ChessBase Engine Performance Tests

Mark Kaprielian

2025-03-30

I. Introduction

This document contains some testing I have performed to try to answer some questions I've often had about the CB Engine such as:

- Does the number of lines Displayed impact the analysis
- How much does using more CPUs appear to improve the results or speed of reaching the same results.
- How repeatable is the analysis for a given position

My hope is that some insight will be gained but not necessarily firm conclusions without much more rigorous testing

The tests and results are presented in the order to help best understand them and some of the tests reference previous tests.

Probably the most interesting and useful results of this document is the last section, **"Best Lines"**.

II. The general test conditions

- CB 18
- Game Analyzed: One move played: 1 d4
- Hash table not cleared unless explicitly mentioned it has:
  - According the CB Help: *Kibitzers use "hash tables" to store positions that have already been examined, in case they turn up again in the search tree. This speed things up tremendously, especially in the endgame.*
- Task Manager: CPU performance window running, No browser windows open or other user activity going on.
- NOTE:
  - Normally we observe the depth grow from 10 to the high 40s over time when analyzing a middle game position
  - During these tests the Depth started out at about 35 and stayed in that range. Apparently, the test conditions being used are not appropriate for testing depth over time. I suspect the complexity of the position would have a large impact and thus make it difficult to profile the behavior.

III. kN/s vs CPUs and Lines viewed

Three test runs were made:

- Pass 1: To establish a base line of measurements, 8 lines being displayed
- Pass 2: Same as Pass 1 to check repeatability
- Pass 3: Change of lines being viewed from 8 to 4 to see what impact it may have

The results appear to be repeatable

There is an anomaly in Pass 1 where a jump in performance occurred then returned to the trend when more CPUs were added. Pass 2 did not show this anomaly.

Pass 3 reduced the number of lines displayed to see if there was any general impact. The anomaly was seen again but with a higher number of CPUs.

I the anomaly was due to some background activity of the OS I would expect that performance would go down when such activity occurred, not up! There are a lot of things in the mix:

- The OS
- CB handling of UCI (Universal Chess Interface) resource management
- Stockfish resource management

CB support had told me in one of my cases that CPU and memory management depends on how the engine creators manage the resources from within the UCI.

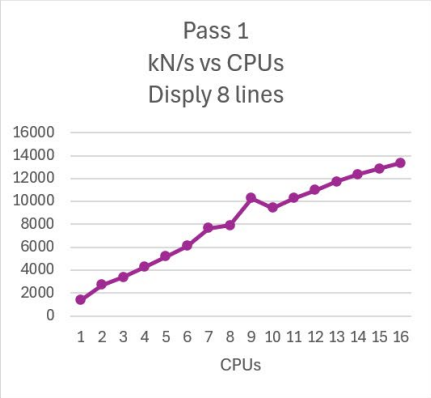

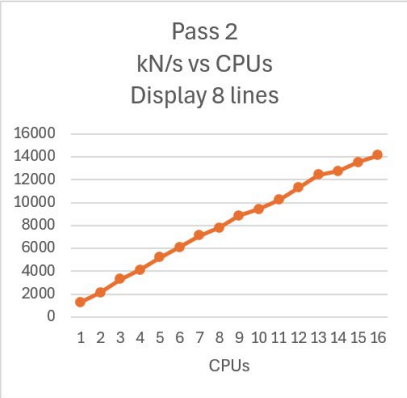

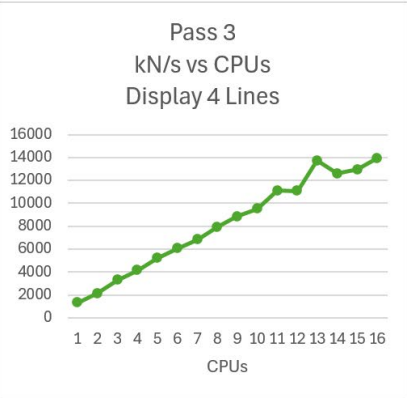

Since the kN/s trends in a straight line as you add CPUs, the following graphs show the relationship.

