Karlsruhe Institute of Technology

Results of HS06 Scaling Studies at GridKa
WLCG Workshop Lisbon 2016-02-01

Manfred Alef

Steinbuch Centre for Computing (SCC)

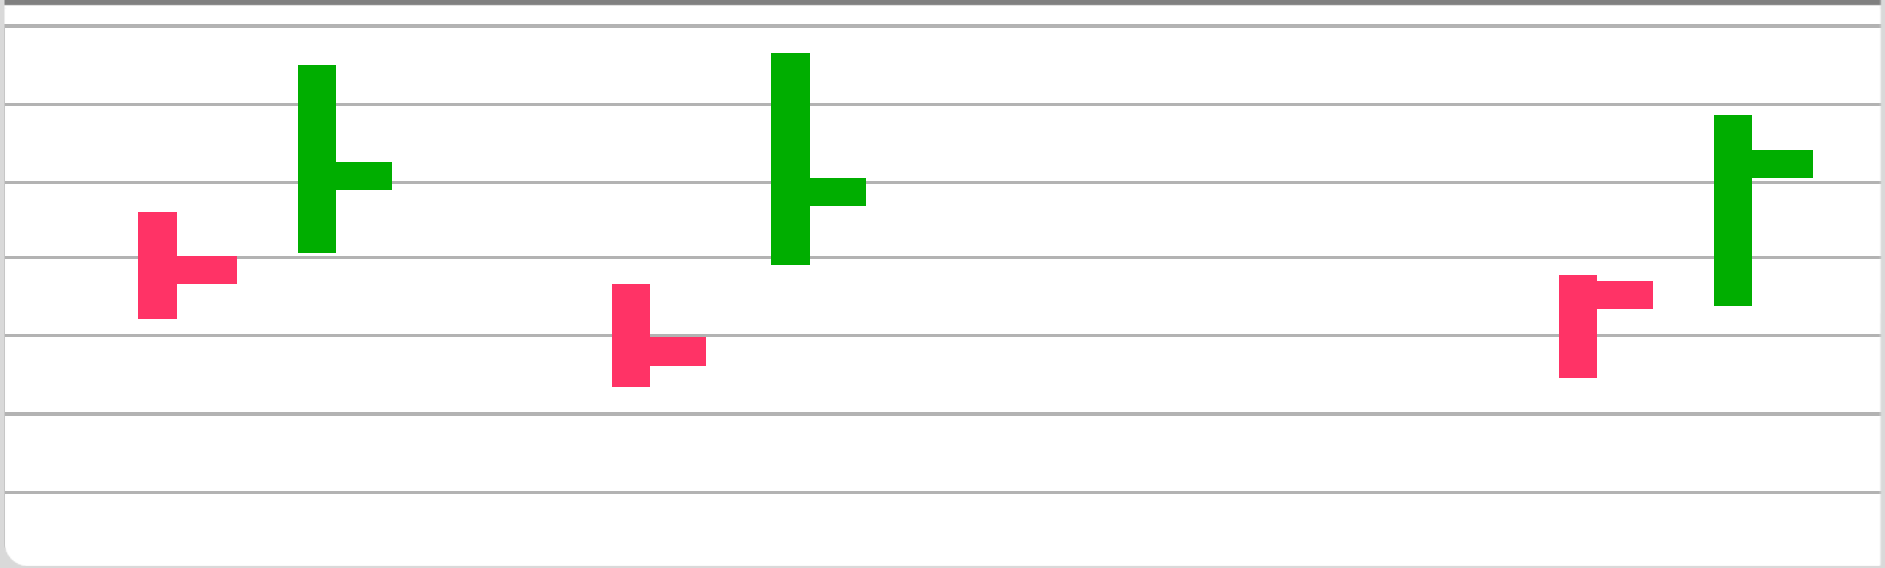

Disclaimer

- HS06 should scale with average HEP applications
  - ➔ 10% initial objective
- **It has <u>not</u> been agreed that HS06 scales with each individual batch job, nor with any class of HEP applications**

