

# Passive benchmarking of ATLAS Tier-0 CPUs

Andrea Sciabà

HEPiX benchmarking WG meeting  
March 10 2017

# Introduction

- Goal
  - Use real jobs to measure the relative speeds of different CPU models
- Method
  - Measure the averages of CPU time/event on the desired sample of jobs
  - Assume that it's inversely proportional to the CPU speed
  - Use the fact that “tasks” contain very similar jobs, so all jobs in the same task can be considered different runs of the same “benchmark”
- Detailed description of the analysis
  - Refer to recent pre-GDB [talk](#)

# Application

- The analysis was applied to jobs run at the ATLAS Tier-0
  - October 2015 - March 2016: SMT on
  - March 2016 - today: SMT off
  - The change was done for efficiency reasons
- Selection cuts
  - Only successful jobs are used
  - No. events per job  $> 1000$  (to avoid a bias from the job initialization/finalization time)

# CPU models

- Only six CPU models were used by the ATLAS Tier-0

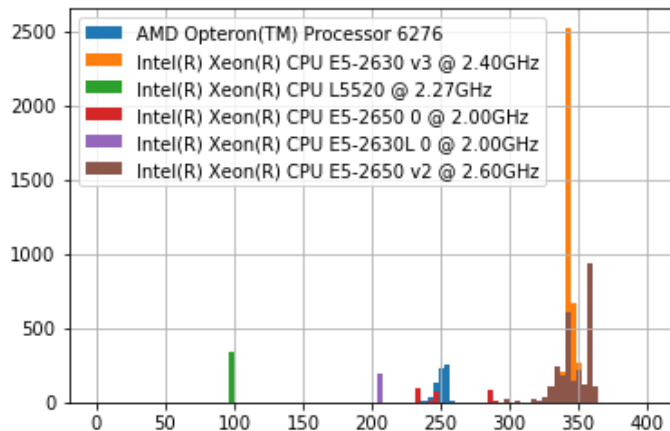
Model name	Architecture	Physical cores	SMT on (1 <sup>st</sup> period)	SMT off (2 <sup>nd</sup> period)
Opteron 6276	Interlagos	2	yes	no
Xeon E5-2630 v3 2.40 GHz	Haswell	8	yes	yes
Xeon L5520 2.27 GHz	Nehalem	4	yes	no
Xeon E5-2650 2.00 GHz	Sandy Bridge	8	yes	no
Xeon E5-2630L	Sandy Bridge	6	yes	yes
Xeon E5-2650 v2 2.60 GHz	Ivy Bridge	8	yes	no

# Question

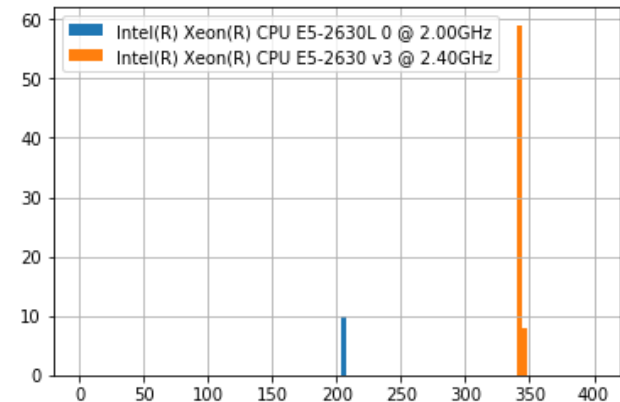
- What I would like to understand from this data is:
  - Do Tier-0 reconstruction jobs scale well with HS06?

# HS06 data

- I use the CPU performance data collected by the CERN procurement team
  - [https://hwcollect.cern.ch:9000/hwinfo/\\_design/hwinfo/\\_view/cpuperformance](https://hwcollect.cern.ch:9000/hwinfo/_design/hwinfo/_view/cpuperformance)
- HEPSPEC06 measured on the bare metal
  - Effect of virtualization not taken into account