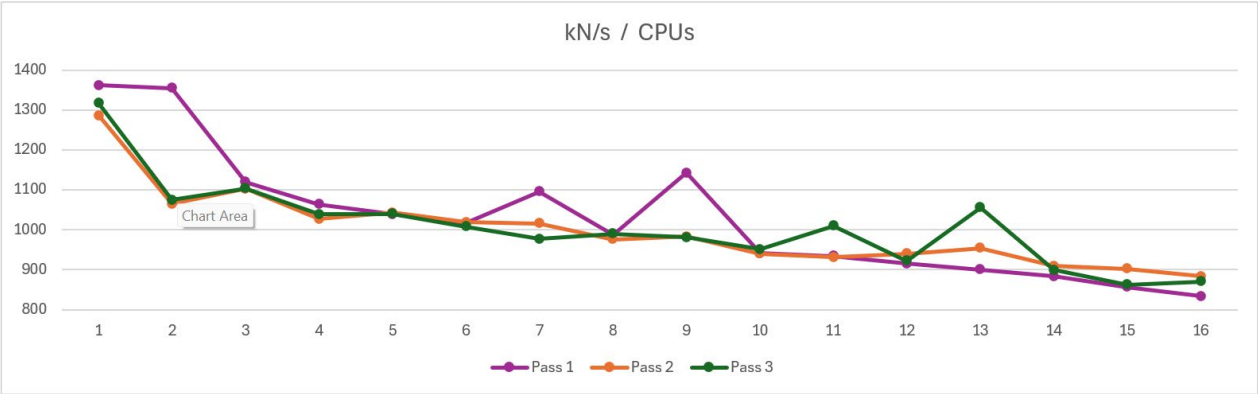

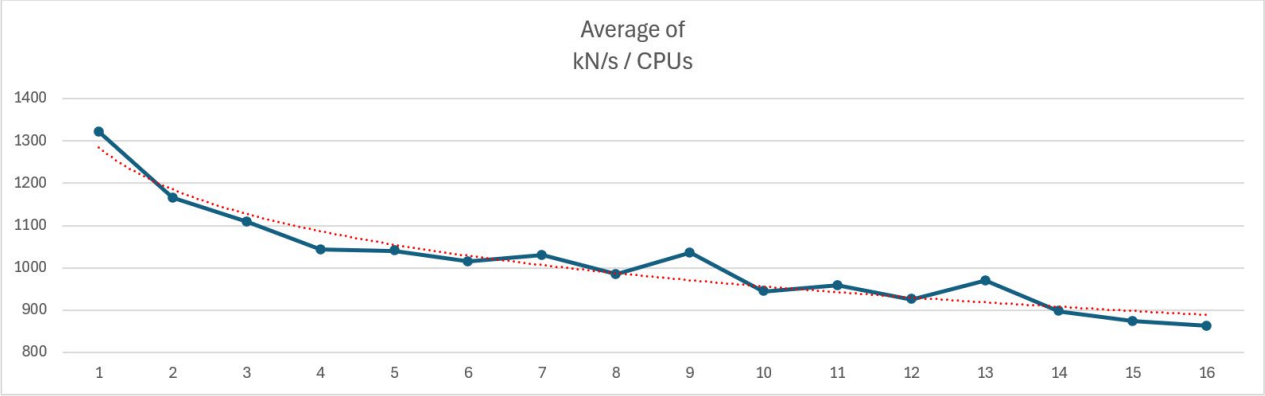

The top graph shows how closely the performance tracked across the three Passes. The spikes at 9 and 13 CPUs are the anomalies noted previously.

The bottom graph is an average of the three Passes

What I find interesting is that the more CPUs you add, the performance drops off. Not by a lot but there seems to be a clear trend. One theory could be that it's going out so much deeper that it has more work to do and thus it slows down. Here is the data from Pass 1 showing Depth.

Notice that the depth does not appear to increase in relation to the number of CPUs going up while performance is clearly going down. Interesting also I think is that the current ply being analyzed for the depth remains about 34 after 1 minute despite the number of CPUs

CPUs	kN/s	Depth		kN/s / CPUs
1	1362			1362
2	2710	35	51	1355
3	3360	33	59	1120
4	4256	33	60	1064
5	5194	34	63	1039
6	6104	34	57	1017
7	7670	38	66	1096
8	7908	34	51	989
9	10283	39	58	1143
10	9419	34	61	942
11	10280	34	71	935
12	10979	35	64	915
13	11705	34	57	900
14	12364	34	63	883
15	12854	34	65	857
16	13335	34	72	833

IV. How you change CPU count appears to impact results

12 CPUs is my preferred sweet spot for my PC where I can still use my PC to do a lot of other things at the same time, I'm letting the engine sit on a position. I based this on fiddling with the number of CPUs after I determined that my optimal Hash table size is 4096. With this size Hash table, the engine loads and unloads in under a second.

To get the measurements I made a video of the lines displayed then played it back in paused mode to get the performance at ten second intervals.