Disclaimer

- HS06 themselves is a mix of 7 benchmark packages (3 INT, 4 FP)
  - ➔ HS06 := geometric mean of the 7 individual performance scores (summed up over all parallel benchmark copies)

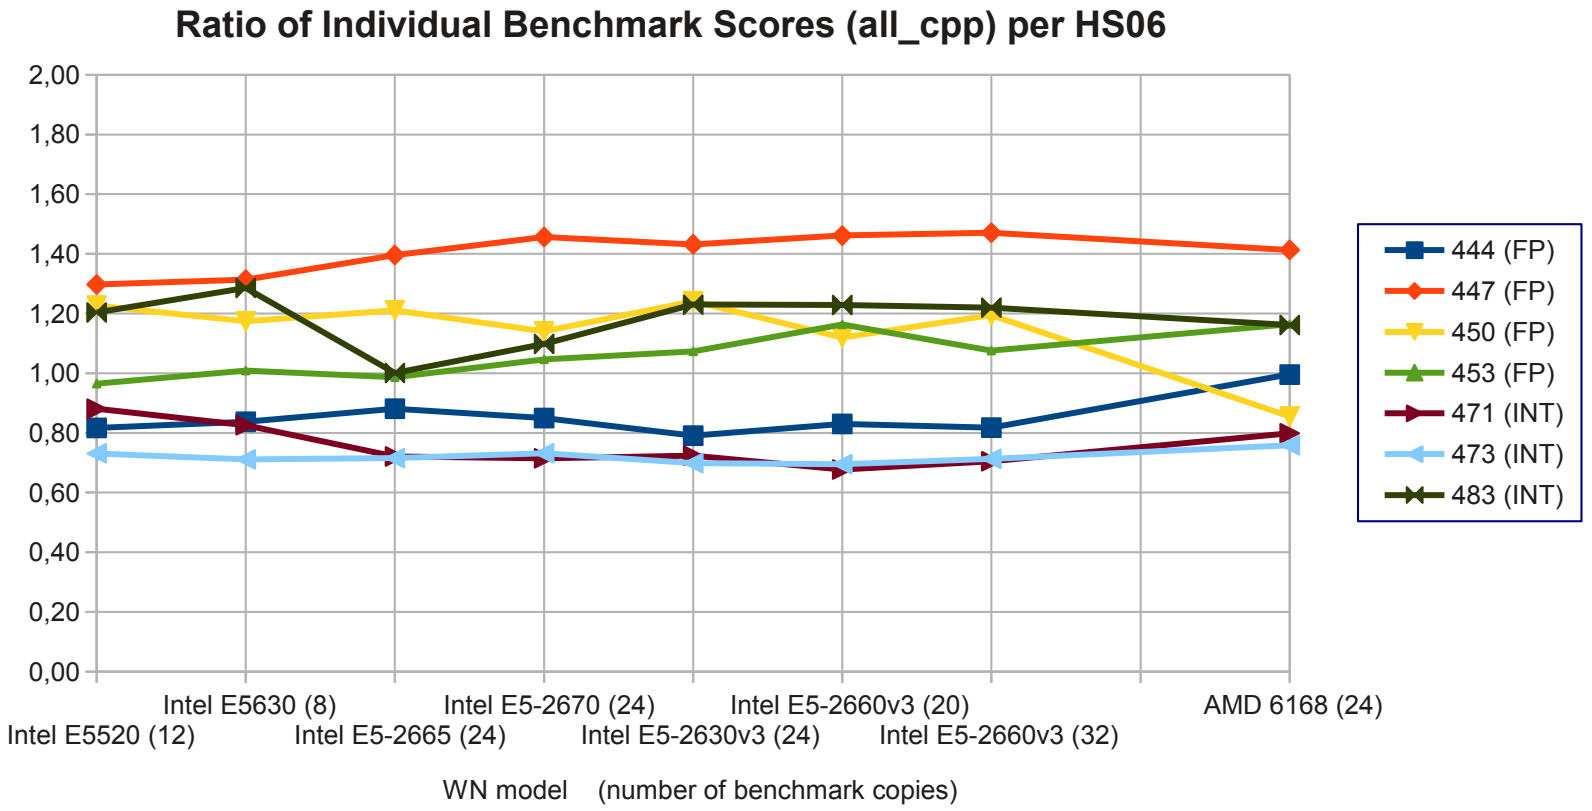

Status 2015

- Philippe Charpentier has found out that LHCb jobs are running some 45% faster on Intel Haswell platform (and around 30% faster on AMD) than estimated by HS06

