

# 1 The maximum number of legal moves of 2 Othello in reachable positions is 33

3 Hiroki Takizawa<sup>1</sup>

4 <sup>1</sup>Preferred Networks, Inc., Chiyoda-ku, Tokyo, Japan

5 Corresponding author:

6 Hiroki Takizawa<sup>1</sup>

7 Email address: contact@hiroki-takizawa.name

## 8 ABSTRACT

9 We prove that the maximum number of moves among all reachable positions in a game of Othello is 33.  
10 We also construct a game record to achieve such a position. Additionally, including positions that are not  
11 reachable from the initial position, we prove that this maximum is 34 and construct such a position. We  
12 utilized Satisfiability (SAT) and Satisfiability Modulo Theories (SMT) solvers to obtain these computational  
13 proofs. These new findings not only clarify properties of the game of Othello, but also provide hints for  
14 the efficient and robust implementation of Othello software, and also present practical examples of using  
15 SAT and SMT solvers.

16 Keywords: Othello, Branching factor, SAT Solver, SMT Solver

## 17 INTRODUCTION

18 Othello is a well-known game, and the development of Othello software has continued for many years.  
19 One of the earliest superhuman-strength programs, Logistello, trained through self-play and overpowered  
20 the human world champion, Takeshi Murakami, in 1997 (Buro, 1999). Even after Othello software  
21 surpassed human performance, efforts to develop even stronger software have persistently progressed,  
22 with Edax serving as a prime example (Delorme, 2021).

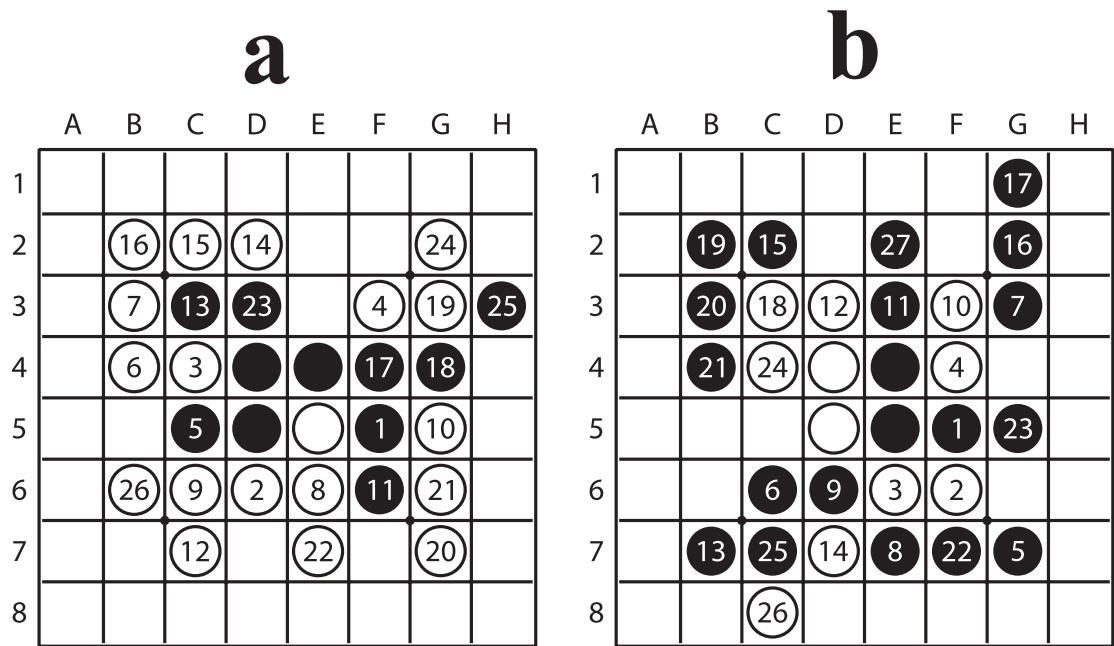
23 In order to enhance Othello software, it is essential to optimize search efficiency and achieve highly  
24 efficient implementation. In alpha-beta search (Abramson, 1989), move ordering is pivotal to the efficiency.  
25 Therefore, a technique is widely adopted that involves enumerating legal moves, quickly evaluating the  
26 potential of each move, and then conducting the search in descending order of each move's potential. If  
27 the number of legal moves, often referred to as the out-degree or branching factor in technical terms, is 32  
28 or fewer in any given position, the array to store them can be set to a length of 32. This can potentially  
29 speed up the software. However, until now, the maximum number of legal moves in Othello has been  
30 unknown.

