Setting up a Computer Clobber Tournament at Tartu University

Jan Willemson[1]

Institute of Computer Science
Department of Mathematics and Informatics
Tartu University, Estonia

1. THE GAME OF CLOBBER

Clobber is a partizan game, created in 2001 by M.H. Albert, J.P. Grossman and R.J. Nowakowski (Albert *et al.*, 2004). It is played on an $m \times n$ rectangular board, where black and white stones are originally placed in a checkerboard fashion. A move consists of moving one's stone to an adjacent cell, provided it is occupied by an opponent's stone, which is then removed from the board ("clobbered"; see Figure 1). Black begins, the players move alternately and the one who can't move, loses.

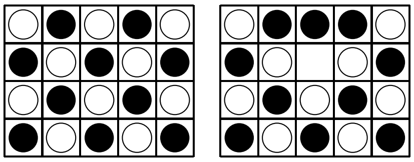

Figure 1: The first move in a game of Clobber

Thus, Clobber is a regular combinatorial game, and since both players have legal moves in every non-terminal position, it is an all-small game (Berlekamp, Conway, and Guy, 2001). Besides that, the game possesses several other properties which make it hard to analyze, despite the very simple rules. Namely, it is not easy to design a good position estimation function for Clobber, since

- in each position both of the players have almost the same number of stones, and
- in each position both of the players have exactly the same number of legal moves.

On the other hand, we see that Clobber positions tend to break down to independent subpositions, enabling classical Conway-style analysis (Conway, 2000).

Solitaire version of Clobber has been solved (Demaine, Demaine, and Fleischer, 2004), but only very little is known about the two-player version (Albert *et al.*, 2004). This makes the game interesting both for research and tournaments (Grossman, 2004; Wolfe, 2002).

Our own research conducted while setting up a Computer Clobber Tournament at Tartu University, Estonia, brought up several other interesting phenomena worth of further study. More specifically, we will see in the rest of the paper that

- the random player is surprisingly good even against some reasonable heuristics,
- the best currently known heuristics for the game try to form large groups of one's stones one way or another, still, these heuristics fail quite often.

[1] email:jan@ut.ee

Board size	Winning probability	Simulated winning probability
2×3	0	0
2×4	0.7637	0.78
3×3	0.8629	0.857
2×5	0.6039	0.586
2×6	0.3544	0.381
3×4	0.3385	$0,319$
4×4	0.5133	0.493
6×6		0.524
6×8		0.508
8×8		0.478

Table 1: Probability of the first player win under random play

2. SETTING UP A COMPUTER CLOBBER TOURNAMENT

The First International Clobber Tournament (with human players) was held at the Dagstuhl Seminar on Algorithmic Combinatorial Game Theory, February 17–22, 2002 (Grossman, 2004). Inspired by that event, a small Clobber programming tournament was held at Tartu University, Estonia, in spring term 2002. As that contest was a success, it was arranged again at the same place in spring 2004. This time it was a compulsory prerequisite to get the credit for the Game Theory course held on that term.

On one hand the obligation to write a program meant a pleasantly large participation – there were 34 submitted AI players. On the other hand, this popularity brought up unexpected concerns. Namely, from pedagogical point of view, it is only right to give credits for good programs that achieve certain goals, and not for something merely compliant with the rules.

Thus the first crucial question of the tournament was how to distinguish between a random player and a "good" player. Since there were many programs submitted, each of them consisting of hundreds of lines of messy student-quality code in different programming languages (Java, C, Pascal), just the code review was not an option. The technique used in Tartu University was to let each student program play against a truly random AI player and declare the program "good" if it would win often enough.

There are several catches hidden in this approach. First, in order for this verification methodology to work properly, we need to assume that if two random players play a game of Clobber, the outcome is 50% of winning chance for both players. This is not necessarily so. It is easy to construct instances of positions (especially small ones), where one of the players (black, white, beginner or the second) has 100% of winning chance under any legal play. We studied this matter further and conducted experiments on the winning probabilities for the first player on small boards having standard initial position with lower left corner being initially occupied by the beginning player.