https://indico.cern.ch/event/319754/session/0/contribution/8/attachments/1202029/1749779/151209-MJFUpdate-LHCb.pdf
https://indico.cern.ch/event/319751/session/0/contribution/6/attachments/1153280/1656518/150909-MJFandBenchmarking-LHCb.pdf

Status 2016

- **Analysis on Atlas jobs**

  - ➔ Many thanks for helpful assistance:
    - Thomas Hartmann
    - Rodney Walker
    - Alessandro di Girolamo
    - Domenico Giordano

Status 2016

- Atlas simulation jobs are running around 35...40% faster on Intel Haswell platform ...

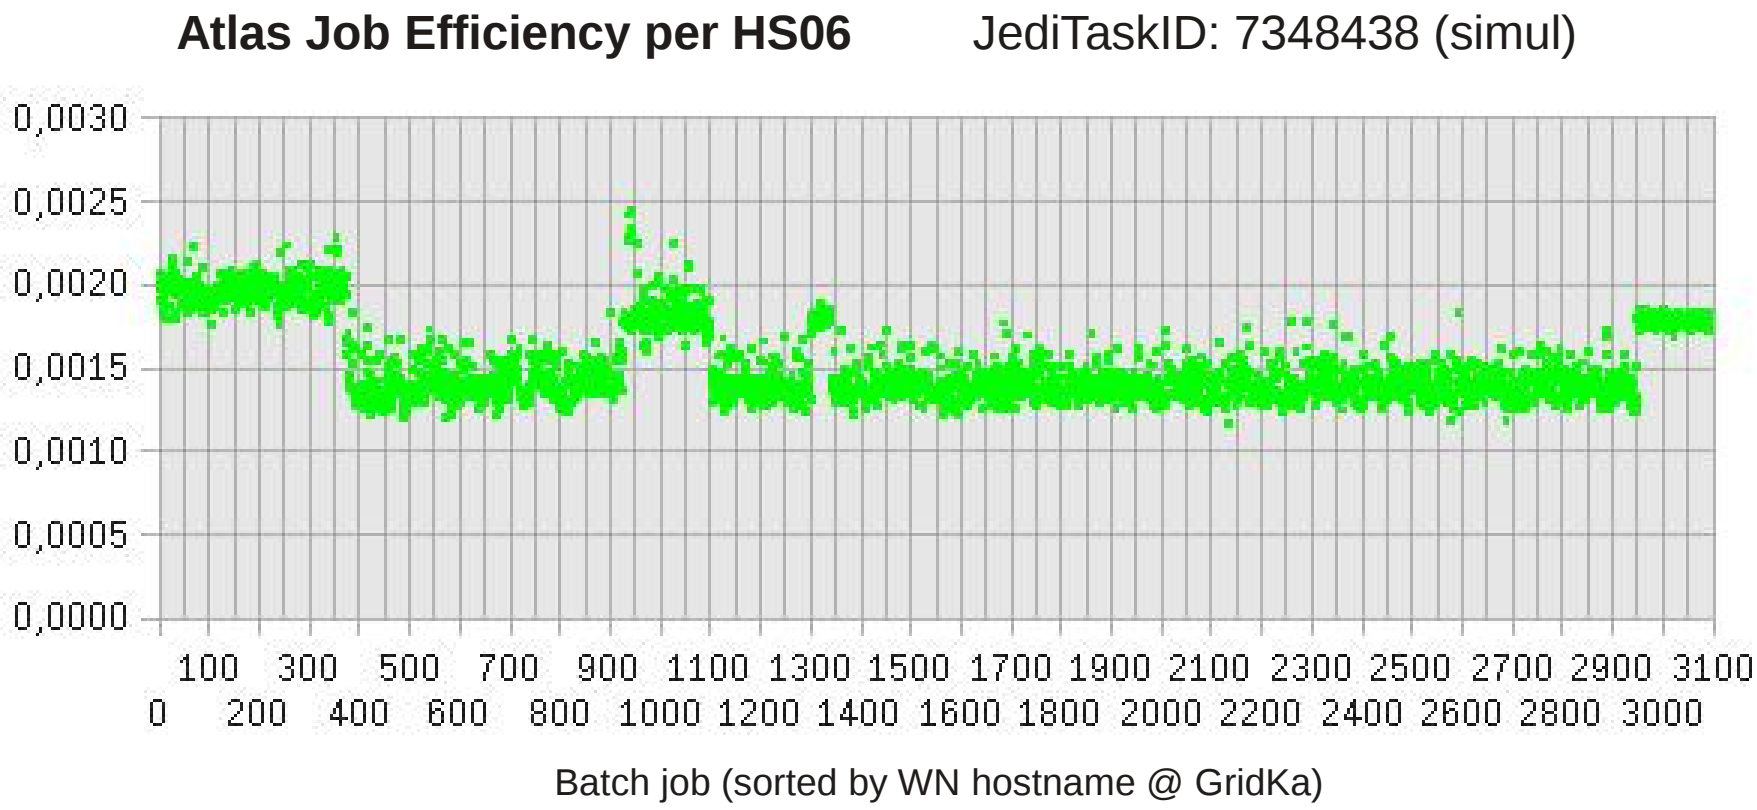

Status 2016

- ... while other Atlas processingtypes are still scaling well:

Atlas Job Efficiency per HS06

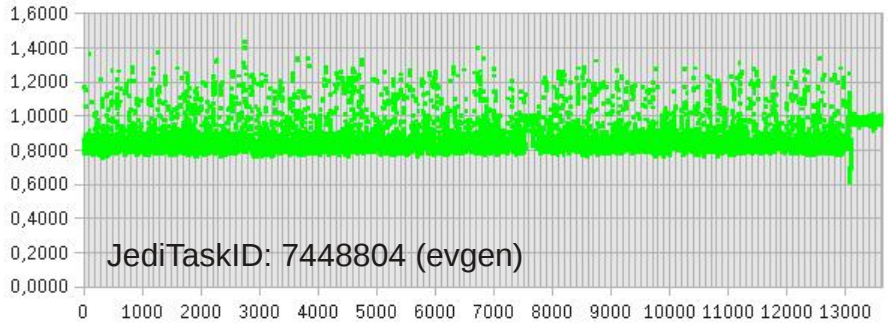

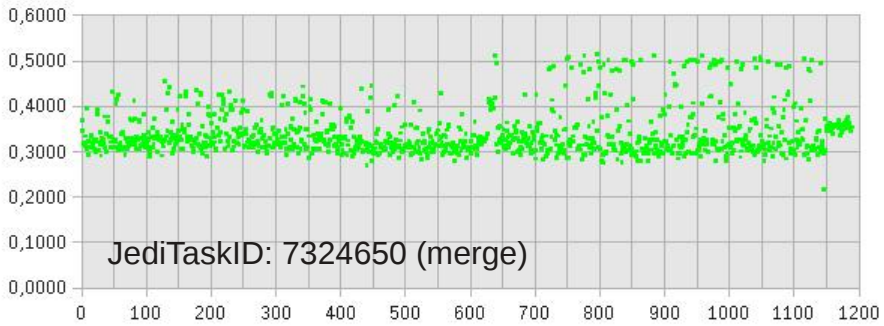

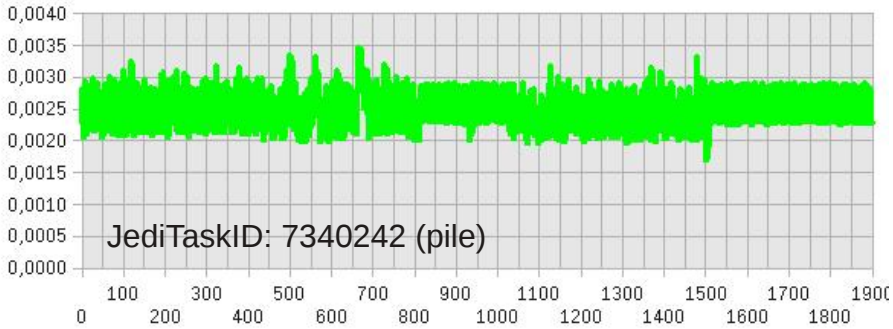

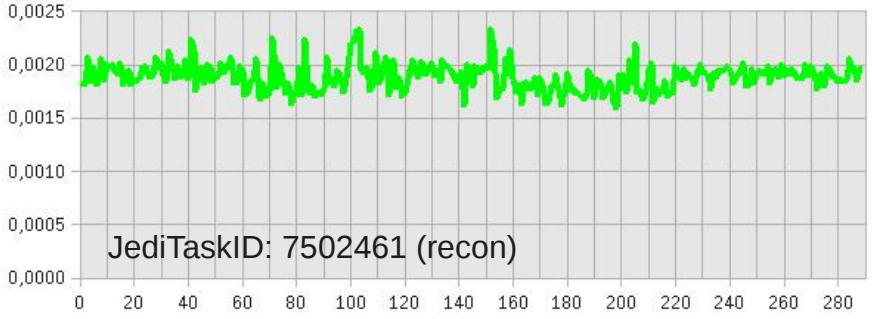

Batch job (sorted by WN hostname @ GridKa)

Status 2016

- In total:

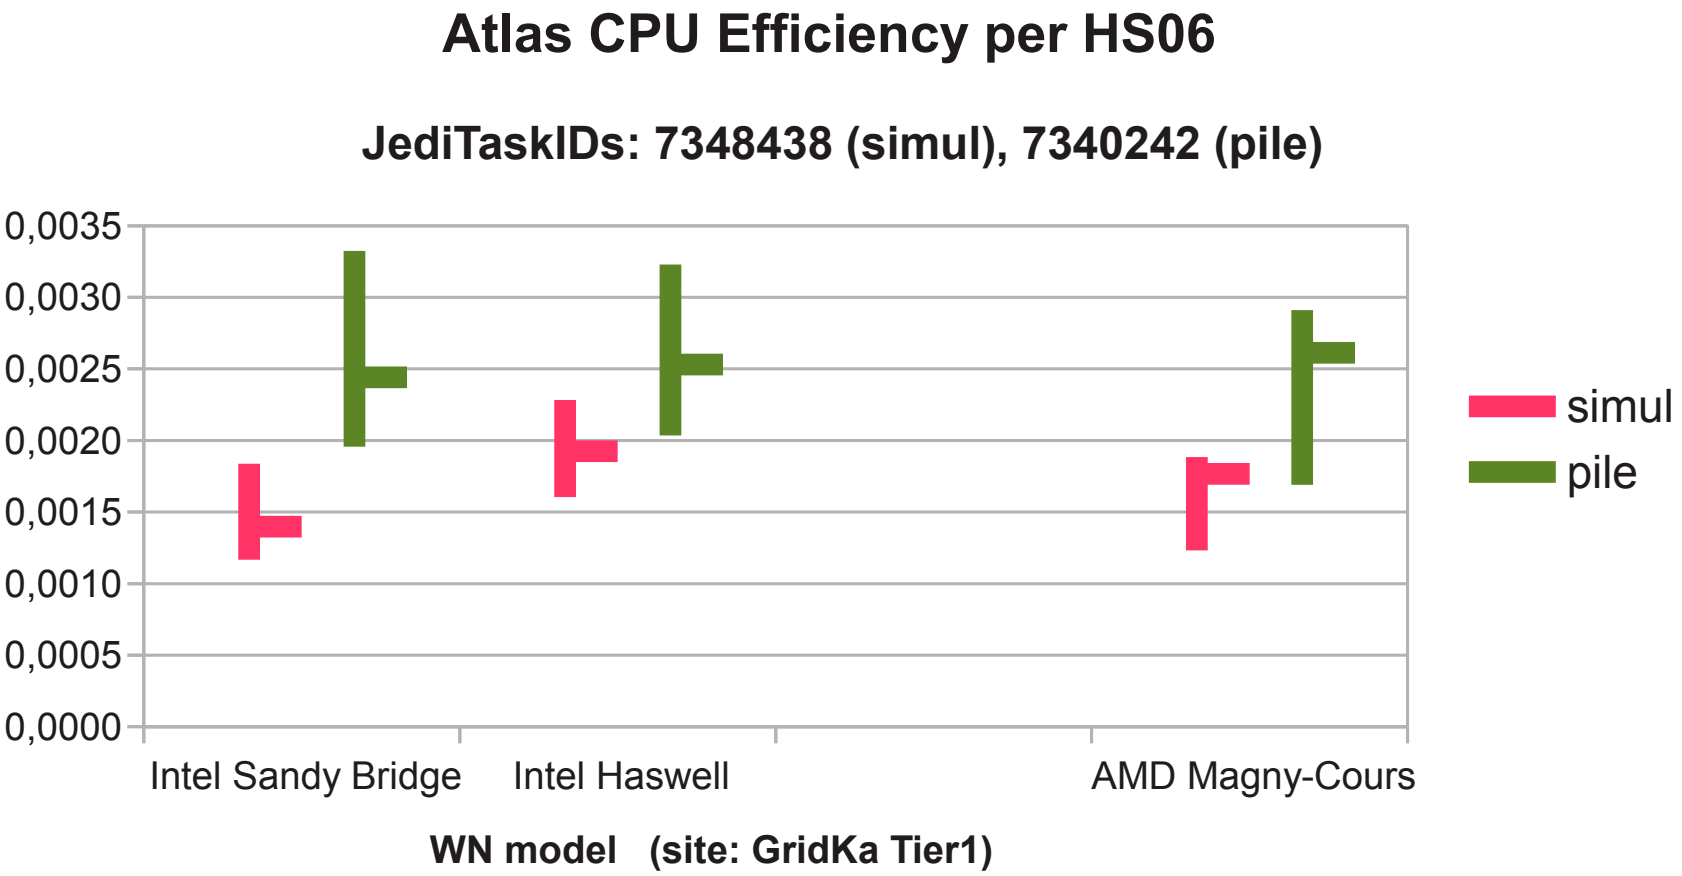

(Interpretation: min – average – max)

Status 2016

- Possible cause of magic boost of jobs landing on Haswell WN:
  - The major improvement of Haswell compared to the previous generations (Sandy Bridge, Ivy Bridge) is the new AVX2 engine (Advanced Vector Extension v2) which is able to speed up special SIMD instructions significantly, e.g.
    c(1:n) := a(1:n) + b(1:n)
  - The SL6 default compiler (GCC 4.4.x) doesn't support AVX2
  - The HS06 compiler flags (-O2 -pthread -fPIC -m32) don't enable AVX2 utilization
  - Atlas and LHCb are using more recent GCC release than 4.4

Fast Benchmarks

- LHCb Dirac fast benchmark looks like LHCb production, and Atlas simulation jobs (Haswell +35%, AMD +15%):