Having set the cores to 12 the performance at one minute should match the three Passes

The previous test passes show at 1 minute the following:

- Pass 1: 10979
- Pass 2: 11280
- Pass 3: 11072

In a different set of 4 test runs, at 12 cores at 1 minute we see:

- New Pass 1: 13994
- New Pass 2: 14003
- New Pass 3: 12452
- New Pass 4: 13882

There was a difference in how the first three Passes were run compared to these later runs.

For the first three Passes:

- The engine was running from the previous CPU test
- The CPU button was pressed which stopped the engine
- The number of CPUs to run was increased by one
- The OK button was pressed and as soon as the display window opened the timer for a minute started.

For the later test runs, the number of CPUs never changed.

- The engine was running then paused
- The video recording was started
- The engine was un-paused, which resets the calculation
- The video is stopped when 5 minutes have been recorded.

So, better performance seems to be due to not having changed the CPUs in the later tests as opposed to the early tests when changing the CPU was required.

Further tests could be done to verify if this was the cause of not making use of the running engines ability to change CPUs but actually changing the default.

V. Performance as you let the engine keep running

Here's a visual of what happens to performance as the longer you let the engine run.

In the graph below, performance settles out to about 11000 at five minutes.

As mentioned earlier, due to the test condition of having the position of 1 d4 being analyzed, no insight on how depth was of significance as the depth starts and stays near 34 right from the start of every test.

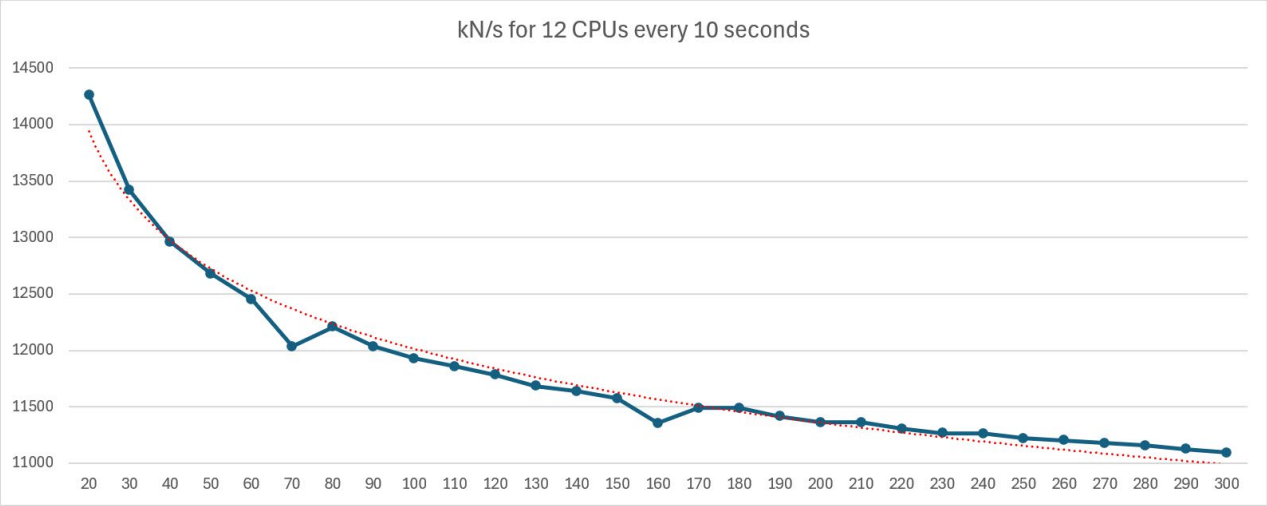

VI. Best Lines – Are we getting the same results every time?

Four Passes were made, in the following order:

- Pass 1: 12 CPUs, 5 minutes, Hash untouched, 8 lines displayed
- Pass 2: 12 CPUs, 5 minutes, Hash untouched, 8 lines displayed
- Pass 3: 12 CPUs, 5 minutes, Hash cleared, 8 lines displayed
- Pass 4: 12 CPUs, 5 minutes, Hash untouched, 4 lines displayed

Pass 1 and Pass 2 were the same scenario and we would hope, provide the same results

Pass 3 had the Hash table reset. As mentioned earlier in the test conditions section, the hash should speed up calculations as it stores previous results. Since we've just tested the same move twice, in theory then, we should get faster / deeper results.

Pass 4 changes the lines displayed to see if it appears to have any impact

The answer appears to be:

- You get about the same results and the lines in the display don't matter if you wait 5 minutes.

Notice the Depth changes. Remember Pass 2 had the benefit of Pass 1 due to the Hash table. Pass 3 had the Hash table cleared so a clean run, actually the cleanest run here, but Pass 4 had the benefit of Pass 3.

And we thought computers should always give the same results.

Here are the results for each pass at the end of the test.