

Wolve Wins Hex Tournament

Broderick Arneson, Ryan Hayward, and Philip Henderson¹

Department of Computing Science, University of Alberta, Canada

Four programs competed in the 2008 Hex competition: Yopt by Abdallah Saffidine and Tristan Cazenave of France; Six (version 0.5.3) by Gábor Melis of Hungary; Wolve by Broderick Arneson, Ryan Hayward, and Philip Henderson of Canada; and MoHex by Philip Henderson, Broderick Arneson, and Ryan Hayward.

Six, the gold medallist in each Hex competition since 2003, uses a two-ply truncated-width alpha-beta search and a Shannon style electric circuit evaluation function in which cell adjacencies are augmented by virtual connections. The virtual connection engine (VCE) uses Anshelevich and/or closure operations. From a list of moves to consider, Six prunes dead cells with low degree; it also prunes cells outside of the virtual connection mustplay region, which are provably losing.

Wolve, the successor to 2006 and 2003 silver medallist Wolve and Mongoose programs² is similar to Six. One difference is that Wolve uses a new inferior³ cell engine (ICE) together with iterative deepening to prune more inferior moves. It also uses an opening book that was generated by self-play and then hand-modified on a few lines.

Yopt is a new UCT Monte Carlo program. It uses two patterns in rollouts, the bridge and the 432, and the all-moves-as-first (RAVE) heuristic. Yopt is the only of the four programs with no VCE.

MoHex is a new UCT Monte Carlo program. MoHex is built on the code base of Fuego, the Go program developed by Martin Mueller, Markus Enzenberger, and others at the University of Alberta. MoHex uses one pattern in rollouts, the bridge, and the all-moves-as-first heuristic. Its VCE is strengthened by the addition of pattern templates and a new combining rule. At the UCT-tree root node and at depth 1, MoHex prunes inferior cells via ICE and the VCE mustplay; rollouts occur on the pruned boards.

MoHex and Wolve share many features. However, several theoretical enhancements which ought to strengthen Hex play proved detrimental when added to one or both programs. For this reason, MoHex has stronger ICE and VCE than Wolve, but does not use Wolve's book.

Gábor Melis could not attend the tournament, so Six was operated by Nathan Sturtevant. All programs⁴ ran on an Intel Core 2 quad Q9559 LGA775 (2.66GHz/1333FSB/12MB) desktop with 2GB RAM.

The tournament had two rounds; the first completed on Wednesday October 1 and the second on Friday October 3. In each round, each player opened once against each opponent.

	Yopt	Six	MoHex	Wolve	total	result
Yopt		2-2	1-3	0-4	3-9	4th
Six	2-2		2-2	0-4	4-8	bronze
MoHex	3-1	2-2		3-1	8-4	silver
Wolve	4-0	4-0	1-3		9-3	gold

Round 1. Game 1:Six-MoHex 1-0 starts 1.Black[i9] 2.swap 3.White[f6] 4.Black[h5] ... Play looks reasonable for both players until 20.B[f8]?; this is perhaps a blunder, as the top left region seems more important here than the bottom right. For example, post-game analysis shows that 20.B[f3] 21.W[f2] 22.B[e3] 23.W[e2] 24.B[c3] would be winning⁵ for MoHex. Instead, Six is winning by 23.W[i3].

¹email:hayward@cs.ualberta.ca

²See previous games competition reports (Melis and Hayward, 2003; Willemson and Björnsson, 2004; Hayward, 2006).

³See (Björnsson *et al.*, 2007) for an introduction to inferior cell analysis.

⁴With the following exception: for operator convenience, on some first round games Six ran on a 2GHz Intel Core Duo with 2GB RAM. Six never comes close to using all its time, so this had no adverse effect on its performance.

⁵By *winning*, we mean in a known winning state, as determined by an automated Hex solver. Our solver uses an ICE and a VCE similar

Game 2:MoHex-Six 1-0 starts 1.B[a4] 2.W[f6] 3.B[d7] 4.W[e5] ... Move 4 looks weak, albeit to the eyes of a human regularly beaten by these programs, as it does not seem to block Black's development on the left side; for example, the Wolve self-play book would play 4.W[d8]. In any event, MoHex is winning by 27.B[b9].

Game 3:Wolve-Yopt 1-0 starts 1.B[f3] 2.swap 3.W[d9] 4.B[g7] ... 1.B[f3] is unorthodox, possibly too close to the centre to be winning after 2.swap. However, Wolve plays reasonably, and is winning by 21.W[c10].