31 In this paper, we prove that the maximum number of legal moves in all position of Othello that are  
32 reachable from the initial position is 33. We also prove that this number increases to 34 if we do not  
33 restrict ourselves to positions that are reachable from the initial position. We utilize Satisfiability (SAT)  
34 and Satisfiability Modulo Theories (SMT) solver to generate these proofs. Additionally, we reconstruct  
35 a game record from our proof, which transitions from the initial position to one with 33 legal moves  
36 shown in Figure 1. To the best of our knowledge, these discoveries and computational proofs are novel  
37 contributions presented in this study.

38 These findings not only offer the first clarification of one aspect of the nature of the game of Othello,  
39 but also represent successful practical applications of solvers for Satisfiability problems.

## 40 RELATED WORKS

41 The concept of computers playing board games can be traced back at least to Claude Shannon's con-  
42 siderations in 1950 (Shannon, 1950). Even during the 20th century, superhuman-strength software had  
43 emerged for many popular board games, including Othello (Buro, 1999), Chess (Campbell et al., 2002),



**Figure 1.** These are positions with 33 legal moves, which were outputted by the SAT solver kissat (Biere and Fleury, 2022), and are reachable from the initial position. (a) A position in which it’s Black’s turn, there are 33 legal moves, and there are 34 empty squares. The game record for this position is: “F5D6C4F3C5B4B3E6C6G5F6C7C3D2C2B2F4G4G3G7G6E7D3G2H3B6”. (b) A position in which it’s White’s turn and there are 33 legal moves with an equal number of empty squares. The game record for this position is: “F5F6E6F4G7C6G3E7D6F3E3D3B7D7C2G2G1C3B2B3B4F7G5C4C7C8E2”.

44 and Checkers (Schaeffer et al., 1996). These were based on alpha-beta search algorithm (Abramson,  
 45 1989), with their strength deriving from the ability to quickly search through a vast number of positions  
 46 via efficient implementation. In recent years, thanks to improvements in computer performance and  
 47 advancements of machine learning technology, software has achieved superhuman-strength in a wider  
 48 range of games, including Go (Silver et al., 2016) and Shogi (Japanese chess) (Kaneko and Takizawa,  
 49 2019; Silver et al., 2018).

## 50 METHODS

### 51 The use of terms ply and move

52 In chess, there is a convention whereby a pair of consecutive moves, one by white and one by black, is  
 53 referred to as a “move” or “full-move”, while a single move is termed a “ply” or “half-move”. However,  
 54 for clarity in this paper, we do not employ the term “ply” and consistently use the term “move” to denote  
 55 a single move.

### 56 Preliminary

57 Here, as a preliminary step, we will elucidate the methods for implementing several functions that form  
 58 the components of the Othello software, in a manner that is conducive to input into an SMT solver.

### 59 Bitboard

60 In this study, an arbitrary Othello position is represented by two 64-bit unsigned integers. Specifically,

- 61 • Each bit of the 64-bit integer is bijectively mapped to each square on the Othello board. The  
 62 least significant bit corresponds to A1, with the A row mapped in sequence to the lowest 8 bits,  
 63 and the most significant bit corresponds to H8. Generally, when  $i = 8j + k$  (where  $0 \leq i < 64$   
 64 and  $0 \leq j, k < 8$ ), the  $i$ -th bit corresponds to the square whose name is the concatenation of the  
 65  $(j + 1)$ -th letter of the alphabet with the number  $(k + 1)$ .

