

Belle-II Benchmark

Tristan Sullivan,
Randall Sobie, Marcus Ebert

University of Victoria
Oct. 23/2020

Belle II at SuperKEKB



Plan to collect **50 ab⁻¹** of collisions at and near $\Upsilon(4S)$
 Successor to Belle at KEKB (1.05 ab⁻¹)


At $\Upsilon(4S)$, $E_{CM} = 10.58$ GeV
 7 GeV e^- (HER; High Energy Ring)
 4 GeV e^+ (LER; Low Energy Ring)

Nano beam scheme

$$\mathcal{L} = \frac{\gamma_{\pm}}{2e r_e} \left(1 + \frac{\sigma_y^*}{\sigma_x^*} \right) \frac{I_{\pm} \xi_{\pm y}}{\beta_y^*} \left(\frac{R_L}{R_y} \right)$$


Belle II detector

5.9 → 0.3 mm
 KEKB SuperKEKB



Physics motivations

- New physics search in B , B_s , D , τ decays
- Direct search for light new particles
- Precise measurement of Standard Model
- Hadron physics



Slide from “Belle II at SuperKEKB” by Toru Iijima, BEAUTY2020

Benchmark Under Development

- Simulation of B0/anti-B0 events, currently background-free
- Includes detector and trigger simulation, track reconstruction
- Containerized using scripts in <https://gitlab.cern.ch/hep-benchmarks/hep-workloads>
- Belle-2 specific part available at <https://github.com/TristanSullivan/Belle2Benchmark>

Compatibility with HEPscore

- Docker container that accepts number of events, copies, and threads as command-line arguments
- Saturates cores of machine by default (have both single-threaded and 4-threaded versions)
- Outputs JSON, including throughput score in total events processed per second

