



Benchmark Working Group

Toward a New HEP CPU Benchmark – Summary of Discussions and Remarks –

Manfred Aef (KIT), Michele Michelotto (INFN-PD)



Benchmark Working Group

GDB 2014-01-15:

- Talk: Toward a New HEP CPU Benchmark
- Lot of comments sent to WLCG GDB mailing list

Today:

- Summary on comments and remarks



Benchmark Working Group

Question #1: Number of parallel benchmark runs?

- Number of cores of the SUT, or
- Number of job slots?

SUT: system under test



Benchmark Working Group

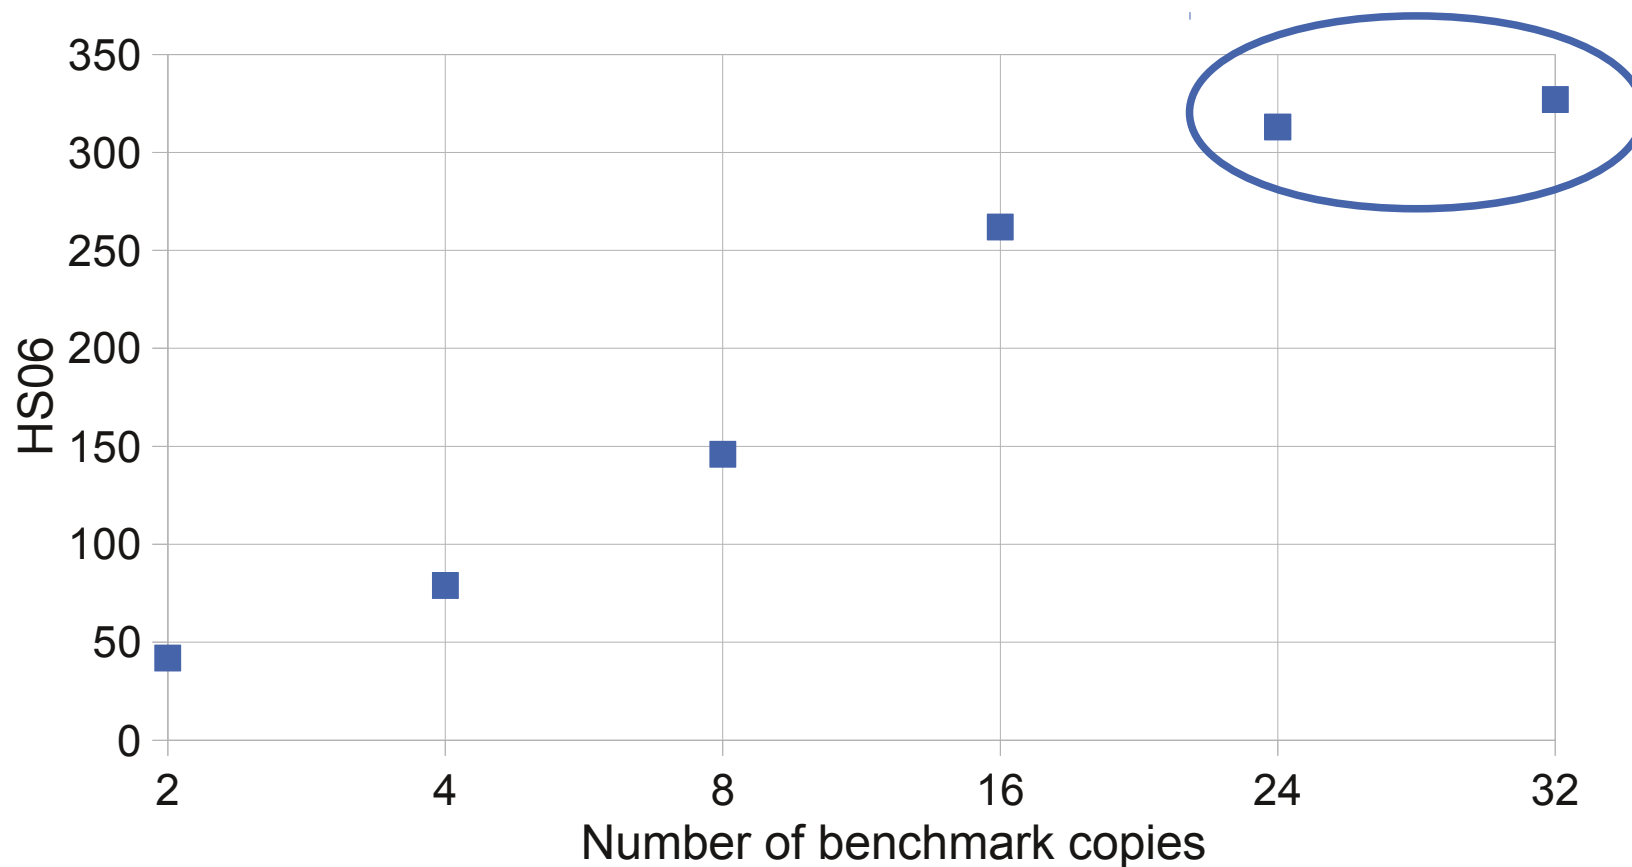
Answer #1:

- One major rule of HS06 benchmarking is to run the benchmark in exactly the same system configuration which is also used for production:
 - If a site runs WNs with hyperthreading disabled, they must run the benchmark on a SUT with HT disabled, and vice versa.
 - The number of HS06 benchmark copies should be set to the same number of job slots like the WNs are configured.
- When the HS06 has been developed, there were only chips without hyperthreading on the market. The first systems with hyperthreading appeared at the end of that project, and therefore the documentations don't reflect that topic.



Benchmark Working Group

Performance (HS06) of 2-way Intel E5-2670 (Sandy Bridge, 8-core)





Benchmark Working Group

Outlook #1:

→ Should become more clear in HS14



Benchmark Working Group

Question #2:

- Matchmaking, correction factors in VM, accounting

Answer #2:

- Be aware that HS06 results should scale with the **average performance of the job mix** which is running in the wild
- HS06 scores measure performance of fully loaded systems; applications may run faster if WN is partially idle
- Accounting uses the worst-case power



Benchmark Working Group

Question #3:

- We should have something that is open source (distributed via HEP_OSlibs) and can run in less than a minute (so admins can add it to the OS or batch system boot sequence)

Answer #3:

- Such a light benchmark tool will most probably not scale very well with the mix of real applications ... :-(



Benchmark Working Group

Question #4:

→ Conversion SI2k to HS06 (and vice versa)

Answer #4:

- The HS06 benchmark working group had proposed a conversion factor of 4 to simplify migration from SI2k to HS06
 - This factor is an average of the HS06 per SI2k ratios of several test systems
- Now the proper way of benchmarking is to run HS06, and to divide by 4 to get the corresponding **kSI2k** score (or to multiply with 250 for the SI2k value ;-)



Benchmark Working Group

Question #5:

→ Why don't HS06 scores scale with my application?

General remarks #5:

- HS06 results (= average of 4 independent floating-point and 3 integer based software packages) shall scale with the **average performance** of the job mix which is running in the wild
 - There is absolutely no warranty that each individual job scales exactly like HS06 nor like the real job mix
 - Please apply appropriate correction factors to be safe



Benchmark Working Group

Answer #5:

→ 3 Examples:

- "Atlas jobs running inefficient on hosts with many cores"
 - ◆ Hardware issue?
- 2 examples of poor performance of code compiled with non-conforming compiler flags running in virtual environments
 - ◆ Code was compiled with 3dnow enabled, but hardware feature not available within the VM
 - ◆ Theoretical chemistry code compiled with XXX compiler, again some hardware features hidden within the VM



Benchmark Working Group

Contact:

- manfred.alef@kit.edu
- michele.michelotto@pd.infn.it