- In one of the 64-bit integers, the standing bits correspond to the squares where black stones are located. In the other 64-bit integer, the bits stand in the same manner, but correspond to the squares where white stones are located.

### 69 **Function to Generate Moves**

70 Assuming that the position is represented by the aforementioned two 64-bit integers, we designate  
 71 the bitboard of the side whose turn it is as “player”, and the bitboard of the opposing side as “oppo-  
 72 nent”. The “get\_moves” function shown in Figure 2 returns a bitboard wherein the set bits correspond  
 73 to all squares where a stone can be legally placed. This “get\_moves” function is implemented us-  
 74 ing only arithmetic operations, bit operations, and logical shifts. This code is available at GitHub:  
 75 <https://github.com/eukaryo/reversi33> .)

```

uint64_t get_some_moves(const uint64_t player, const uint64_t mask, const int dir)
{
    uint64_t flip_l, flip_r;
    uint64_t mask_l, mask_r;
    const int dir2 = dir + dir;

    flip_l = mask & (player << dir);    flip_r = mask & (player >> dir);
    flip_l |= mask & (flip_l << dir);    flip_r |= mask & (flip_r >> dir);
    mask_l = mask & (mask << dir);      mask_r = mask & (mask >> dir);
    flip_l |= mask_l & (flip_l << dir2); flip_r |= mask_r & (flip_r >> dir2);
    flip_l |= mask_l & (flip_l << dir2); flip_r |= mask_r & (flip_r >> dir2);

    return (flip_l << dir) | (flip_r >> dir);
}

uint64_t get_moves(const uint64_t player, const uint64_t opponent)
{
    const uint64_t mask = opponent & 0x7E7E7E7E7E7E7Eull;

    return (get_some_moves(player, mask, 1)
            | get_some_moves(player, opponent, 8)
            | get_some_moves(player, mask, 7)
            | get_some_moves(player, mask, 9))
            & ~(player | opponent);
}

```

**Figure 2.** A C++ implementation of a function to generate moves. This is based on Edax 4.4 (Delorme, 2021) with modifications by the author. Edax 4.4 is licenced under GNU General Public License v3.0 (Free Software Foundation, 2007).

### 76 **Function for Population Count**

77 A function that counts the number of set bits in a 64-bit integer can be implemented using only arithmetic  
 78 and bit operations, as exemplified by the “bit\_count” function in Figure 3. In Figure 3, 64-bit integer  
 79 multiplication is utilized. This choice is motivated by the fact that some recent CPUs have very low  
 80 latency for integer multiplication. This calculation has been known for a long time, and an implementation  
 81 without using integer multiplication is also feasible (Reingold et al., 1977). This code is also available at  
 82 GitHub: <https://github.com/eukaryo/reversi33> .)

### 83 **Function to Identify the Positions of Stones to Be Flipped**

84 The function that takes the bitboards representing the current player’s stones (where “player” denotes the  
 85 one whose turn it is) and the opponent’s stones, along with an integer specifying the move coordinates,  
 86 and generates the bitboard corresponding to the stones that will be flipped (hereafter referred to as the “flip  
 87 function”), is essential for Othello software. The efficiency of the flip function significantly influences  
 88 the overall computation speed of the search algorithm. In Edax, an array of flip functions is implemented,  
 89 allowing the appropriate function to be selected via compilation options. It is worth noting that all of the  
 90 flip functions implemented in Edax require index access to precomputed tables (except for the most naïve  
 91 implementation using slow for-loops), making them unsuitable for direct transcription into SMT solvers.

```

92 int bit_count(const uint64_t b)
93 {
94     uint64_t c = b
95     - ((b >> 1) & 0x77777777777777ULL)
96     - ((b >> 2) & 0x33333333333333ULL)
97     - ((b >> 3) & 0x11111111111111ULL);
98     c = ((c + (c >> 4)) & 0x0F0F0F0F0F0F0FULL) * 0x01010101010101ULL;
99     return (int)(c >> 56);
100 }

```

**Figure 3.** A C++ implementation of a function to population count. This is based on Edax 4.4 (Delorme, 2021) with modifications by the author. Edax 4.4 is licenced under GNU General Public License v3.0 (Free Software Foundation, 2007).

