	0.0580	0.0314	0.0246	0.0121	0.0077	0.0059	0.0045

# samples	3	10	33	100	333	1000	3333	10000	33333
$1/2\sqrt{k}$	0.2887	0.1581	0.0870	0.0500	0.0274	0.0158	0.0087	0.0050	0.0027
14 subgames	0.1893	0.0977	0.0655	0.0322	0.0190	0.0107	0.0053	0.0034	0.0019
17 subgames	0.1923	0.1093	0.0619	0.0419	0.0206	0.0125	0.0064	0.0038	0.0019
20 subgames	0.2091	0.1154	0.0634	0.0329	0.0199	0.0130	0.0064	0.0034	0.0019

The experimental results match the prediction very well when there is a fine granularity for randomly generated probability, but the model is not appropriate for coarse granularity when large sample sizes are used.

The term *Monte-Carlo terminal evaluation* (MCTE) is used to refer to the Monte-Carlo method discussed in this subsection. MCTE estimates the winning probability of a terminal node.

3.2 Monte-Carlo Sampling for Heuristic Interior Node Evaluation

Left's winning probability at an interior node, which is represented by its Left, Right and open sets, can be computed by a complete minimax search, using the exact evaluation of Section 3.1 in all terminal nodes of the search. However, when full search is too slow, Monte-Carlo sampling is useful for improving heuristic evaluation as well. Such methods are popular in games with incomplete information such as Poker [4], and also in games with complete information such as Go [6]. Abramson's expected-outcome evaluation [1] evaluates a node in a search tree by averaging the values of terminal nodes reached from it through random play. This method is adapted here. Similar ideas are developed in [10,2].

From an interior node, the sequence of alternate moves by both players to reach a terminal node is simulated by randomly choosing each move of the sequence among all its legal choices with equal probability. Such a simulation is iterated k times, and the average winning probability of the k sampled terminal nodes is an approximation of the winning probability of the interior node. The winning probability at terminal nodes can be either accurately computed by using Formula (1) in Section 3.1, or approximated by methods such as MCTE.

In SPCG, the order of moves chosen by the same player is irrelevant. For efficiency, a move sequence simulation can be replaced by uniformly randomly selecting k nodes out of all descendant terminal nodes below this interior node. Such sampling is only meaningful if k is (much) smaller than the total number of descendant terminal nodes of an interior node.

Experiments measure the difference between P_w and P'_w. For an interior node, its exact game value P_w and the approximate value P'_w estimated by the average winning probability of all its descendant terminal nodes are computed. The test set consists of starting positions from 100 randomly generated games with $n = 14$ subgames. The mean of $|P_w - P'_w|$ is 0.148 in the experiments. Since the difference is large, P_w can not be approximated well using P'_w. However, for a set of interior nodes at the same depth, their errors are highly correlated, and the relative order of their P'_w values is a very good approximation of the order of their P_w values. Section 5.4 gives detailed results. P'_w is an excellent move ordering heuristic.

In contrast to MCTE, the term *Monte-Carlo interior evaluation* (MCIE) is used to refer to the Monte-Carlo sampling discussed in this subsection. MCIE estimates the winning probability at an interior node. *Monte-Carlo move ordering* uses MCIE for move ordering at interior nodes in a search.

Monte-Carlo sampling is used in both the Monte-Carlo terminal evaluation and the interior evaluation, but they serve different purposes, and have different control parameters. These two types of evaluation are used in an exact SPCG solver and a heuristic SPCG player, which will be discussed in the next two sections.

4 Exact Solver and Heuristic Player for SPCG

In order to compute the exact winning probability of a game, a complete solver was implemented based on alpha-beta search. Standard enhancements including transposition tables and move ordering using the history heuristic [11] are used.

The optimized solver spends more than 90% of its time on the accurate evaluation of terminal nodes Since a SPCG game tree has far more interior than terminal nodes, most terminal nodes must be evaluated in order to solve the game. Unless a more efficient method can be found to evaluate terminal nodes, it is difficult to further improve the overall performance of the solver.