Two kinds of experiments were carried out. First, for smaller boards (up to size 4×4), the whole game tree was generated and traversed, computing the precise probability. For larger boards, the tree was too big, and simulations-based approach was used (Abramson, 1990). 1000 random games were played on all (including smaller) boards and the results are displayed in Table 1.

The results show that simulations adequately approximate the known small board probabilities, giving some confidence in the results on larger boards as well. The latter ones vary closely around 50%, which shows applicability of our approach for distinguishing random players from "good" ones.

The second question to answer is how often should an AI player win the random one in order to be considered non-random oneself. For the tournament held in Tartu University, the threshold of 11 wins out of 15 games was used. Let us compute the probability that a random player that is expected to win every match with 50% chance, will win 11 or more games. This probability is equal to $\sum_{i=11}^{15} \binom{15}{i} \cdot (1/2)^i \cdot (1/2)^{15-i} \approx 0.059$. Thus, with confidence level of higher than 94% we can claim that a program that is able to win at least 11 games against the random player is statistically significantly better than the random one.

	O. Košik	E. Sutt	M. Kree	I. Lorvi	Sum	Place
Oleg Košik		2:0	2:0	1:1	5	I
Eneli Sutt	0:2		2:0	1:1	3	II
Mihkel Kree	0:2	0:2		2:0	2	III-IV
Imre Lorvi	1:1	1:1	0:2		2	III-IV

Table 2: Results of the final round of Clobber tournament

3. COMPUTER CLOBBER TOURNAMENT AT TARTU UNIVERSITY

The full tournament was arranged in four rounds. First, every submitted program was playing 15 rounds against the random player on 8×8 boards. All students whose programs won 11 or more matches were passed and got the credit. It is interesting to note that there was also one program that won four (!) games out of 15, i.e. behaved statistically significantly *worse* than the random player. This is an interesting phenomenon by itself. Let us compare random Clobber to famous game RoShamBo (Rock-Paper-Scissors) (Donkers, 2003). In RoShamBo no-one can beat the random player, but the random player can beat no-one either. In Clobber however, the random player seems to be surprisingly strong so that even quite reasonable heuristics turn out to fail against it. For example, the heuristic of always choosing the most top-left stone and selecting the move direction in predefined order of preference has resulted in weak AI players. However, this phenomenon has not been studied enough and such a study would be of great interest in the future.

At the end of the first round, there were 9 players who managed to win all 15 games against the random opponent. Those players were let further to the second round, 6 of them made it further to the semifinals. Two semifinals of three players were held, where each contestant had to play two games with different colors against all its opponents in the same semifinal. Two winners of each semifinal were allowed to enter the final round that was held by the same rules as the semifinals except for the used board size 6×8. Results of the final round are given in Table 2.

It is interesting to note that Oleg Košik's program won every game in every round except for its final game against Imre Lorvi's program (see Appendix A for the game record). As mentioned above, no good position estimation heuristics are known for Clobber. However, the top ones were created by the winners of the tournament, Oleg Košik, Eneli Sutt and Imre Lorvi.

Oleg Košik's and Eneli Sutt's heuristic is based on a very simple observation that it is favorable for a player to have a group of stones where there are many stones of his/her own and a few of the opponent's (Albert *et al.*, 2004). Thus their programs counted the number of neighboring pairs of friendly stones for both players and tried to maximize the difference between these counts. Imre Lorvi's heuristic had a somewhat similar motivation, but he tried to reduce the number of his groups, which would (hopefully) result in increasing the remaining ones.

Both approaches allowed the programs to create rather big groups, but there are weak points in such enlargement strategies. On one hand, if there will be no opposing stones adjacent to the formed group, this group becomes useless. On the other hand, large groups can often be easily separated, as happened in the game between Imre Lorvi's and Oleg Košik's programs (see Appendix A).