As shown in Figure 4, We found an algorithm in a separate GitHub repository (primenumber, 2016) that does not use index access to tables at all. This can be easily transcribed for use with SMT solvers. The code shown in Figure 4 is also available at GitHub: <https://github.com/eukaryo/reversi33> .)

### Definition of the Othello Variant

We defined the rules for the Othello variant as follows.

- The game does not end even in the case of two consecutive passes.
- The game ends after the 42nd move, including passes.
- The value of the game record equals to the maximum number of legal moves in all observed board positions, including the final position at game's end.
- Excluding the above stipulations, the rules align with those of traditional Othello.

Given the above definition, for  $n = 33$  and  $n = 34$ , we can assert that a game record with a value of  $n$  or greater exists if and only if there is a position that is reachable from the initial position and offers  $n$  legal moves. The reasons for this are as follows.

As a preliminary experiment, we conducted an exhaustive search during the opening phase and confirmed that up to 2 passes can occur by the time the number of stones on the board increases from 17 to 18. Consequently, the maximum number of passes that can occur while the number of stones on the board increases from 30 to 31 is 15.

In order to achieve 33 legal moves, the total number of stones must not exceed 31. Consequently, the value must reach 33 at the latest immediately following a pass after the placement of the 31st stone. As such, consideration of 42 moves (comprising 26 stone placements, a maximum of 15 interspersed passes, and a final pass) is sufficient. Similarly, for cases where there are 34 legal moves, consideration of 42 moves is sufficient.

## RESULTS

### Proven to be 34 Under Unrestricted Conditions

Initially, we disregarded the constraint regarding whether a position is reachable from the initial position. Assuming that the central four squares (namely D4, D5, E4, and E5) always contain stones, we proceeded to derive a proof for the maximum number of legal moves.

We used SMT solver z3 (De Moura and Bjørner, 2008) version 4.12.2 for our proof. In z3, it is possible to perform arithmetic operations on bitvectors, as well as bit manipulation and logical shifts. To input the problem into z3, the function to generate moves and the function to perform population count (described in Methods section above) were implemented using only the operations available in z3.

As a result, the maximum number of legal moves was determined to be 34. We proved that it is unsatisfiable to have 35 or more legal moves, regardless of the number of empty squares. This value can

```

125 _mm256i flipVertical(const _mm256i dbd) {
126     const _mm256i byteswap_table = _mm256_set_epi8(
127         8, 9, 10, 11, 12, 13, 14, 15, 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 0, 1, 2, 3, 4, 5, 6, 7
128     );
129     return _mm256_shuffle_epi8(dbd, byteswap_table);
130 }
131
132 _mm256i upperbit(_mm256i p) {
133     p = _mm256_or_si256(p, _mm256_srli_epi64(p, 1));
134     p = _mm256_or_si256(p, _mm256_srli_epi64(p, 2));
135     p = _mm256_or_si256(p, _mm256_srli_epi64(p, 4));
136     p = _mm256_andnot_si256(_mm256_srli_epi64(p, 1), p);
137     p = flipVertical(p);
138     p = _mm256_and_si256(p, _mm256_sub_epi64(_mm256_setzero_si256(), p));
139     return flipVertical(p);
140 }
141
142 uint64_t hor(const _mm256i lhs) {
143     _mm128i lhs_xz_yw = _mm_or_si128(_mm256_castsi256_si128(lhs), _mm256_extractf128_si256(lhs, 1));
144     return uint64_t(_mm_cvtsi128_si64(lhs_xz_yw) | _mm_extract_epi64(lhs_xz_yw, 1));
145 }
146
147 uint64_t flip(const uint64_t player, const uint64_t opponent, const int pos) {
148     const _mm256i pppp = _mm256_set1_epi64x(player);
149     const _mm256i oooo = _mm256_set1_epi64x(opponent);
150     const _mm256i OM = _mm256_and_si256(oooo, _mm256_set_epi64x(
151         0xFFFFFFFFFFFFFFFFULL, 0x7E7E7E7E7E7E7EULL, 0x7E7E7E7E7E7E7EULL, 0x7E7E7E7E7E7E7EULL));
152     _mm256i flipped, outflank, mask;
153     mask = _mm256_srli_epi64(_mm256_set_epi64x(
154         0x0080808080808080ULL, 0x7F00000000000000ULL, 0x0102040810204000ULL, 0x0040201008040201ULL), 63 - pos);
155     outflank = _mm256_and_si256(upperbit(_mm256_andnot_si256(OM, mask)), pppp);
156     flipped = _mm256_and_si256(_mm256_slli_epi64(_mm256_sub_epi64(_mm256_setzero_si256(), outflank), 1), mask);
157     mask = _mm256_slli_epi64(_mm256_set_epi64x(
158         0x0101010101010100ULL, 0x0000000000000000FEULL, 0x0002040810204080ULL, 0x8040201008040200ULL), pos);
159     outflank = _mm256_and_si256(_mm256_and_si256(mask, _mm256_add_epi64(_mm256_or_si256(
160         OM, _mm256_andnot_si256(mask, _mm256_set1_epi8(0xFF))), _mm256_set1_epi64x(1))), pppp);
161     flipped = _mm256_or_si256(flipped, _mm256_and_si256(_mm256_sub_epi64(outflank, _mm256_add_epi64(
162         _mm256_cmpeq_epi64(outflank, _mm256_setzero_si256()), _mm256_set1_epi64x(1))), mask));
163     return hor(flipped);
164 }