For cases when it is infeasible to solve a SPCG, or when it is desirable to play the game fast with reasonable strength, a heuristic player was designed as follows: A fixed depth alpha-beta search is performed. Nonterminal frontier

nodes of the search are evaluated using MCIE. Within MCIE, terminal nodes are evaluated by MCTE for efficiency.

A heuristic Monte-Carlo player with search depth of 1 is the same as the expected-outcome player in [1]. It performed well in experiments. This player chooses its move as follows: it finds all legal moves from the starting position, generates a depth 1 interior node for each move, and compares these nodes using MCIE. The move that leads to the highest-valued depth 1 node is chosen. This simple technique performs very well in SPCG. An advantage of this 1-ply Monte-Carlo sampling method is that it is friendly to time control. Each sampling process costs almost the same amount of time, and it can be performed continuously until time runs out. Experimental results for the solver and heuristic player are given in the next section.

5 Experimental Results and Analysis

This section summarizes experimental results for different configurations of the SPCG solver and the Monte-Carlo SPCG player, and measures how Monte-Carlo move ordering performs in games.

5.1 Experimental Setup

All experiments were run on Linux workstations with AMD 2400MHz CPUs. The compiler was gcc 3.4.2. The transposition table has 2^{20} entries, using about 34MB memory. Since the evaluation of terminal nodes is expensive, they can never be overwritten by interior nodes in the transposition table.

A set of 100 randomly generated games with $n = 14$ subgames are used for testing. Probabilities were uniformly generated between 0% and 100% with a granularity of 1% and values between -1000 to 1000 with a granularity of 1. Table 3 contains statistics for solving these instances. Each cell is of the form: mean value $\pm$ standard deviation. The first player has a big advantage in SPCG, so the average winning probability is much larger than 50%.

Table 3. Statistics of solving the 100 random games

| Time (sec)	Total nodes visited	Terminal nodes visited	Cache hits	Winning prob.
8.2 ± 0.8	$(2.2 \pm 1.2) \times 10^5$	$(2.5 \pm 0.2) \times 10^3$	$(1.3 \pm 0.8) \times 10^4$	$70.0\% \pm 25.8\%$

5.2 Performance of the Exact Solver

Monte-Carlo move ordering has two important parameters: the percentage of the total number of descendant terminal nodes to be sampled, n_t (see Section 3.2), and the number of value combinations to be sampled for approximate evaluation

of each terminal node, n_c (see Section 3.1). The Monte-Carlo move ordering uses MCIE, which contains two levels of Monte-Carlo sampling: on the top level, a number of an interior node's descendant terminal nodes are randomly chosen. On the bottom level, each of these terminal nodes is evaluated by the average result of a number of value combinations sampled from the value distributions of the node. Finally, the evaluation of the interior node is the average evaluation of those terminal nodes sampled.

The solver has a depth limit for Monte-Carlo move ordering (d_m). Nodes at depth up to d_m use MCIE for move ordering, and others use the history heuristic. The total time for solving the 100 test games for different parameter combinations is given in Figure 2.
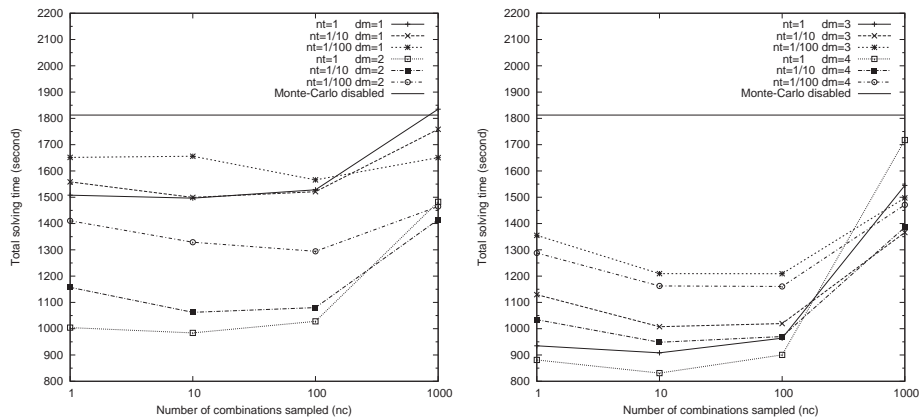
**Fig. 2.** Solving time with respect to parameters in Monte-Carlo move ordering

