ICGA Journal Vol. 33 No. 3 Sept 2010 pp 181-186 *JAIST Computer Olympiad, 2010 Kanazawa.*

MOHEX WINS HEX TOURNAMENT

Broderick Arneson, Ryan Hayward, Philip Henderson[1]

Edmonton, Canada

1. THE TOURNAMENT

Four programs competed in the 2010 Hex competition: MIMHEX by Jakub Pawlewicz, Lukasz Lew, and 11 students from Poland, YOPT by Abdallah Saffidine and Tristan Cazenave from France; WOLVE by Broderick Arneson, Ryan Hayward, and Philip Henderson from Canada; and MOHEX by Philip Henderson, Broderick Arneson, and Ryan Hayward. YOPT, MOHEX, and WOLVE competed in previous Olympiads. A fifth program, BITaHex from China, registered but was unable to participate due to visa problems related to the recent dispute between Japan and China.

MIMHEX is a Monte Carlo tree search program developed as part of a course taught by Pawlewicz and Lew on AI and Games in the Faculty of Mathematics, Informatics, and Mechanics at the University of Warsaw in Poland in 2009/10. Each assignment corresponded to a key part of the final program. After each assignment, a tournament was run to determine the best version of the program, which was then used as the code base for the next assignment. Each design decision is based on extensive testing.

MIMHEX uses the RAVE UCT formula (Gelly and Silver, 2007). In the rollout (game simulation) policy, the tournment version of MIMHEX uses only the bridge pattern; due to last-minute technical difficulties, a stronger version — with machine-learned local 7-cell (each board position and its 6 neighbors) patterns — was not used.

MIMHEX has no opening book. It spends about 10% of the remaining time on each move. This results in strong opening play — perhaps the strongest in this tournament — but also leads to occasional endgame difficulties. MIMHEX makes 50K rollouts per second. Thus an early move, which takes about 150 seconds, uses about 7500K rollouts.

YOPT is a Monte Carlo tree search program that uses the RAVE UCT formula, the bridge and 432 patterns in rollouts, some dead cell analysis, and 90K rollouts per move. YOPT performs a brief H-search at the start of each move; if no winning move is found, the tree search execution begins. A non-Olympiad version of YOPT computes virtual connection information in the UCT tree (Cazenave and Saffidine, 2009).

WOLVE, the 2009 silver medallist (Arneson, Hayward, and Henderson, 2009a)[2], uses a 4-ply truncated-width alpha-beta search, a Shannon style electric circuit evaluation function (with cell adjacencies augmented by virtual connections), and significant pruning of inferior cells. On average, one 4-ply move takes under 1 minute; the variance is large. To save time, WOLVE uses a book built by caching 6-ply moves. Once outside the book, usually after 2 or 3 moves, WOLVE reverts to 4-ply search.

MOHEX, the 2009 gold medallist (Arneson *et al.*, 2009a), is a Monte Carlo tree search program built on the code base of FUEGO, the Go program developed by Martin Müller, Markus Enzenberger, and others at the University of Alberta. FUEGO uses lock-free parallelization (Enzenberger and Müller, 2009), and backs up virtual losses for better parallelization. MOHEX performs virtual connection computation and inferior cell analysis in the UCT tree, and is only slightly modified since 2009, except that this year it ran on 16-cores instead of 8, yielding about 2000K rollouts per move.

MOHEX and WOLVE share many features, including SOLVER, a single-threaded depth-first proof number search solver that runs in parallel with the rest of the program, and produces perfect play whenever it solves the position

[1]Department of Computing Science, University of Alberta, Canada. Email:hayward@ualberta.ca

[2]See also previous games competition reports (Melis and Hayward, 2003; Willemson and Björnsson, 2004; Hayward, 2006; Arneson, Hayward, and Henderson, 2008; Arneson, Hayward, and Henderson, 2009b).

within the time allocated for a move. YOPT, MIMHEX, WOLVE, and MOHEX used 1, 1, 2, and 16 threads respectively.