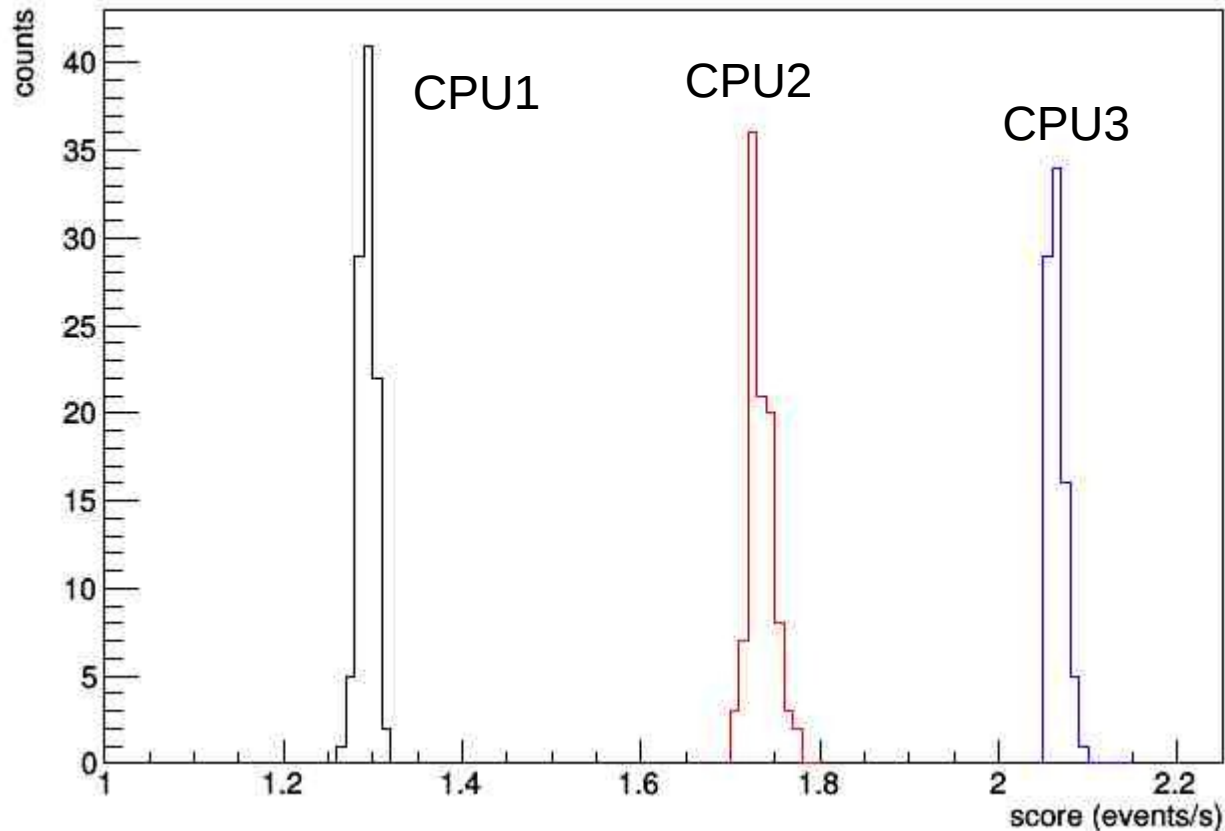
Results

The benchmark was run on a small Openstack cloud at UVic, used for testing

8 core VMs, 2 GB RAM / core

No other VMs on the hypervisors (dedicated machines)