In the figure, the straight lines parallel to the x-axis denote the solving time when Monte-Carlo move ordering is disabled, and the solver uses only the history heuristic for move ordering. The large difference between these lines and the best solving time demonstrates that Monte-Carlo move ordering is superior to the history heuristic.

Even when only a small number of value combinations are sampled ($n_c =$ 1, 10, 100), the terminal node evaluation seems to provide a reasonably good estimate of the approximate value of interior nodes. $n_c = 1000$ is too slow to be practical.

In the range of n_c from 1 to 100, for each depth limit d_m the solving time almost always increases when n_t decreases. This indicates that the more descendant terminal nodes sampled, the better the performance will be. Thus it is best to sample all those terminal nodes whenever possible.

The depth limit for Monte-Carlo move ordering is an intricate parameter that needs a trade-off between search efficiency and overhead computation. Figure 3

shows the solver's performance for d_m from 0 to 14 for fixed $n_t = 1$ and $n_c = 10$. With $d_m = 0$ only the history heuristic is used for move ordering, and with $d_m = 14$ only Monte-Carlo move ordering is used.
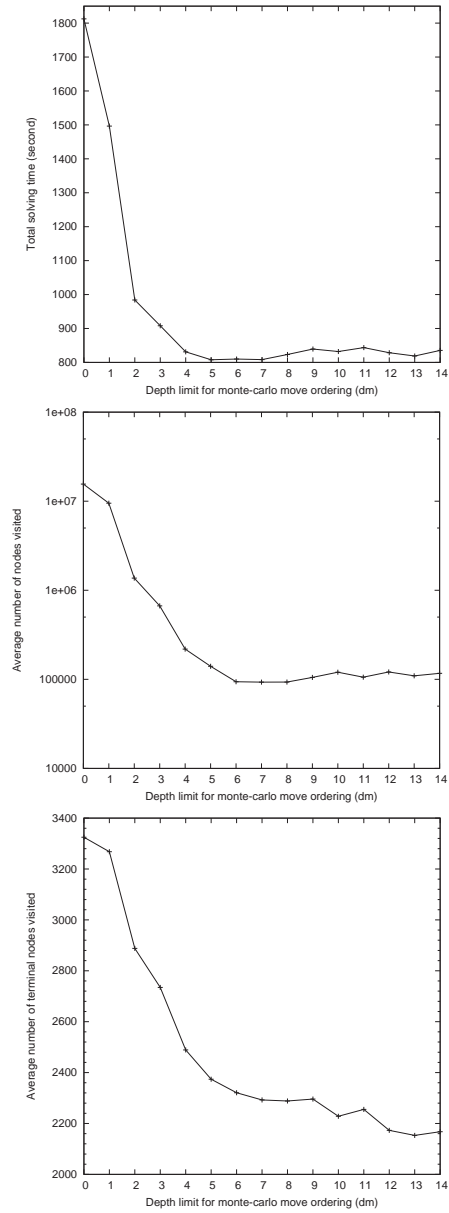

Fig. 3. Statistics of the solver for different values of d_m (n_t=1, n_c=10)

Figure 3 clearly shows that for $d_m \geq 4$, the solver's performance is relatively stable in terms of solving time and number of nodes visited. This fact suggests that the Monte-Carlo move ordering is more powerful for nodes close to the root of the game tree. It is also evident in the rightmost graph of Figure 3 that the number of terminal nodes gradually decreases as d_m increases, which proves the Monte-Carlo interior evaluation is more accurate than the history heuristic for move ordering in this game.