Game 4:Yopt-Wolve 0-1 starts 1.B[b2] 2.W[d8] 3.B[f7] 4.W[f6] ... Wolve is in its book for its first four moves, and is winning by 16.W[d7]. However, Yopt has no VCE, a considerable end game disadvantage.

Game 5:Wolve-Six 1-0 starts 1.B[a2] 2.swap 3.W[d8] 4.B[g6] ... Six is winning with 32, 34, 36, and 38 but blunders with 40.B[e9]?; for example, 40.B[e8] 41.W[e9] 42.B[e10] 43.W[f9] 44.B[d9] 45.W[e7] 46.B[g7] 47.W[g8] 48.B[e6] 49.W[e5] 50.B[c6] 51.W[d5] 52.B[b5] 53.W[b6] 54.B[c5] is a winning line. Wolve is lucky to win.

Game 6:Six-Wolve 0-1 starts 1.B[i9] 2.swap 3.W[f6] 4.[h5] ... Wolve is winning by 44.B[e6].

Game 7:MoHex-Yopt 1-0 starts 1.B[d2] 2.swap 3.W[b10] 4.B[d9] ... Yopt is losing after 22.B[c5]. This move, which is outside of the mustplay computed by MoHex, might have been avoided if Yopt had a VCE.

Game 8:Yopt-MoHex 1-0 starts 1.B[b2] 2.swap 3.W[e8] 4.B[d8] ... Yopt is winning by 21.W[g8].

Game 9:Six-Yopt 1-0 starts 1.B[i9] 2.swap 3.W[f6] 4.B[d7] ... Six is winning after 26.B[b4].

Game 10:Yopt-Six 1-0 starts 1.B[b2] 2.swap 3.W[e7] 4.B[h6] ... Six is winning by 26.B[e2], but Yopt plays a solid defensive game, and wins after Six blunders with 36.B[i7]? (36.B[h8] wins). Even after this blunder, Six might have saved itself if it played differently after recognizing an opponent win: the current Six version seems to play to prevent the longest win, whereas playing to prevent the shortest win might yield more opponent blunders.

Game 11:MoHex-Wolve 1-0 starts 1.B[d2] 2.W[f6] 3.B[d7] 4.W[e5] ... 22.W[c3] 23.B[b4] is a bad exchange for Wolve, as it gains nothing for Wolve but leaves MoHex with greater influence; this happens frequently with Wolve's circuit evaluation function. MoHex is winning by 27.B[h7].

Game 12:Wolve-MoHex 0-1 starts 1.B[a2] 2.swap 3.W[d8] 4.B[f7] ... MoHex is winning by 30.B[i8].

Players took advantage of the time between rounds to run tests and tweak programs. Tests showed MoHex and Wolve occasionally blundering near the end of the game, so for the second round both programs added a simple end game solver which, after move 15, ran for 15 seconds prior to the search phase.

Round 2. Game 13:Six-MoHex 0-1 starts 1.B[j9] 2.swap 3.W[f6] 4.B[e6] ... MoHex is winning by 30.B[h9].

Game 14:MoHex-Six 0-1 starts 1.B[i2] 2.swap 3.W[h4] 4.B[e6] ... 21.W[b10]? appears to be a MoHex blunder, as 21.W[b7] seems likely to win. For example, 21.W[b7] 22.B[c6] 23.W[c9] leaves MoHex winning. Six is winning by 28.B[d9].

Game 15:Wolve-Yopt 1-0 starts 1.B[a3] 2.swap 3.W[e7] 4.B[c8] ... Wolve is winning by 23.W[i9].

Game 16:Yopt-Wolve 0-1 starts 1.B[b2] 2.W[d8] 3.B[g7] 4.W[g6] ... Wolve is winning by 34.W[d6].

Game 17:Wolve-Six 1-0 starts 1.B[a8] 2.W[c9] 3.B[a10] 4.W[c10] ... Wolve is winning by 31.B[d7].

Game 18:Six-Wolve 0-1 starts 1.B[a3] 2.swap 3.W[f6] 4.B[g6] ... Wolve is winning by 50.B[e9].

Game 19:MoHex-Yopt 1-0 starts 1.B[a4] 2.swap 3.W[b10] 4.B[e9] ... MoHex is winning by 31.W[h6].

Game 20:Yopt-MoHex 0-1 starts 1.B[b2] 2.swap 3.W[f7] 4.B[e7] ... MoHex is winning by 30.B[j3].

Game 21:Six-Yopt 0-1 starts 1.B[i9] 2.swap 3.W[f6] 4.B[g6] ... Yopt is winning by 22.B[c8].

