

ALICE CPU benchmarks

Costin.Grigoras@cern.ch

Benchmark considerations

- **~70% of the Grid time is taken by simulation jobs**
- **A benchmark has to be simple to find and to run**
- **Short execution time relative to the job duration**
 - For automatic benchmarking of nodes
- **Reflecting the experiment's software performance on the hardware**
- **Simplified method to collect summarize and share the results**
- **No licensing concerns**

MC simulation vs benchmarks

(old results from Sep 2015)

- **Reference production:**

- “*pp 13 TeV, new PYTHIA6(Perugia-2011) min.bias, LHC15f anchors*”
- 200 ev/job, avg(8h) running time, CPU-intensive
- Blanket production, 76 sites

- **Benchmarks:**


- ROOT's */test/stress* (O(30s))
- *condor_kflops* from ATLAS' repository (if found) (O(15s))

- **Each benchmark ran twice after the simulation**


- To fill in the CVMFS cache and load the libraries in mem
- Recording the second iteration only

Results at a glance

ALICE pp event simulation time vs ROOT `stress` marks and Condor KFlops




Events/s vs KFlops



No correlation between Kflops and simulation performance, probably because of small ratio of floating point operations in it.


Events/s vs rootmarks



**Rootmarks
scale ~better
with the
simulation time
in the Grid
environment**


Is there anything better still?

- Running in a controlled environment
 - ALICE Central Services machines (~50 hosts)
- ROOT stress test results don't look that good




Sysbench ?

- Available by default on many Linux variants
- ~30s to run
- But it doesn't scale well...




GeekBench

- **Commercial product**
- **Evaluation license for the 64b version**
- **~2min to run**
- **Good results in the test environment**