5.3 Performance of the Monte-Carlo Player

The relative strength of two players is tested by playing 100 random games twice, switching colors for the second game. A player's winning probability of a game is the average of its winning probabilities in the two rounds. The perfect player always achieves 50% winning probability against itself, and at least 50% against any player.

The Monte-Carlo player has two parameters: the search depth (at least 1), and n_c for the approximate evaluation of terminal nodes. It was tested with different configurations of these two parameters, and against the perfect player. Figure 4 illustrates the winning probability of the perfect player as well as the ratio of the time the perfect player spent to that of its Monte-Carlo counterpart. The parameter d_m is set to the search depth (meaning only MCIE is used), and $n_t = 1$. In the figure, $n_c = all$ means that all terminal nodes are evaluated exactly.
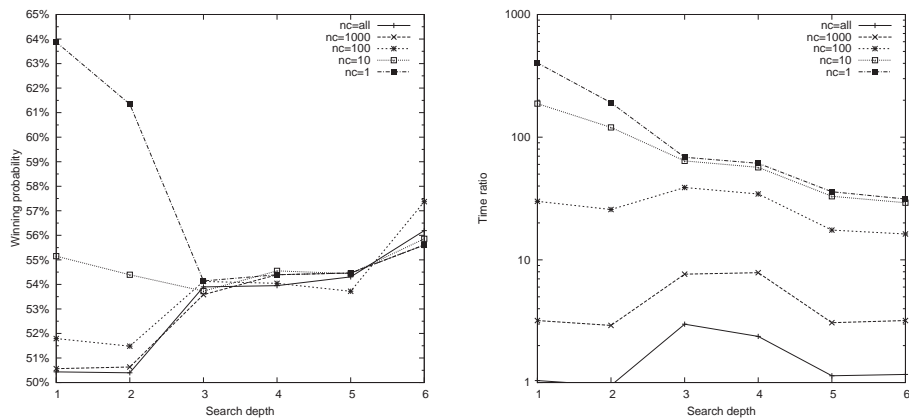
Fig. 4. The perfect player against the Monte-Carlo player: winning probability and time ratio

The left graph in Figure 4 reveals an interesting property of the Monte-Carlo player: increased search depth does not necessarily lead to better winning

probability, and sometimes even weakens the player. The Monte-Carlo player with accurate terminal node evaluation is best against the perfect player, but it uses almost the same amount of time, so it is not practical.

When the number of value combinations sampled, n_c, is from 100 to 1000 while the search depth is very shallow (1 or 2), the Monte-Carlo player performs very well. Especially when $n_c = 1000$ and the search depth is 1, the perfect player only achieves a winning probability of 50.56%. The Monte-Carlo player is just about 1% weaker.

The Monte-Carlo player with $n_c = 1000$ needs time comparable to the perfect player (about 25% to 30%), and is much slower than the player with $n_c = 100$. The latter player seems to offer a good compromise between time and accuracy: it plays games quickly, using 4% of the perfect player's time, and is only about 3% weaker.

It is not surprising that players with small n_c do not perform well. Insufficient sampling leads to large errors. The graph on the right side in Figure 4 suggests that a deeper search can make up for insufficient sampling to some extent.

### 5.4   Monte-Carlo Move Ordering

Since MCIE provides good estimates, it is important to provide a quantitative measurement to show how good it is with regard to move ordering.

As an experiment, one interior node at each depth is randomly selected from each of the 100 test games. For each node, all pairs of legal moves are compared to test if the order of their exact values is the same as the order suggested by MCIE. If it is different, then Monte-Carlo evaluation makes a mistake, and the winning probability difference of the two moves is recorded. This value denotes the winning probability lost due to the mistake. The *average probability error*, defined as the average of this value over all move pairs, measures the quality of move ordering for this node. The influence of move ordering on choosing the best move is measured by the *worst probability error*, the winning probability difference between the best move and the move suggested by MCIE.