Belle-2 Benchmark



CPU1 = Intel Xeon CPU E5-2670, 2.6 Ghz, family 6, model 45, stepping 7

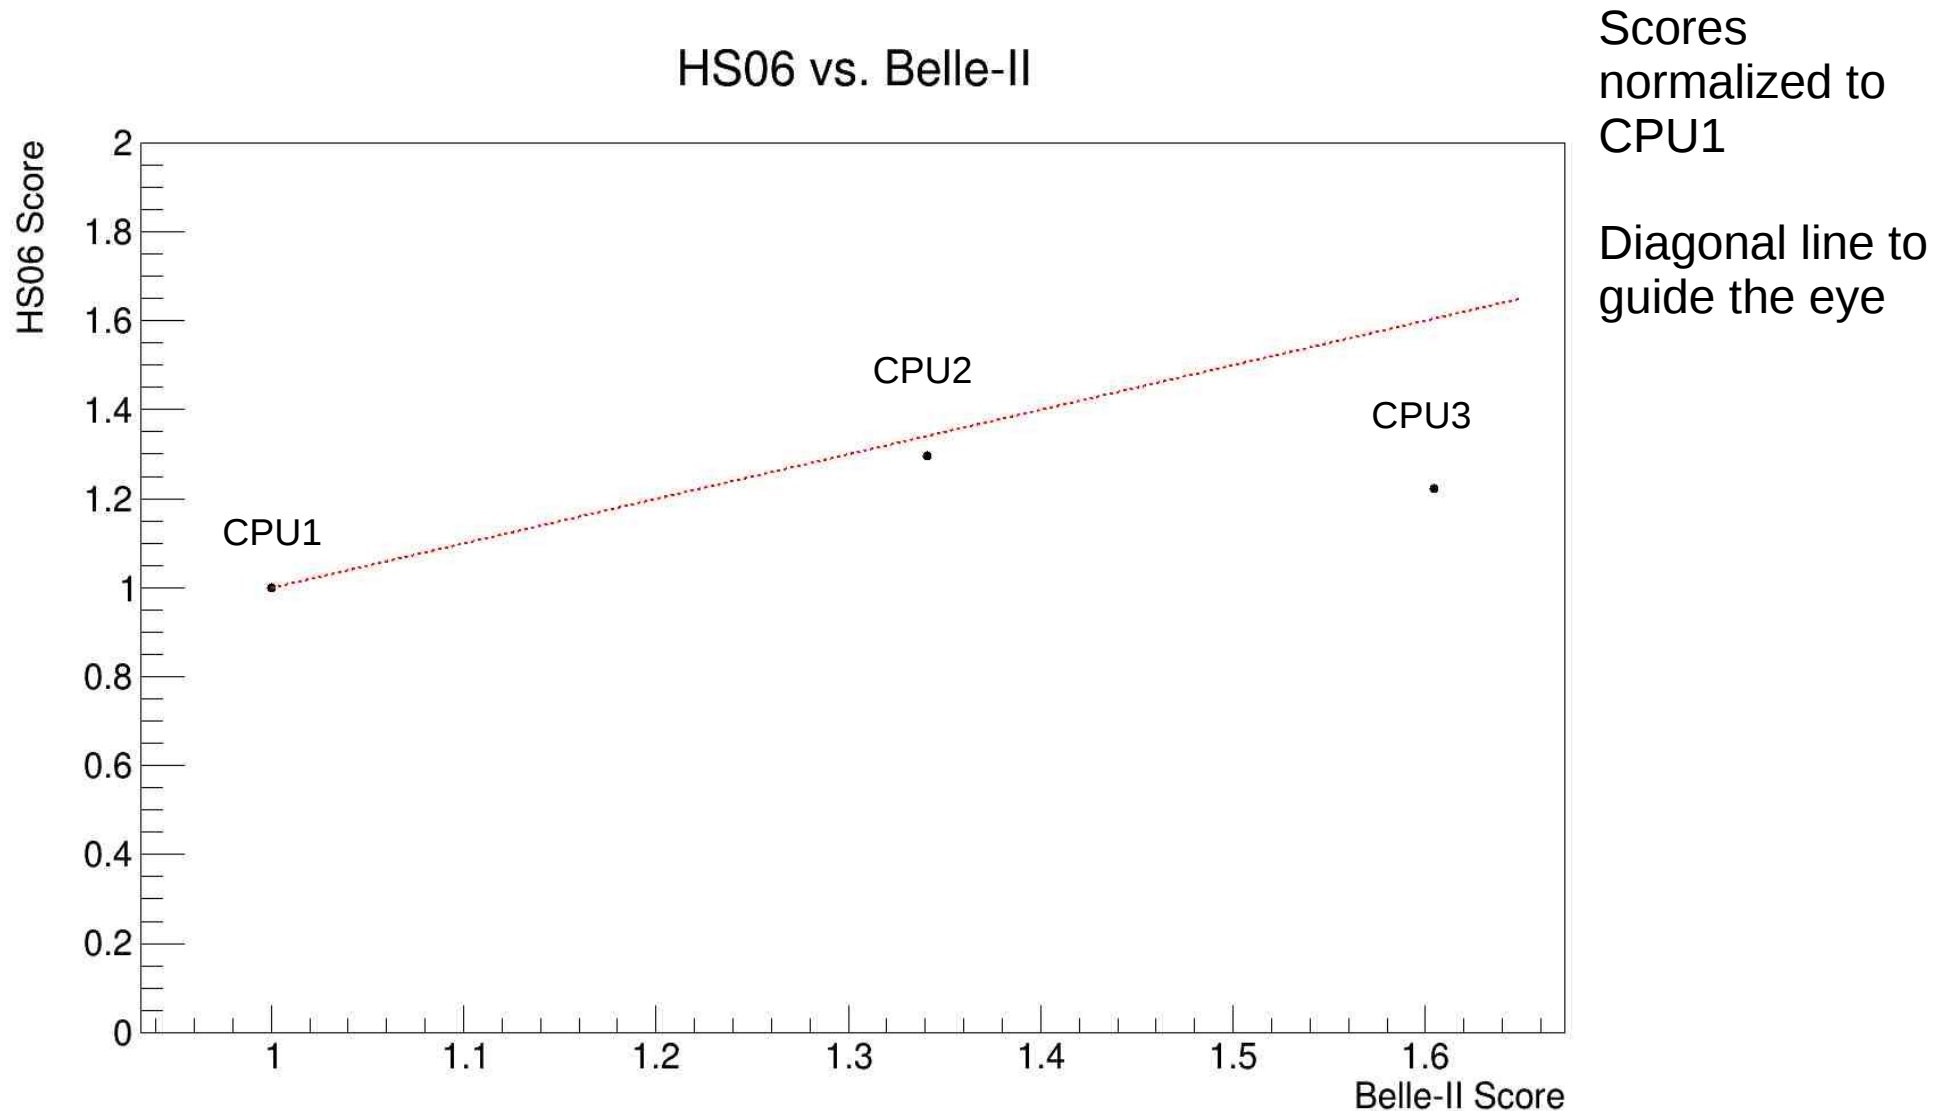
CPU2 = Intel Xeon Gold 6226 CPU, 2.7 GHz, family 6, model 85, stepping 7