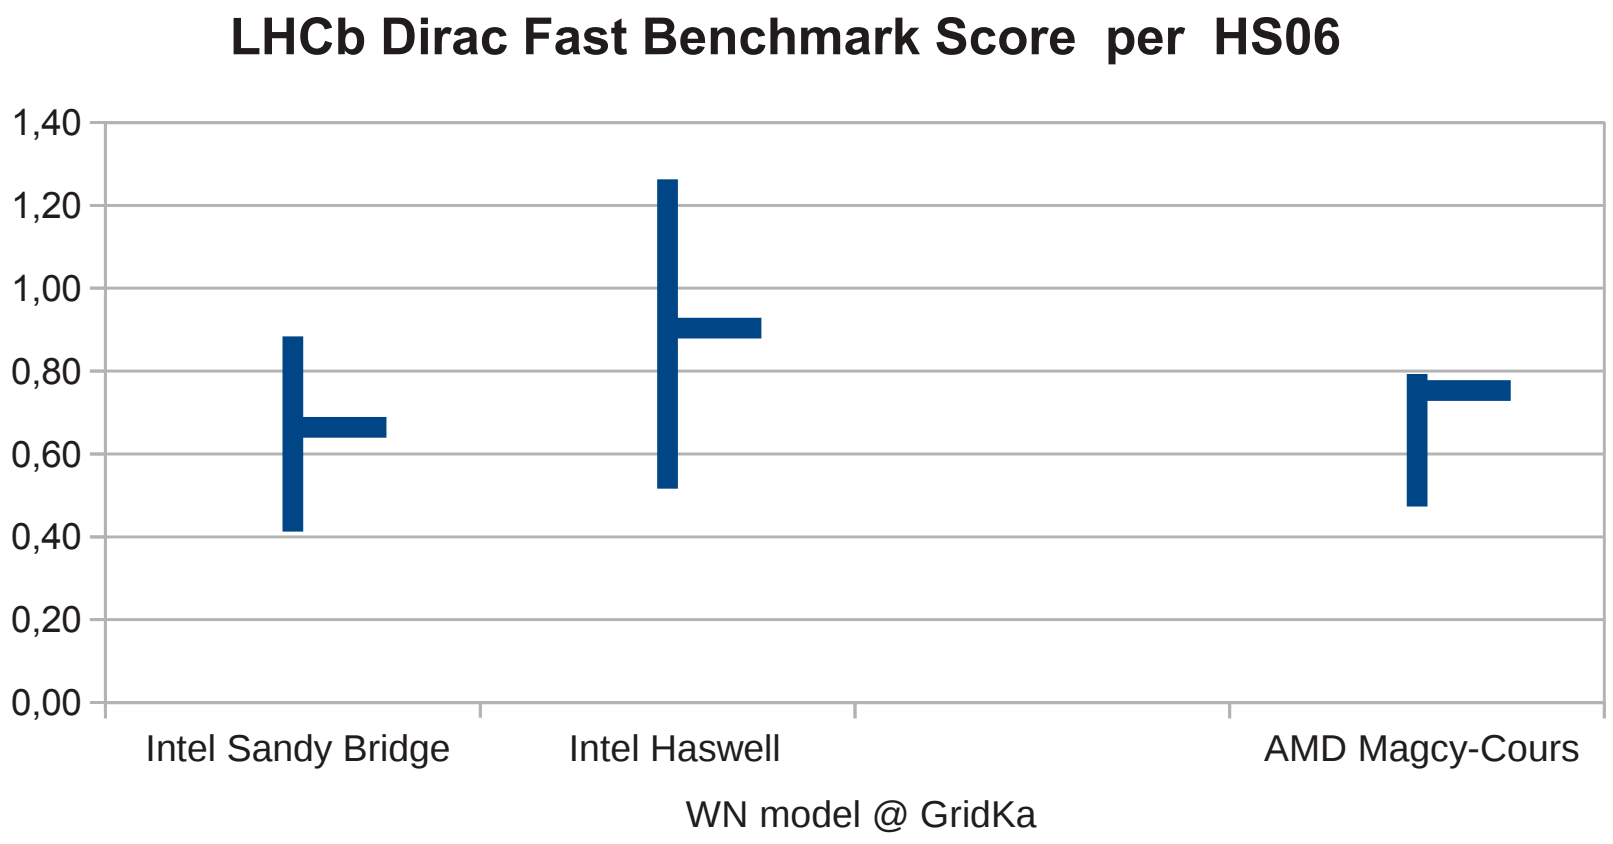

Fast Benchmarks

- Whetstone (and Dhrystone, too) still scaling with HS06:
  (UnixBench: https://github.com/kdlucas/byte-unixbench)