Figure 5 illustrates the influence of n_c on the quality of Monte-Carlo move ordering. Data points with probability error of 0 are not shown due to the logarithmic scale used. Again, $n_c =$ all means that accurate terminal node evaluation is performed. The two graphs in Figure 5 clearly show that Monte-Carlo evaluation provides nearly perfect move ordering with marginal error if there are a substantial number of value combinations sampled. Even when n_c is small, 10 or 100, the worst probability error is still less than 2%. The error becomes larger when the node is closer to terminal nodes.

6   Conclusions and Future Work

This paper investigates SPCG, sums of 1-level probabilistic combinatorial games, and discusses methods to solve them, as well as strong heuristic players. Properties of this game are analyzed, and experimental results clearly show that

Fig. 5. Probability error due to incorrect move ordering

Monte-Carlo methods are useful for evaluation of both terminal and interior nodes.

Very interestingly, comparing the average winning probability of all descendant terminal nodes of an interior node is a good indicator of relative value as measured by winning probability. Experimental results show convincingly that this heuristic performs very well in move ordering for this game, and that is why it could bring a big performance increase to the complete solver and also why the Monte-Carlo player based on it performs well against the perfect player. It is interesting to investigate when and why Abramson's simple expected-outcome evaluation provides a good heuristic in games, and SPCG seems to be a good abstract model to study.

The SPCG solver and the Monte-Carlo player still have room for improvement. The bottleneck of the solver lies in the accurate evaluation of terminal nodes. An improvement on sampling strategy would be to incorporate progress pruning and selective search [4,6,5]. In the current Monte-Carlo player, each legal move is sampled at the same frequency. However, it is more efficient to use an adaptive strategy such that most of the effort is spent on those moves that have a high chance of being selected.

A SPCG solver could be incorporated in a Go program. The solver would be used on the high level to instruct the program to maximize its winning probability. As proposed in [7], such an approach might help to improve the playing strength of current Go programs. It would present a significant new application of Monte-Carlo methods.

References

1. B. Abramson. Expected-outcome: A general model of static evaluation. *IEEE transactions on PAMI*, 12:182–193, February 1990.
2. E. Baum and W. Smith. A bayesian approach to relevance in game-playing. *Artificial Intelligence*, 97(1-2):195–242, 1997.
3. E. R. Berlekamp, J. H. Conway, and R. K. Guy. *Winning Ways for your Mathematical Plays*. Academic Press, 1982.
4. D. Billings, A. Davidson, J. Schaeffer, and D. Szafron. The challenge of poker. *Artificial Intelligence*, 134(1-2):201–240, 2002.
5. B. Bouzy. Associating shallow and selective global tree search with Monte Carlo for 9x9 Go. In *4th Computer and Games Conference*, 2004.
6. B. Bouzy and B. Helmstetter. Monte Carlo Go developments. In H. J. van den Herik, H. Iida, and E. A. Heinz, editors, *Advances in Computer Games conference (ACG-10)*, pages 159–174, 2003.
7. K. Chen. Maximizing the chance of winning in searching Go game trees. *Information Sciences*, 2005. To Appear.
8. M. Müller. *Computer Go as a Sum of Local Games: An Application of Combinatorial Game Theory*. PhD thesis, ETH Zürich, 1995. Diss. ETH Nr. 11.006.
9. M. Müller. Partial order bounding: A new approach to evaluation in game tree search. *Artificial Intelligence*, 129(1-2):279–311, 2001.
10. A. Palay. *Search with Probabilities*. Morgan Kaufman, 1985.
11. J. Schaeffer. The history heuristic and the performance of alpha-beta enhancements. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(11):1203–1212, 1989.

Acknowledgements