

ALICE CPU benchmarks

Costin.Grigoras@cern.ch



Benchmark considerations

- **Simple to find and to run**
- **Short execution time relative to the job duration**
- **Reflecting the experiment's sw performance on the hardware**
- **No licensing concerns**

MC simulation vs benchmarks

- **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:**

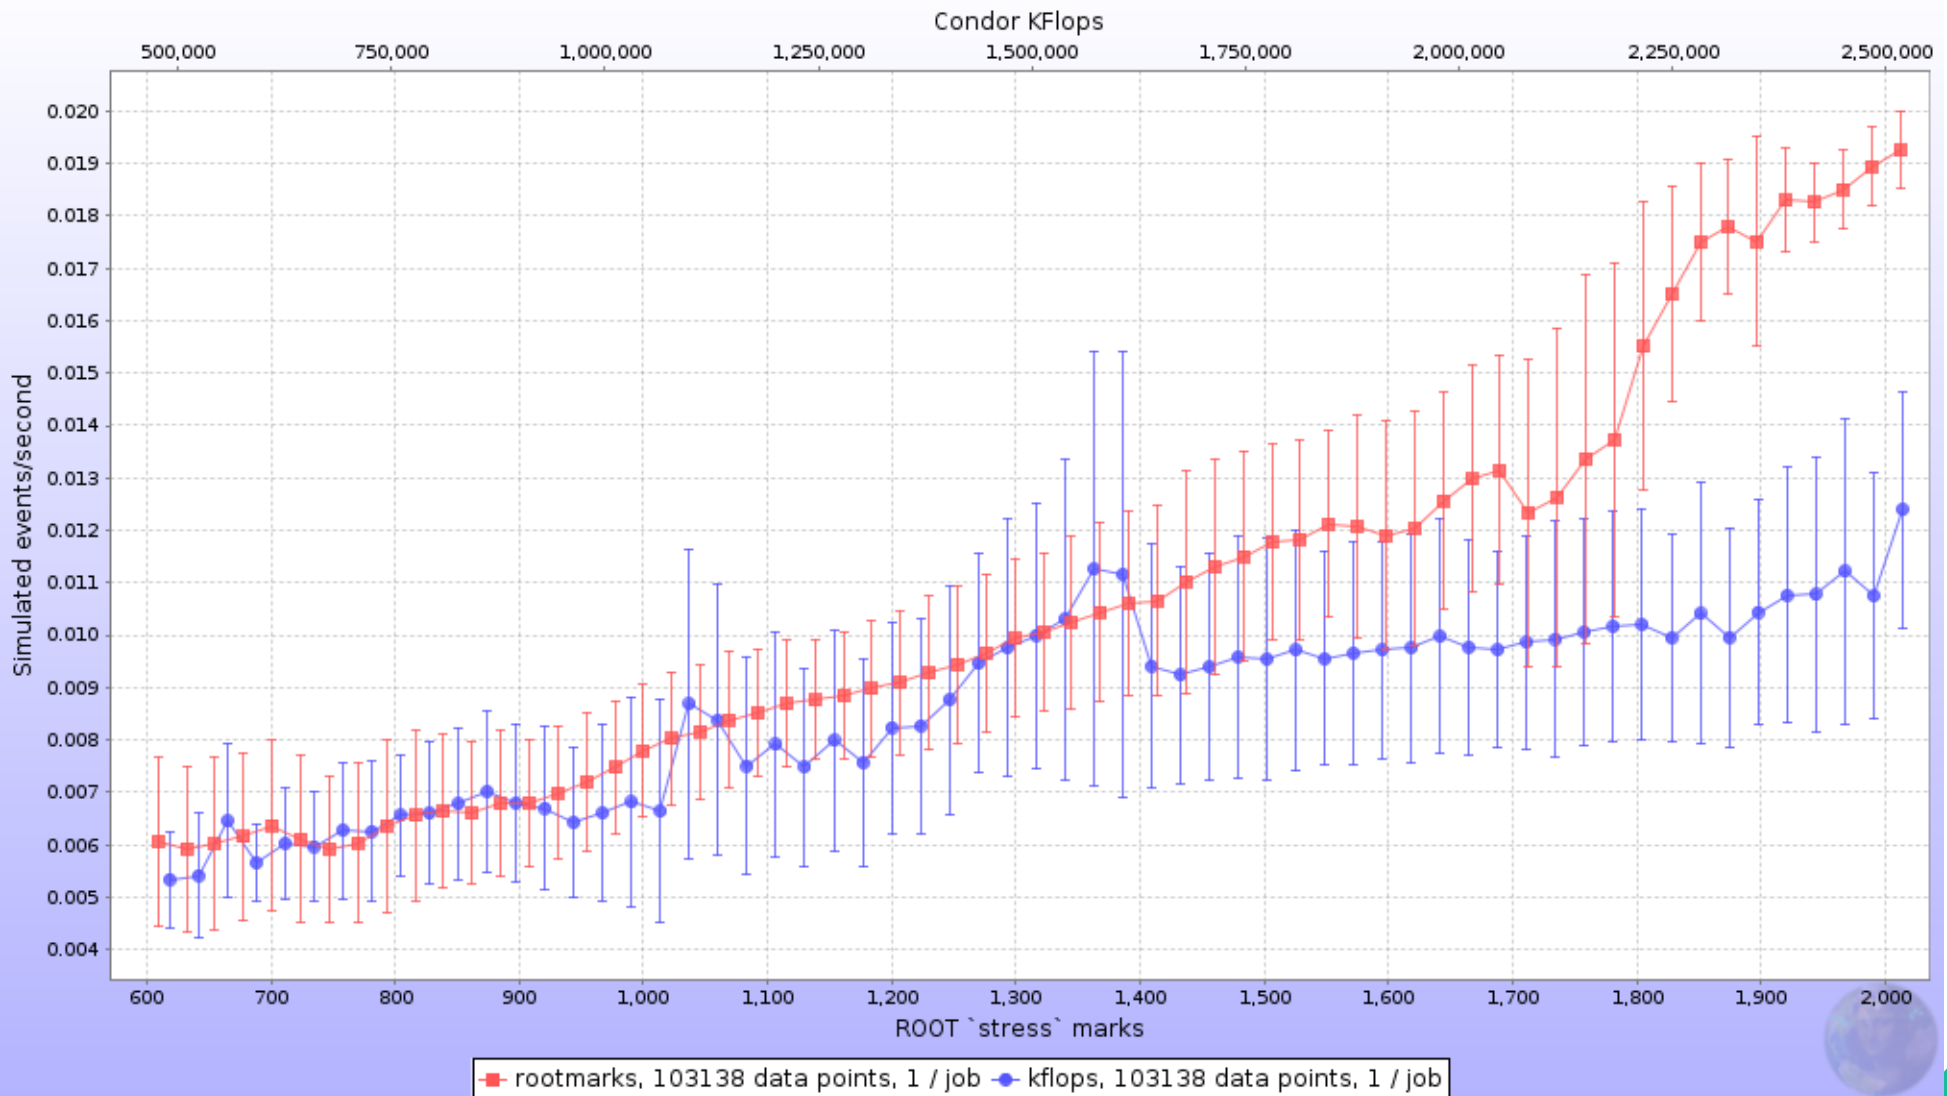
- ROOT's `/test/stress` ($O(30s)$)
- `condor_kflops` from ATLAS' repository (if found) ($O(15s)$)

- **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 rootmarks


ALICE pp event simulation time vs ROOT marks



Rootmarks scale linearly with the simulation efficiency

Events/s vs KFlops

ALICE pp event simulation time vs Condor KFlops



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

Summary and plans

- **Coverage**
 - 11783 nodes from 76 sites
 - 106 distinct CPU models
- **ROOT test good enough to estimate experiment software performance**
 - How does this scale with other VO's sw?
 - Does not scale with the few HS06 results we found
- **Planing to build a per node database of benchmark results**
 - Current one based on CPU model
 - Better even if this was a common service across VOs