```

**Figure 4.** A C++ implementation of the flip function (using Intel intrinsics for AVX2). This is based on issen (primenumber, 2016) with modifications by the author. The GitHub repository is licenced under GNU General Public License v3.0 (Free Software Foundation, 2007).

125 be considered as the upper bound for the maximum number of legal moves that can be achieved from the  
126 initial position through a series of legal moves.

127 Another discovery was that among the positions with 33 legal moves, the maximum number of empty  
128 squares was 38 (for example, in the position shown in Figure 5a). For positions with 34 legal moves, the  
129 maximum was 37 (for example, in the position shown in Figure 5b). We proved that if a value exceeding  
130 these maximums is given as a constraint, it becomes unsatisfiable. (As declared in the data availability  
131 section below, the source codes and the outputs of analyses in this subsection are available at GitHub:  
132 <https://github.com/eukaryo/reversi33> .)


### 133 **The Generated Positions Were Found to Be Unreachable from the Initial Position**

134 Next, we investigated whether the generated positions could be legally reached from the initial position.  
135 We implemented a function that takes a position as input and enumerates all positions that could directly  
136 precede the given position, followed by conducting a depth-first search.

137 To enhance computational efficiency, we pre-enumerated all positions reachable from the initial  
138 position that contain 15 stones and have at least one legal move (19,785,690 variations when considering  
139 symmetric positions as identical), storing them in a hash table. When the aforementioned backward  
140 depth-first search arrives at a position with 15 stones, we can conclude that the original position is  
141 reachable from the initial position if and only if that position exists in the hash table.

142 It is worth noting that a pruning technique can be applied to the aforementioned backward depth-first  
143 search. If a position is reachable from the initial position, the graph formed by treating the stones as  
144 nodes and connecting stones that exist within a Moore neighborhood (i.e., differences in x-coordinate  
145 and y-coordinate are either one or zero) will always yield a connected graph. Therefore, during the  
146 aforementioned backward depth-first search, pruning can be performed immediately once such a graph  
147 becomes disconnected.

148 As a result, we determined that the generated positions were unreachable from the initial position. By  
149 iteratively adding a constraint that prevented the generation of identical positions, we produced over 1000



**Figure 5.** Positions with 33 or more legal moves that the SMT solver z3 outputted, without considering whether they are reachable from the initial position. Both of the following are positions with black to move. (a) A position with 33 legal moves and 38 empty squares. (b) A position with 34 legal moves and 37 empty squares.