Each player opened twice against each opponent. The tournament started on Saturday September 25 and finished on Tuesday September 28.

	MOHEX	WOLVE	MIMHEX	YOPT	total	result
MOHEX		3-1	3-1	4-0	10-2	gold
WOLVE	1-3		4-0	3-1	8-4	silver
MIMHEX	1-3	0-4		4-0	5-7	bronze
YOPT	0-4	1-3	0-4		1-11	4th

2. THE GAMES

The games are shown at the end of this article.

To obtain data for the following analysis, we ran SOLVER on positions from each game. For each game, we found the earliest position that SOLVER could solve in the allotted time, which varied depending on our interest from two hours to more than one day.

At each node in its search tree, SOLVER performs virtual-connection and inferior cell analysis. The latter allows many inferior moves to be pruned from the search; the former allows the computation of a *mustplay* region: any move outside this region is a loss for the player to move next. SOLVER has the same virtual-connection and inferior cell engines as MOHEX and WOLVE.

Round 1. Double round robin, Games 1–12.

Game 1. MIMHEX-MOHEX 0-1. White winning 23+ (from move 23 on).
Move 31.B[e3] by MIMHEX is outside SOLVER's mustplay, so SOLVER knows that this move loses as soon as the initial virtual connection computations finish. This move is technically not a blunder, as all other moves lose; however, some of these other moves offer more resistance, taking SOLVER about 400 seconds to refute.

Game 2. MIMHEX-WOLVE 0-1. White winning: 29+.
Move 29.B[i3] by MIMHEX is weak, allowing 30.W[h6], which captures all of {g6, i4, i5, i6}, rendering the previous move essentially useless; h3 offers more resistance to SOLVER than i3.

Game 3. MIMHEX-YOPT 1-0. White winning: 22+.
Move 32.B[a5] is *vulnerable*, that is, can be killed by the opponent's next move, and 34.B[h4] is outside the mustplay.

Game 4. MOHEX-MIMHEX 0-1. Black winning: 18. White winning: 19+.
Move 19.B[d7] by MOHEX is technically a blunder, as j2 is winning. However, it takes many hours for SOLVER to confirm that j2 wins.

Game 5. MOHEX-WOLVE 0-1. Black winning: 25+.

Game 6. MOHEX-YOPT 1-0. Black winning: 20+.
Move 36.W[f3] by YOPT is outside the mustplay, but the only move inside the mustplay is f11, and it loses.

Game 7. WOLVE-MIMHEX 1-0. Black winning: 21–35. White winning: 36+.
Move 36.B[h8] is a blunder by MIMHEX; it is outside the mustplay, and SOLVER needs only 25 seconds to show that b3 wins.

Game 8. WOLVE-MOHEX 0-1. White winning: 19+.
Move 19.B[g7] offers SOLVER less resistance than c4; it seems that WOLVE defends the wrong side of the board.

Game 9. WOLVE-YOPT 1-0. Black winning: 20+.
Move 22.W[a10] by YOPT is outside the mustplay, allowing WOLVE an easy win. However, SOLVER takes about 1.5 hours to solve the previous position, which is a Black win, so YOPT had no winning move to choose from.

Game 10. YOPT-MIMHEX 0-1. Black winning: 19+.
Move 27.W[f6] by YOPT is vulnerable, and 31.W[c10] by YOPT is outside the mustplay.

Game 11. YOPT-MOHEX 0-1. White winning: 15+.

Game 12. YOPT-WOLVE 0-1. Black winning: 16–22, 24. White winning: 23, 25+.
Move 23.B[i2] is technically a blunder by YOPT, as f6 wins, although it takes SOLVER about 10.5 minutes to show this. 24.W[g5] is a blunder by WOLVE, as i4 wins; it takes SOLVER 23 seconds to show this. 25.B[i6] by YOPT is also a blunder; it is outside the mustplay, and the two moves inside (j4 and i5) are winning.

Round 2. Double round robin, Games 13-24.