Game 22:Yopt-Six 0-1 starts 1.B[b2] 2.swap 3.W[c9] 4.B[g6] ... Yopt is winning with 45,...,53 but blunders with 55.W[c6]? instead of 55.W[b7]. This game determines the bronze medal; a VCE, end game solver, or ICE (since 55 was a probe to the wrong side of a bridge) might have allowed Yopt to preserve the victory that was in its grasp.

Game 23:MoHex-Wolve 0-1 starts 1.B[b4] 2.swap 3.W[b10] 4.B[c9] ... The opening moves for MoHex look weak, especially after 9.W[j2]. MoHex often seems to get stuck on a bad move, often playing b10 and j2 at

unusual times. Wolve is winning by 22.B[i2]. This game clinches the gold medal for Wolve.

Game 24:Wolve-MoHex 0-1 starts 1.B[a9] 2.swap 3.W[b10] 4.B[c10] . . . Even with the 15-second solver, Wolve makes end game mistakes. A post-game run of solver shows that Wolve blunders with 43.W[i10]?, and that 43.W[h9] would have won. MoHex is winning by 44.B[i7].


Conclusions. The level of Hex play has improved significantly since the previous Hex competition in 2006, when Six soundly defeated Wolve and HexKriger for the gold medal at the 11th Computer Games Olympiad in Torino. Yopt usually plays elegant opening and mid-game moves but is hurt in the end game by its lack of a VCE. Six and Wolve usually play well but make a surprising number of end game blunders, for example in games 5, 10, and 24. MoHex also usually plays well, provided it avoids catastrophic openings, such as in games 8 and 23.

Although occasionally overlooking human-recognizable virtual connections, top Hex programs now play at a strong intermediate level that is almost expert in the end game. Opening and midgame play is still weak; mistimed moves and influence-losing exchanges are common. Improvements can still be made (e.g. MoHex might blunder less with UCT root parallelization; Wolve might play a stronger middle game with an influence-rewarding evaluation function), but a significant strength increase seems unlikely in the near future (UCT programs probably need to incorporate dynamic virtual connection information into rollouts; circuit evaluation programs probably need a better evaluation function and a deeper search; both tasks seem difficult). It could be many years before computers can play 11×11 Hex almost perfectly.


Acknowledgements. The success of Wolve and MoHex is due to the support of many people. We thank NSERC, Martin Mueller, Jonathan Schaeffer, and the UofA GAMES group for financial support; Lorna Stewart for loaning her new computer; the UofA technical support team for the short notice purchase and installation; Martin Mueller and Markus Enzenberger for sharing the Fuego code base; Paul Lu for frequent technical advice; and Yngvi Björnsson, Andrea Buchfink, Laurie Charpentier, Teri Drummond, Leah Hackman, Mike Johanson, Morgan Kan, Maryia Kazekevich, Martha Lednicky, Nathan Po, Jack van Rijswijck, and Geoff Ryan for work on previous versions of Wolve, Mongoose, and our Hex solver.

1. REFERENCES

- Björnsson, Y., Hayward, R., Johanson, M., and Rijswijck, J. van (2007). Dead Cell Analysis in Hex and the Shannon Game. *Graph Theory in Paris: Proceedings of a Conference in Memory of Claude Berge (GT04 Paris)* (eds. A. Bondy, J. Fonlupt, J.-L. Fouquet, J.-C. Fournier, and J. L. R. Alfonsin), pp. 45–60. Birkhäuser.
- 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. Game 1:Six-Mohex 1-0, Game 2:MoHex-Six 1-0 (top row). Game 3:Wolve-Yopt 1-0, Game 4:Yopt-Wolve 0-1 (2nd row). 5:W-S 1-0, 6:S-W 0-1 (3rd row). 7:M-Y 1-0, 8:Y-M 1-0 (4th row). 9:S-Y 1-0, 10:Y-S 1-0 (5th row). 11:M-W 1-0, 12:W-M 0-1 (bottom row). In some games the operator of the losing program resigned once the outcome was obvious.



Round 2. Game 13:Six-Mohex 0-1, Game 14: Mohex-Six 0-1 (top row). Game 15:Wolve-Yopt 1-0, Game 16:Yopt-Wolve 0-1 (2nd row). 17:W-S 1-0, 18:S-W 0-1 (3rd row). 19:M-Y 1-0, 20:Y-M 0-1 (4th row). 21:S-Y 0-1, 22:Y-S 0-1 (5th row). 23:M-W 0-1, 24:W-M 0-1 (bottom row).