



HEP-SPEC06 on Intel Xeon Scalable Processors (Skylake)

23.02.2018 - HEPiX Benchmarking Working Group

Luca Atzori (luca.atzori@cern.ch)

Goal

- Compare HS06 scores on different Skylake SKUs.
- Approach:
 - swapping CPUs on the same bare-metal node 🤖;
 - using openlab scripts ⇒ 64 bit compilation.
- Qualitative analysis:
 - interested in overall behavior;
 - average of 2 / 3 runs per CPU model;
 - main focus on scalability.

Environment

Hardware & BIOS	
Family	Skylake
CPU	Dual Socket
RAM	192 GB
SMT	ON
TurboBoost	ON


Software	
OS	CentOS 7.3
Kernel	see next slide
SPEC 2006 version	1.2
Compiler	gcc 4.8.5 20150623

Tested SKUs

CPU Model	Cores / Threads (per socket)	Base / Turbo CLK (GHz)	Kernel version
Silver 4114	10 / 20	2.20 / 3.00	3.10.0-693.17.1
Gold 5122	4 / 8	3.60 / 3.70	3.10.0-693.5.2
Gold 6154	18 / 36	3.00 / 3.70	3.10.0-693.11.6
Platinum 8170	26 / 52	2.10 / 3.70	3.10.0-693.5.2

HS06: raw scores

	P8170	G6154	G5122	S4114
1	37.86	37.08	35.36	29.02
2	75.61	73.59	70.45	57.82
4		135.96	130.29	106.52
8		256.32	243.38	194.02
10				234.44
16	462.52	478.09	<u>316.57</u>	360.05
18		529.66		
20				425.46
32		856.86		
36		918.78		
40				<u>541.45</u>
52	1089.22			
72		<u>1098.42</u>		
104	<u>1261.01</u>			



HS06: normalized scores

All frequencies normalized to 2.7GHz.


Added to comparison:

- E5-2683v3: Haswell, 14c, 2.0-3.0GHz
- E5-2630v4: Broadwell, 10c, 2.2-3.1GHz
- E5-2699v4: Broadwell, 22c, 2.2-3.6GHz

Would need more fine-grained samples for the Platinum 8170.

Unfortunately, no data for the 88 processes score for the E5-2699v4 in our DB. Would have been interesting to see if same behaviour of Gold 6154.

Platinum 8170 shows best performance: bigger L3 cache compared to other SKL models and bigger L2 compared to previous gen (1M vs 256K).




HS06: scalability

Removed the E5-2699v4 and the Platinum 8170 (too high end!) and Gold 5122 (too few cores!).

- HS06 scalability doesn't seem to have changed a lot respect to previous gens.

- E5-2630v4 vs Silver 4114: much closer (but BDW performs better!). Should investigate more.



Conclusion

General considerations

- Use of HyperThreading never leads to worst scalability (HT=OFF is not convenient?);
- No use of AVX512, which is a key component of Skylake;
- TurboBoost doesn't look like a big issue for this kind of comparisons.
- Different kernel versions!
 - No security fixes for Gold 5122 and Platinum 8170 (3.10.0-693.5.2)
 - Available for Gold 6154 (3.10.0-693.11.6)
 - Spectre v2 removed on Silver 4114 (3.10.0-693.17.1)

Future work

- Power consumption benchmarks
- Fast benchmarks, SPEC 2017, ...
- Accurate HS06 analysis (maybe focusing on less SKUs)

FULL REPORT: https://luatzori.web.cern.ch/docs/Skylake_HS06_Full_Report.html