Game 13. MIMHEX-MOHEX 0-1. Black winning: 22–28. White winning: 29+.
Move 29.B[c11] by MIMHEX is a blunder: it takes SOLVER less than 2 seconds to show that f8 wins.

Game 14. MIMHEX-WOLVE 0-1. Black winning: 28. White winning: 29+.
Move 29.B[d7] by MIMHEX is technically a blunder, as c6 wins; however, it takes SOLVER several hours to show this.

Game 15. MIMHEX-YOPT 1-0. Black winning: 22+.
Move 14.W[h8] by YOPT is vulnerable, and is killed by 15.B[i7]. This exchange is bad for White, as it helps Black connect more easily to the top left edge.

Game 16. MOHEX-MIMHEX 1-0. Black winning: 20+.
Move 20.W[h4] by MIMHEX seems weak, offering SOLVER far less resistance than g6.

Game 17. MOHEX-WOLVE 1-0. Black winning: 15. White winning: 16+.
Move 16.B[d5] by WOLVE is technically a blunder, as h6 wins; however, it takes SOLVER 20 minutes to show this.

Game 18. MOHEX-YOPT 1-0. Black winning: 34–46, 48+. White winning: 33, 47.
Move 48.W[d3] by YOPT is a blunder, as i8 wins; it takes SOLVER 5 minutes to show this. The previous move 47.B[d4] by MOHEX is also a blunder, as k7 is winning; it takes SOLVER 10 minutes to show this. 34.W[f3] by YOPT is also a blunder, as g2 is winning.

Game 19. WOLVE-MIMHEX 1-0. Black winning: 24+.

Game 20. WOLVE-MOHEX 0-1. Black winning: 21+.

Game 21. WOLVE-YOPT 0-1. Black winning: 23+.

Game 22. YOPT-MIMHEX 0-1. Black winning: 33+.
Move 37.W[b4] by YOPT is outside the mustplay. It takes more than 10 hours to solve the previous position.

Game 23. YOPT-MOHEX 0-1. Black winning: 15+.
Move 17.W[f3] by YOPT is outside the mustplay; d5 offers more resistance.

Game 24. YOPT-WOLVE 0-1. Black winning: 23+.
Move 27.W[c5] by YOPT is outside the mustplay.

3. OBSERVATIONS

SOLVER found several errors in play by YOPT (8 moves outside mustplay, 3 vulnerable moves, 4 blunders), and a few by MIMHEX (2 outside mustplay, 3 blunders), WOLVE (2 blunders), and MOHEX (2 blunders). Many of the moves made out of the mustplay had no winning alternative; however, any such move leads to a quick loss if the opponent has a virtual connection engine, as do WOLVE and MOHEX.

YOPT and WOLVE had the quickest losses (after only 15 moves, twice by YOPT and once by WOLVE); all three losses were against MOHEX.

YOPT played in the two toughest games to solve, as measured by the number of stones required before SOLVER found the winner: MOHEX-YOPT-2 (34 stones) and YOPT-MIMHEX-2 (33-stones).

MIMHEX seems to have had the strongest opening play; for example, YOPT had winning positions against

MOHEX in one game, and against WOLVE in two games, but never — as far as SOLVER could tell — against MIMHEX.

With the exception of its two obvious blunders — one against MOHEX, one against WOLVE — MIMHEX played solidly, and seems close to challenging the recent Olympiad domination of WOLVE and MOHEX.

4. CONCLUSIONS

Monte Carlo Tree Search programs now dominate in the number of entries to the Hex competition. It appears they might soon dominate in strength as well.

MIMHEX's performance as a new entry this year was surprising strong, especially since its strength seems to derive only from Monte Carlo tree search. MIMHEX and YOPT might benefit by adding or strengthening virtual connection computation, as each made moves outside the mustplay. Adding vulnerable cell pruning and other forms of inferior cell analysis might also help.

We look forward to seeing BITaHex, and hopefully more new entries, at next year's competition.