CPU3 = Intel Xeon CPU E5-2687W, 3.0 GHz, family 6, model 79, stepping 1

CPU2 is the newest

CPU3 has highest clock frequency

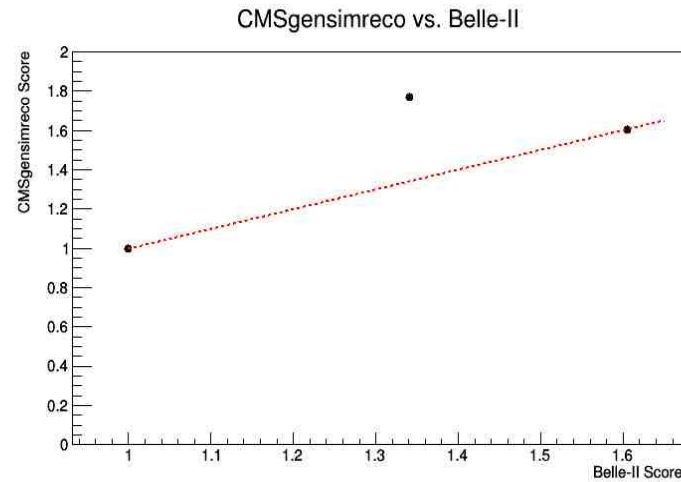
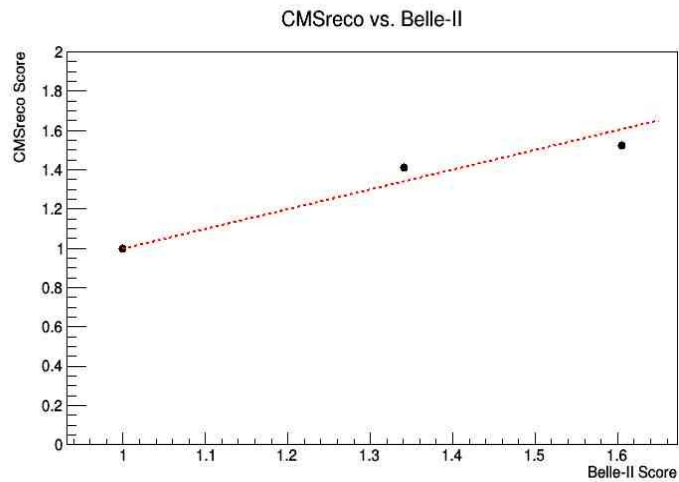
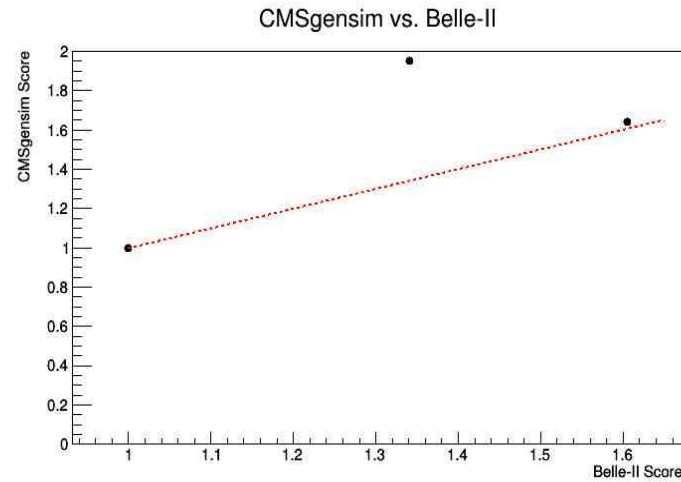
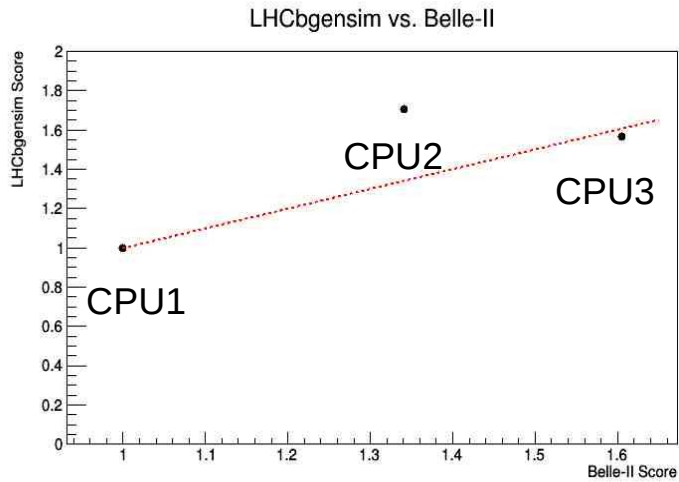
Benchmark Comparison