CPU found when SMT on



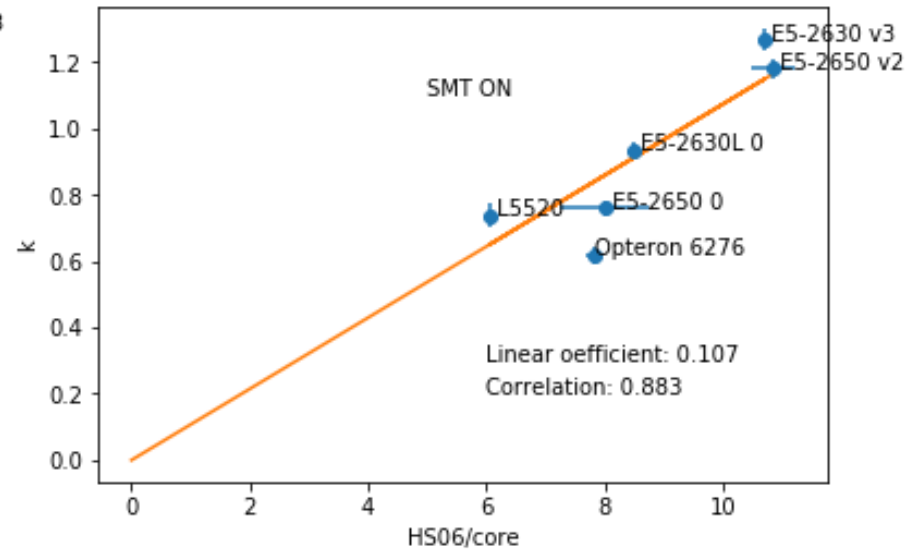
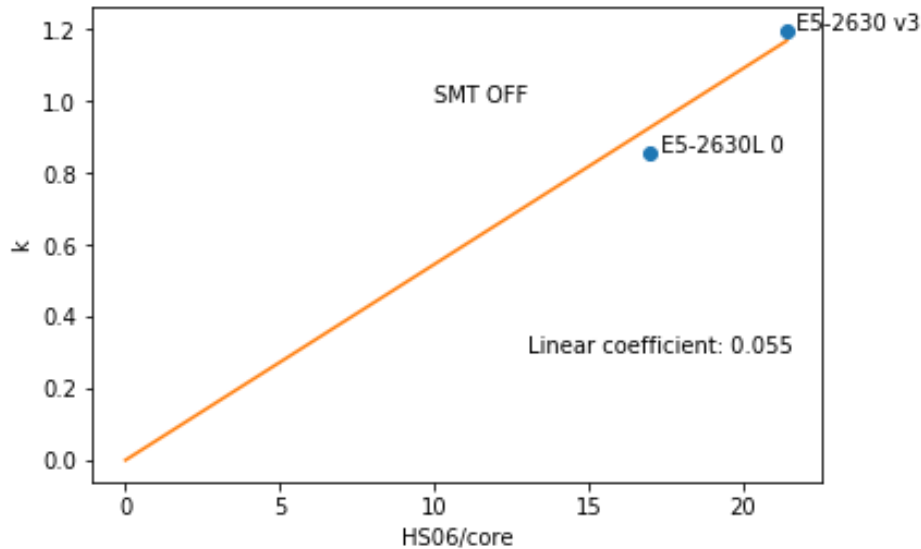
CPU found when SMT off

# HS06 per core (or HW thread)

Model name	Hardware threads	Physical cores	SMT on (1 <sup>st</sup> period)	SMT off (2 <sup>nd</sup> period)
Opteron 6276	16	8	$7.82 \pm 0.13$	
Xeon E5-2630 v3 2.40 GHz	8/16	8	$10.72 \pm 0.09$	$21.42 \pm 0.06$
Xeon L5520 2.27 GHz	8	4	$6.065 \pm 0.004$	
Xeon E5-2650 2.00 GHz	16	8	$8.00 \pm 0.74$	
Xeon E5-2630L	6/12	6	$8.49 \pm 0.08$	$17.0 \pm 0.0$
Xeon E5-2650 v2 2.60 GHz	16	8	$10.85 \pm 0.36$	

Always 2 CPUs per node

# Speed factor k – HS06 correlation



- Scaling generally good, two exceptions
  - Opteron way off
    - HS06 measures better performance than reco jobs w.r.t. Xeons
  - Haswell tends to perform better than expected than Sandy Bridge
    - If “expectation” is given by HS06



# Sandy Bridge vs. Haswell

- Let's compare the values of the coefficient between k and HS06/core for SB and Haswell
  - $k = \alpha \times \text{HS06/core}$

	Sandy Bridge	Haswell	Ratio
SMT ON	0.109	0.118	1.08
SMT OFF	0.0504	0.0559	1.11

Values of alpha

In both cases, Haswell shows an improvement in throughput which is ~10% larger than what HS06 would predict

# Conclusions

- Looking at ATLAS Tier-0 CPU consumption, scalability of HS06 with performance is reasonable, but
  - It is not good for AMD processors
  - It hints at a mismatch from SB to Haswell that goes in the direction of what was seen for DB12, but with a rather lower effect

# References

- Analysis notebooks available at
  - <https://github.com/sciaba/ATLAS-CPU-time-analysis>
  - Not very well commented, sorry

# Acknowledgements

- Many thanks to Jaroslav Guenther, Mario Lassnig and all the ATLAS Tier-0 team
- And to Marco Guerri for making the procurement benchmarking data available