150 positions, none of which were reachable from the initial position (data not shown).

### 151 Conversion of Problems to SAT Solver-Compatible Format and Their Resolution

152 Building on the results above, we sought to enable the SAT solver to simultaneously identify positions  
 153 reachable from the initial position that feature 33 or more legal moves, along with the corresponding  
 154 game records. As a preparatory step, we defined the rules for the Othello variant as described in Methods  
 155 section.

156 First, we implemented the entire aforementioned Othello variant for the z3 SMT solver. However, the  
 157 solution did not complete within 24 hours. Furthermore, to the best of our knowledge, z3 does not have a  
 158 feature that permits the export of problem input in a format compatible with an external SAT solver, nor  
 159 does it allow for the input of solutions provided by an external SAT solver.

160 Therefore, we developed software that, given a problem setting, converting it into Conjunctive Normal  
 161 Form (CNF) which is equivalent to the original problem in terms of satisfiability, and outputs it in  
 162 DIMACS CNF format. In addition, for the case that the solution was found to be satisfiable, we also  
 163 developed software to reconstruct the game record from the assignment of Boolean values to variables of  
 164 the original problem. This approach enabled us to solve the aforementioned Othello variant using the  
 165 latest SAT solvers.

166 We utilized the SAT solver kissat (Biere and Fleury, 2022), which demonstrated superior performance  
 167 in the SAT Competition 2022, to solve the problem. The solver features a “target” option. It accepts  
 168 the assumption whether the solution is satisfiable or unsatisfiable, and the solver will solve the problem  
 169 efficiently if the assumption is correct; even if the assumption is incorrect, it will still reach the correct  
 170 conclusion, albeit potentially taking more time than when not using this option. In this experiment, both  
 171 assumptions were tried in parallel on a Ryzen 5950X CPU.

172 As a result, we obtained a position where there were 33 legal moves (as shown in Figure 1b) and the  
 173 game record from the initial position. When we imposed a constraint that the number of legal moves must  
 174 be 34 or more, we found this constraint to be unsatisfiable. Additionally, we introduced an extra rule to  
 175 the aforementioned Othello variant that only considers the game record when it’s Black’s turn to move.  
 176 As a result, we obtained a position where there were 33 legal moves for Black (as shown in Figure 1a)  
 177 and the game record from the initial position. All three solutions were acquired within a 24-hour timeframe.

178 (As declared in the data availability section below, the source codes and the outputs of analyses in this  
179 subsection are available at GitHub: <https://github.com/eukaryo/reversi33> .)

## 180 **DISCUSSION**

181 In this study, we proved that among the positions reachable from the initial position in Othello, there  
182 exists a position with 33 legal moves, and that 33 is the maximum number of legal moves. Additionally,  
183 we also proved that, without considering whether the position is reachable from the initial position, the  
184 maximum number of legal moves increases to 34. This study not only elucidates of one aspect of Othello  
185 but also provides valuable information for the development of memory-safe and robust Othello software.

186 In the course of this study, we successfully constructed positions that fulfilled the specified conditions,  
187 thereby demonstrating that the maximum numbers act as a lower bound. Conversely, we verified that these  
188 maximum numbers also act as a tight upper bound by demonstrating that an increment by one results in  
189 the SAT solver returning “unsatisfiable”. To validate this proof, it is crucial to ensure the absence of bugs  
190 both in the code we developed to convert the Othello problem to CNF for the SAT solver and within the  
191 SAT solver itself. Although we carefully checked our code for bugs, the potential introduction of formal  
192 verification to heighten certainty may be considered. This represents a limitation of our study. However,  
193 the primary aim of our study was to determine the existence of a position that is reachable from the initial  
194 state and in which the number of legal moves exceeds 32. Therefore, this was deemed beyond the scope  
195 of our study.