GeekBench – correlation with MC jobs

ALICE PbPb event simulation time vs Geekbench score



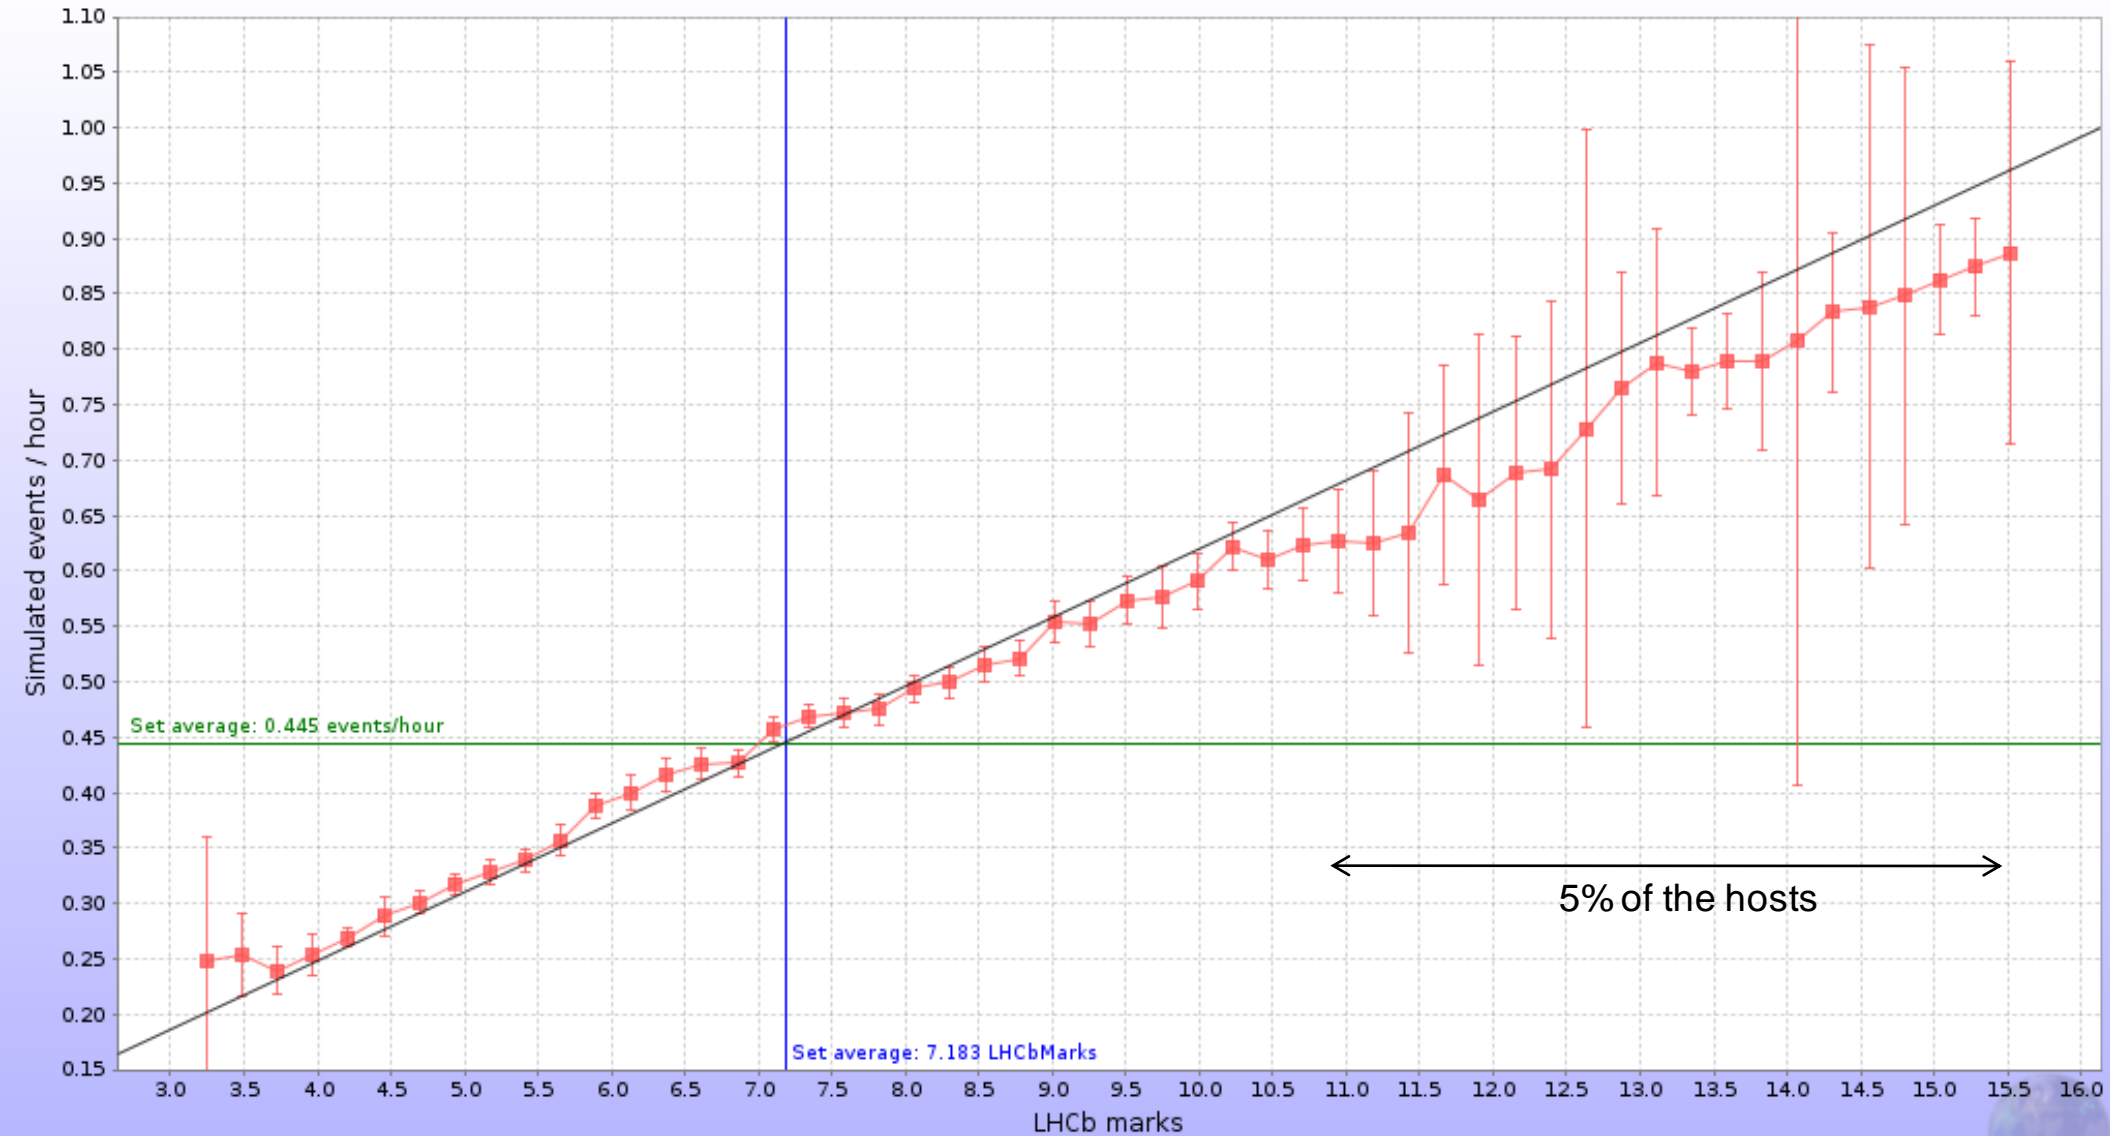
LHCb's test

- Simple *python* script
- ~1 min to run
- Used to estimate how many events the job will be able to generate in a fixed amount of time
- Very good results on the CS machines



