

Version 3.0

A lot has changed since the last official version (2.5). Not all that much can be seen from the outside though.

Strategy seeking

- Way better method of parallelizing implemented. This basically solves the problem that some seekers can run out of work to do. See the documentation for more details.
- Reduced memory footprint for seeking artifacts that are sometimes needed in the millions. Especially useful if there are many promoted pieces on the board.

Strategy analysis

- As always, more collision detection added here and there.
- Internal change that fixes issues with very long SPGs (> 47.0 moves). This was a requirement for trying to test the length-record problems.

Parameters

- `strategySeekingSyncDepth`: The recursion depth in the new strategy seeking algorithm at which the strategy seekers synchronize. Lower numbers mean that work chunks are bigger, higher numbers mean that work chunks are smaller. In case the strategy seeking tree is very skewed, it might be useful to use a value up to 18. Going too high (say 100) effectively means that there is no parallelization gain at all, since all seekers then find the whole seeking tree. Going too low can have the effect that some seekers run out of work to do in case the tree is too skewed. See the documentation for more details.

Internals

- A healthy dose of cleanup.
- A lot of refactorings were needed for the strategy seeking changes.
- 3 solver bugs fixed. Nothing too serious. Rare or fail-fast (i.e. internal error) cases.