## 196 **CONCLUSIONS**

197 This study presented a scenario where a problem that proved intractable with the SMT solver z3 was solved  
198 within a reasonable time frame by converting it to CNF format and utilizing the latest SAT solver kissat  
199 (Biere and Fleury, 2022). Given the rapid advancements in SAT solvers through annual competitions, this  
200 case appears to support a perspective that challenges the design of SMT solvers closely integrated with  
201 older SAT solvers. However, further research is warranted on this matter.

202 We hope this research proves beneficial to readers endeavoring to develop Othello software, understand  
203 the nature of Othello and other games, and those working towards the enhancement of SAT solvers and  
204 other theoretical solvers.

## 205 **ADDITIONAL INFORMATION AND DECLARATIONS**

### 206 **Competing Interests**

207 The author declares that there are no competing interests.

### 208 **Author Contributions**

209 Hiroki Takizawa conceived and designed the research, implemented and performed the computational  
210 experiments, analyzed the data, prepared figures, authored drafts of the paper, and approved the final draft.

### 211 **Data Availability**

212 The source codes and the outputs of analyses are available at GitHub: <https://github.com/eukaryo/reversi33>  
213 .

### 214 **Funding**

215 The author has received no funding for this study.

## 216 **REFERENCES**

- 217 Abramson, B. (1989). Control strategies for two-player games. *ACM Comput. Surv.*, 21(2):137–161.  
218 Biere, A. and Fleury, M. (2022). Gimsatul, IsaSAT and Kissat entering the SAT Competition 2022. In  
219 Balyo, T., Heule, M., Iser, M., Järvisalo, M., and Suda, M., editors, *Proc. of SAT Competition 2022 –  
220 Solver and Benchmark Descriptions*, volume B-2022-1 of *Department of Computer Science Series of  
221 Publications B*, pages 10–11. University of Helsinki.  
222 Buro, M. (1999). How machines have rearned to play othello. *IEEE intelligent systems & their applications*,  
223 14(6):12–14.



- 224 Campbell, M., Hoane, A., and hsiung Hsu, F. (2002). Deep blue. *Artificial Intelligence*, 134(1):57–83.
- 225 De Moura, L. and Bjørner, N. (2008). Z3: An efficient smt solver. In *International conference on Tools*  
226 *and Algorithms for the Construction and Analysis of Systems*, pages 337–340. Springer.
- 227 Delorme, R. (2021). edax-reversi. <https://github.com/abulmo/edax-reversi>. Last retrieved 2023-07-07.
- 228 Free Software Foundation, I. (2007). Gnu general public license, version 3.  
229 <http://www.gnu.org/licenses/gpl.html>. Last retrieved 2020-01-01.
- 230 Kaneko, T. and Takizawa, T. (2019). Computer shogi tournaments and techniques. *IEEE Transactions on*  
231 *Games*, 11(3):267–274.
- 232 primenumber (2016). issen/src/move\_generator.cpp. [https://github.com/primenumber/issen/blob/](https://github.com/primenumber/issen/blob/72f450256878094ffe90b75f8674599e6869c238/src/move_generator.cpp)  
233 [72f450256878094ffe90b75f8674599e6869c238/src/move\\_generator.cpp](https://github.com/primenumber/issen/blob/72f450256878094ffe90b75f8674599e6869c238/src/move_generator.cpp). Last retrieved 2023-07-07.
- 234 Reingold, E. M., Nievergelt, J., and Deo, N. (1977). *Combinatorial algorithms: theory and practice*.  
235 Prentice Hall College Div.
- 236 Schaeffer, J., Lake, R., Lu, P., and Bryant, M. (1996). Chinook the world man-machine checkers champion.  
237 *AI magazine*, 17(1):21–21.
- 238 Shannon, C. E. (1950). Xxii. programming a computer for playing chess. *The London, Edinburgh, and*  
239 *Dublin Philosophical Magazine and Journal of Science*, 41(314):256–275.
- 240 Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J.,  
241 Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner,  
242 N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., and Hassabis, D. (2016).  
243 Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489.
- 244 Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran,  
245 D., Graepel, T., Lillicrap, T., Simonyan, K., and Hassabis, D. (2018). A general reinforcement learning  
246 algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144.