Best correlation so far


ALICE PbPb event simulation time vs LHCbMarks



■ 13310 host averages, average distance to (0, 0, 20.0, 1.239609) = 0.02767875

HepSpec06 from MJF

ALICE PbPb event simulation time vs HepSpec06



NERSC

Intel(R) Xeon(R) CPU E5-2698 v3 @ 2.30GHz = 0.834

HIP, CERN

Torino

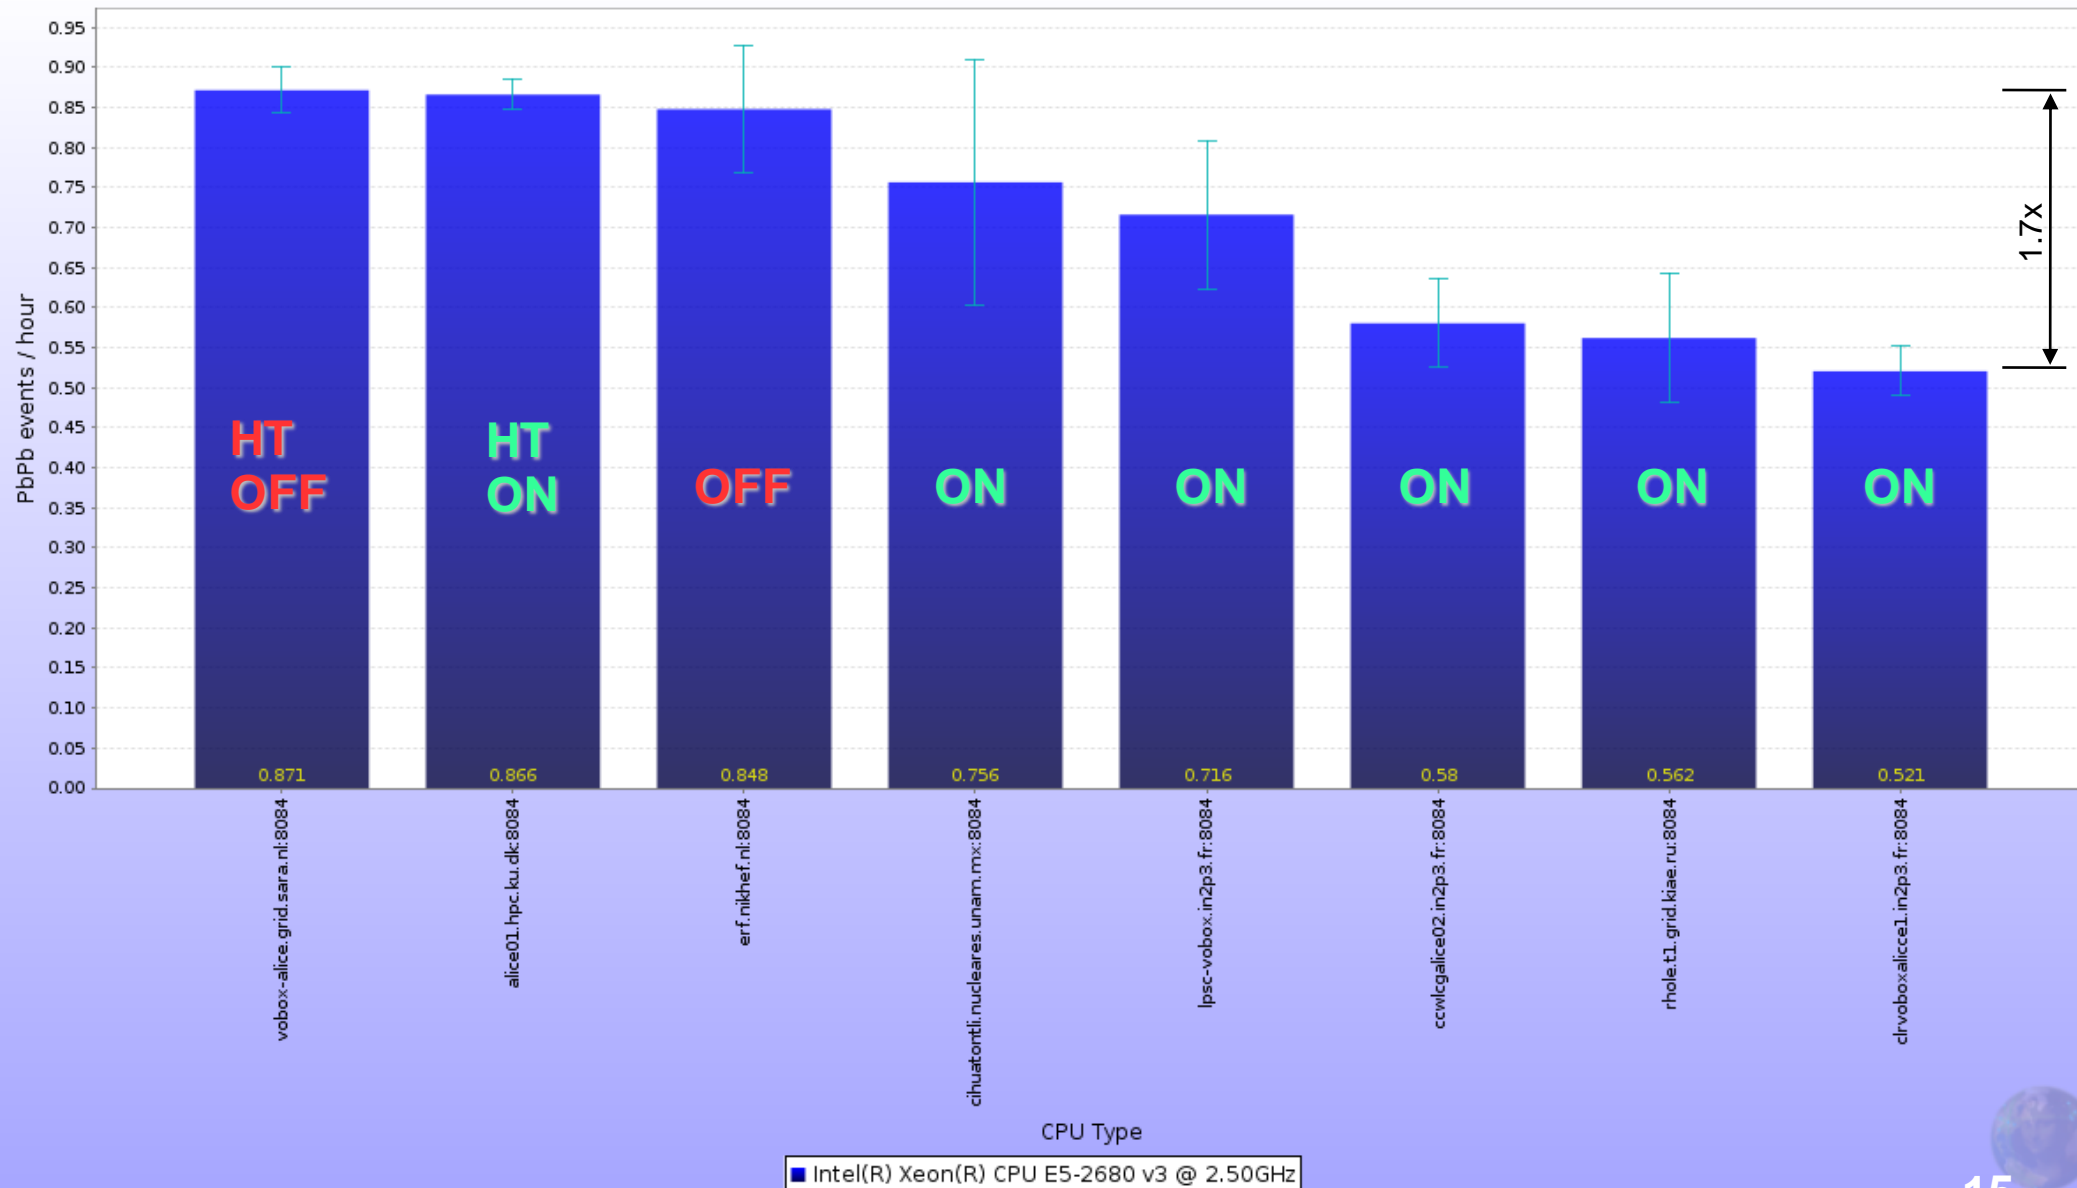
HLT



Site-specific configurations

HT on/off, mem type, #of slots / machine, ...

Intel(R) Xeon(R) CPU E5-2680 v3 @ 2.50GHz



Summary

- **Extensive database of Grid nodes' actual performance**
 - Can be now queried/populated by calling:
[http://alimonitor.cern.ch/marks/?cpumodel=M&hostname=H\(&site=S\)\(&lhcmarks=L\)](http://alimonitor.cern.ch/marks/?cpumodel=M&hostname=H(&site=S)(&lhcmarks=L))
- **Lots of distinct configurations**
 - ~14K hosts
 - 114 CPU models (155 combinations with HT on/off)
- **Can compare any benchmark in production**
 - LHCb's looks very good

Plans

- **Use this database to get the slot performance in each pilot job**
 - fast benchmark run + the history of the same (or similar) nodes
- **Account for the used CPU in this unit**
 - A common version of the fast benchmark that also queries the database for more stable values
- **Maybe even feed back actual execution performance (events/second/unique job type)**
 - Automatic measurement, also good for cloud resources
 - Fair accounting of provided CPU power