Acknowledgements. We thank NSERC, iCORE, Martin Müller, Jonathan Schaeffer, and the UofA GAMES group for financial support; Lorna Stewart and Joe Culberson for loaning computers; Martin Müller for sharing the Fuego code base; and the many others who have worked on previous versions of WOLVE, MOHEX, and SOLVER.

5. REFERENCES

Arneson, B., Hayward, R. B., and Henderson, P. (2008). Wolve Wins Hex Tournament. *ICGA*, Vol. 31, No. 4.

Arneson, B., Hayward, R. B., and Henderson, P. (2009a). MoHex Wins Hex Tournament. *ICGA*, Vol. 32, No. 2, pp. 114–116.

Arneson, B., Hayward, R. B., and Henderson, P. (2009b). Wolve 2008 Wins Hex Tournament. *ICGA*, Vol. 32, No. 1, pp. 49–53.

Cazenave, T. and Saffidine, A. (2009). Utilisation de la recherche arborescente Monte-Carlo au Hex. *Revue d'Intelligence Artificielle*, Vol. 23, Nos. 2–3, pp. 183–202.

Enzenberger, M. and Müller, M. (2009). A Lock-free Multithreaded Monte-Carlo Tree Search. *Advances in Computer Games 12* (eds. H. J. van den Herik and P. Spronck), Vol. 666 of *Lecture Notes in Computer Science*, pp. 1–12. Springer.

Gelly, S. and Silver, D. (2007). Combining online and offline knowledge in UCT. *ICML '07: Proc. 24th international conference on machine learning*, pp. 273–280.

Hayward, R. B. (2006). Six Wins Hex Tournament. *ICGA*, Vol. 29, No. 3, pp. 163–165.

Melis, G. and Hayward, R. (2003). Six Wins Hex Tournament. *ICGA*, Vol. 26, No. 4, pp. 277–280.

Willemson, J. and Björnsson, Y. (2004). Six Wins Hex Tournament. *ICGA*, Vol. 27, No. 3, p. 180.