Thus Clobber can still be considered far from being solved and further research will definitely be very interesting.

4. REFERENCES

Abramson, B. (1990). Expected-outcome: a general model of static evaluation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 12, No. 2, pp. 182–193.

Albert, M., Grossman, J., Nowakowski, R., and Wolfe, D. (2004). An Introduction to Clobber. Submitted to INTEGERS.

Berlekamp, E. R., Conway, J. H., and Guy, R. K. (2001). *Winning Ways for Your Mathematical Plays*, Vol. 1. A. K. Peters.

Conway, J. H. (2000). *On Numbers and Games.* A. K. Peters.

Demaine, E. D., Demaine, M. L., and Fleischer, R. (2004). Solitaire Clobber. *Theoretical Computer Science*, Vol. 313, No. 3, pp. 325–338.

Donkers, J. (2003), Open Roshambo Competition Olympiad 2003. http://www.cs.unimaas.nl/~donkers/games/roshambo03/.

Grossman, J. P. (2004). Report on the first international Clobber tournament. *Theoretical Computer Science*, Vol. 313, No. 3, pp. 533–537. ISSN 0304–3975.

Siegel, A. (2004), Combinatorial Game Suite. http://cgsuite.sourceforge.net/.

Wolfe, D. (2002), Clobber research. http://homepages.gac.edu/~wolfe/games/clobber/.

5. APPENDICES

APPENDIX A: IMRE LORVI VS. OLEG KOŠIK

The game took place on June 17th 2004 in Tartu, Estonia.

1.	d2-d3	a2-b2
2.	f6-e6	a4-b4
3.	a7-b7	a6-b6
4.	e5-d5	c8-c7
5.	e1-e2	e4-d4
6.	f8-e8	f7-e7
7.	f2-f3	d1-c1

White's strategy is now apparent – he has managed to create a white group consisting of 14 stones of his own (maximum what one can do in 7 moves is 15 stones in a group).

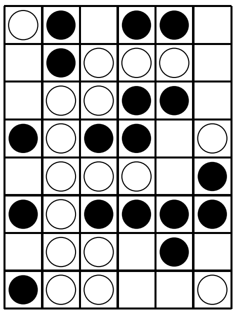

This might seem like a huge advantage, however black's strategy of reducing the number of his groups with the next moves a5-b5 and a3-b3 will (luckily for him) also break white's position. At the same time white tries to reconnect his a8 stone to the big group, actually removing one more little black group and thus unwillingly helping black.

8.	a5-b5	a8-b8
9.	a3-b3	b8-b7

White's large group has been divided into three parts and now black possesses the largest group of 7 stones.

10.	b5-b4	c4-c3
11.	d3-d4	…

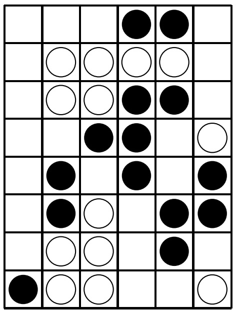

The value of the position can now be computed to be equal to $\Uparrow$ (e.g. with the help of Aaron Siegel's Combinatorial Game Suite, see (Siegel, 2004)) and white will not be able to avoid the inevitable. The key to black's success is the group on the right side of the board having the value $3\cdot\uparrow$, whereas the groups on top and in the lower-left corner have values 0 and $\downarrow$, respectively (see (Berlekamp *et al.*, 2001) for more detailed discussion on the position values).

11.	...	c3-b3
12.	b4-b3	c6-d6
13.	e6-d6	f5-f4
14.	d8-d7	e7-d7
15.	a1-b1	f4-f3
16.	b1-b2	f3-e3
17.	e2-e3	c2-b2
18.	b3-b2	d7-d6
19.	d5-d6	B+