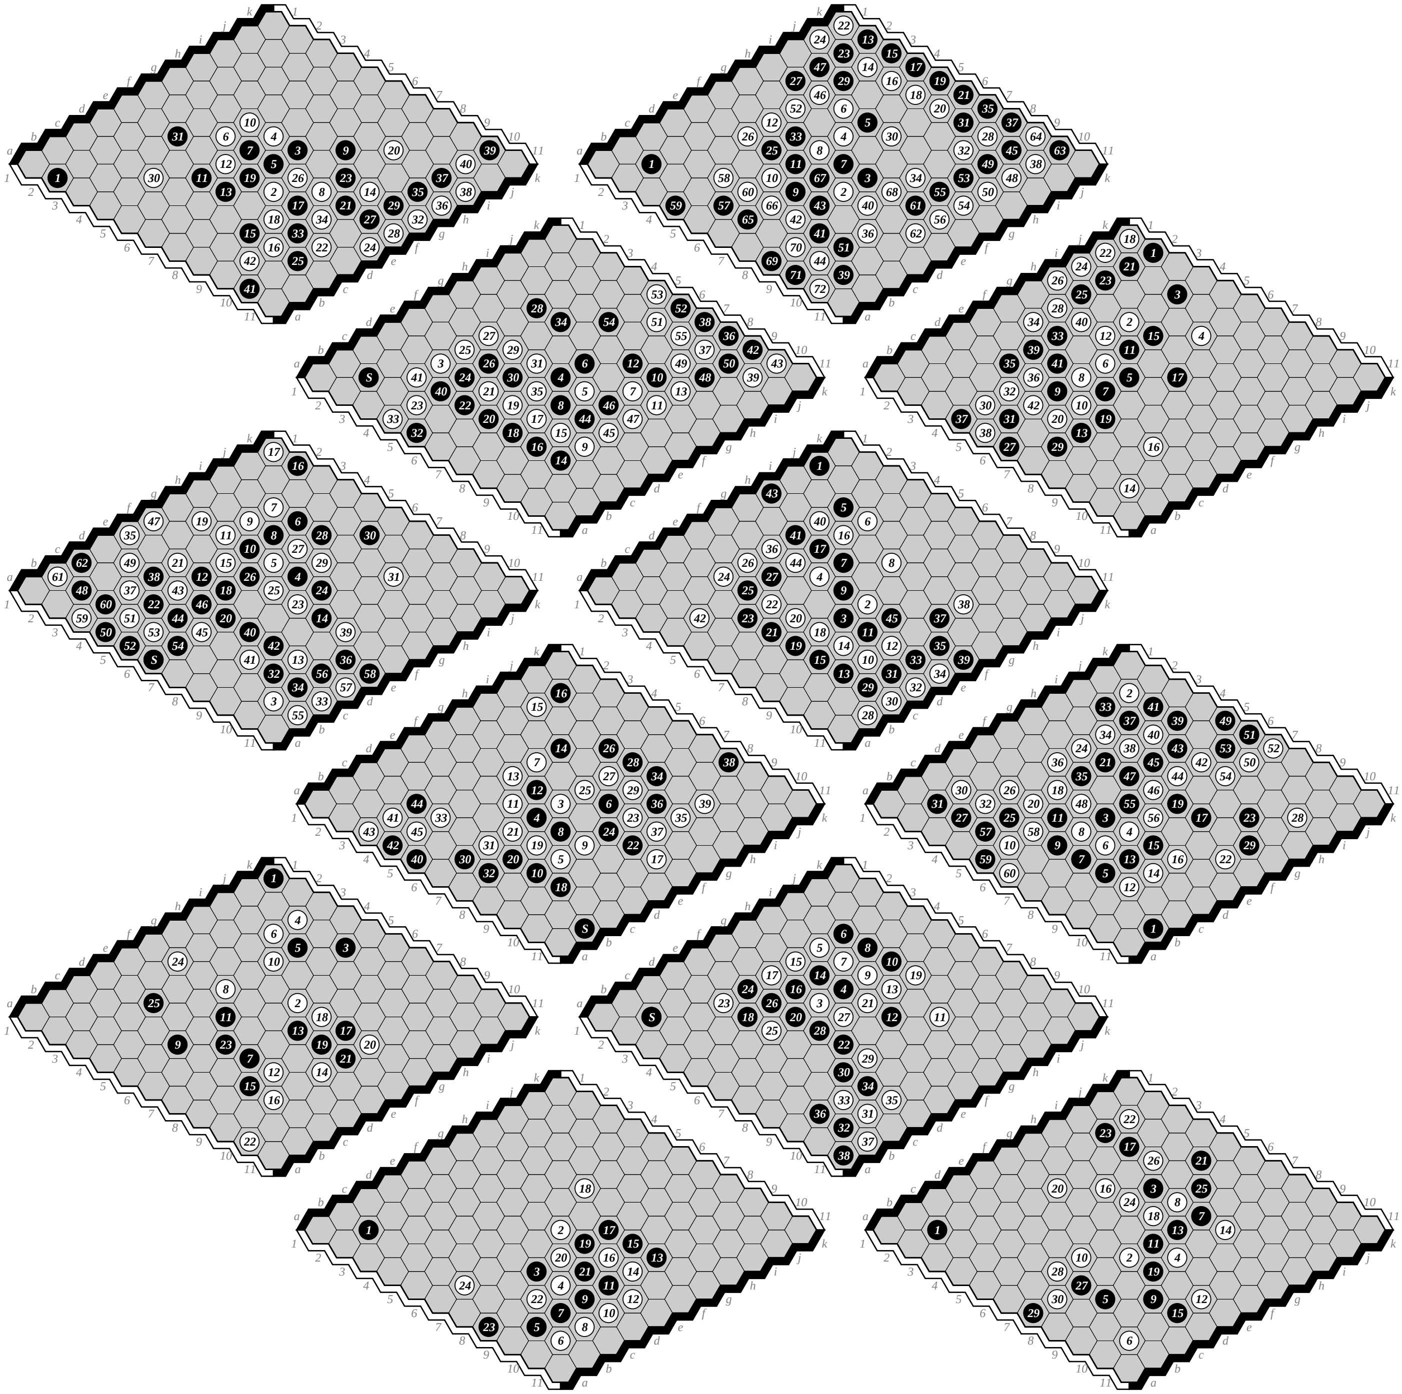

Round 1.
MIMHEX-MOHEX 0-1, MIMHEX-WOLVE 0-0 (top row).
MIMHEX-YOPT 1-0, MOHEX-MIMHEX 0-1 (2nd row).
MOHEX-WOLVE 0-1, MOHEX-YOPT 1-0 (3rd row).
WOLVE-MIMHEX 1-0, WOLVE-MOHEX 0-1 (4th row).
WOLVE-YOPT 1-0, YOPT-MIMHEX 0-1 (5th row).
YOPT-MOHEX 0-1, YOPT-WOLVE 0-1 (bottom row).
In some games the operator of the losing program resigned once the outcome was obvious.

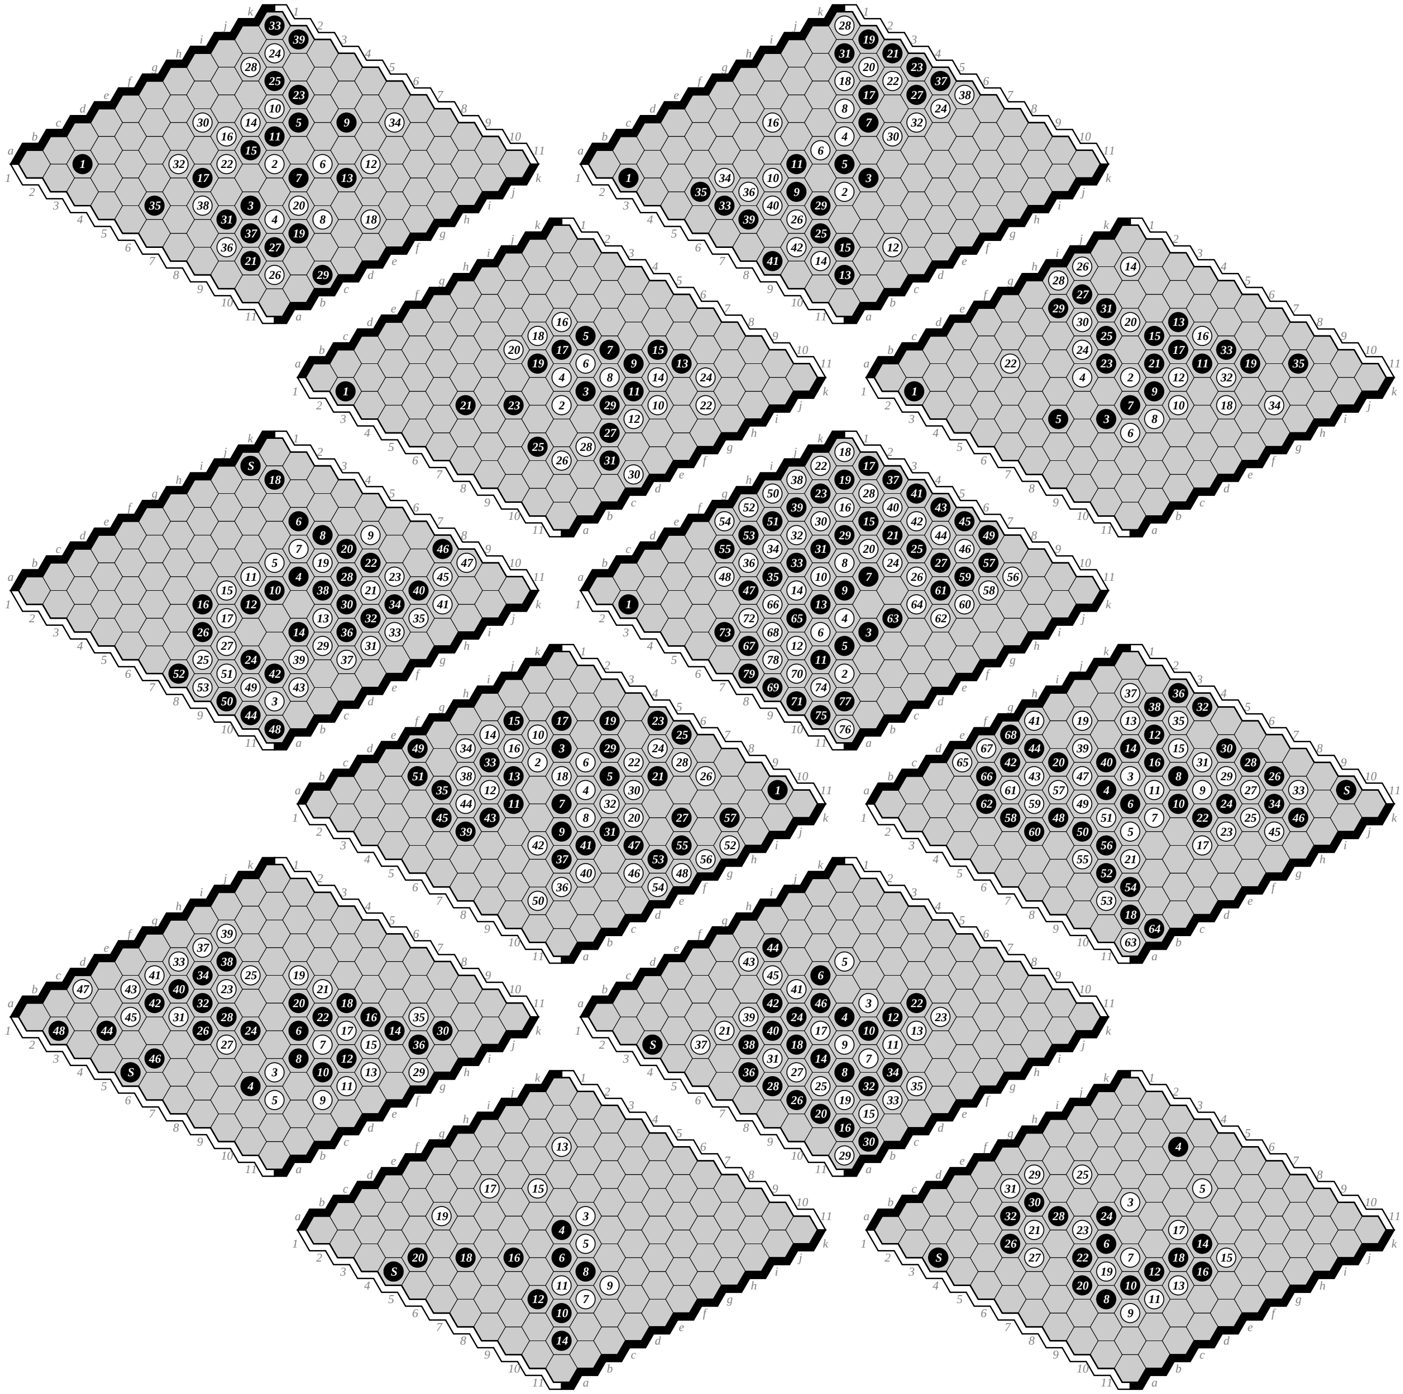

Round 2.
MIMHEX-MOHEX 0-1, MIMHEX-WOLVE 0-1 (top row).
MIMHEX-YOPT 1-0, MOHEX-MIMHEX 1-0 (2nd row).
MOHEX-WOLVE 1-0, MOHEX-YOPT 1-0 (3rd row).
WOLVE-MIMHEX 1-0, WOLVE-MOHEX 0-1 (4th row).
WOLVE-YOPT 0-1, YOPT-MIMHEX 0-1 (5th row).
YOPT-MOHEX 0-1, YOPT-WOLVE 0-1 (bottom row).