Benchmark Comparison

Scores
normalized to
CPU1

Diagonal lines to
guide the eye



Summary

- All results preliminary; exact choice of benchmark pending formal approval by Belle-2 collaboration (event type, whether to include background and reconstruction)
- Workload seems to behave reasonably so far; running on more CPU types would clarify scaling relative to the other workloads
- Ready to test integration into HEPscore?

Supplementary Material

Detailed Machine Information

- CPU1 HV:

- 32 cores with HT on
- Two numa nodes
- 10 x 8192 MB RAM, DDR3, 1600 MT/s
- L1d cache: 32 K
- L1i cache: 32K
- L2 cache: 256 K
- L3 cache: 20480 K

- CPU2 HV:

- 48 cores with HT on
- Two numa nodes
- 6 x 16384 MB RAM, DDR4, 2933 MT/s
- L1d cache: 32 K
- L1i cache: 32K
- L2 cache: 1024 K
- L3 cache: 19712 K

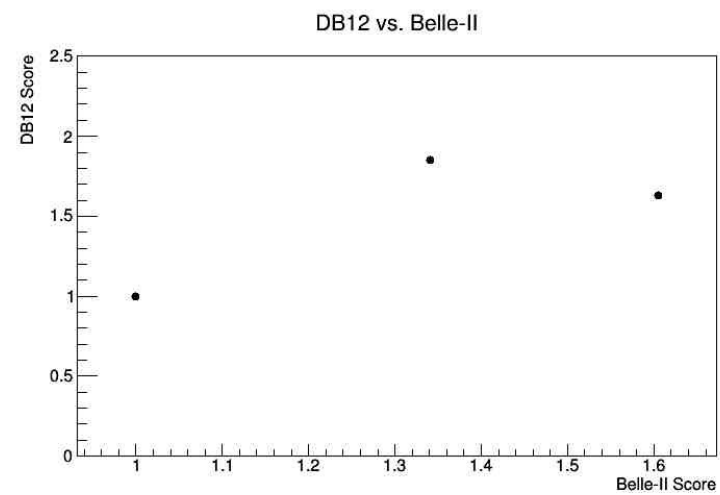
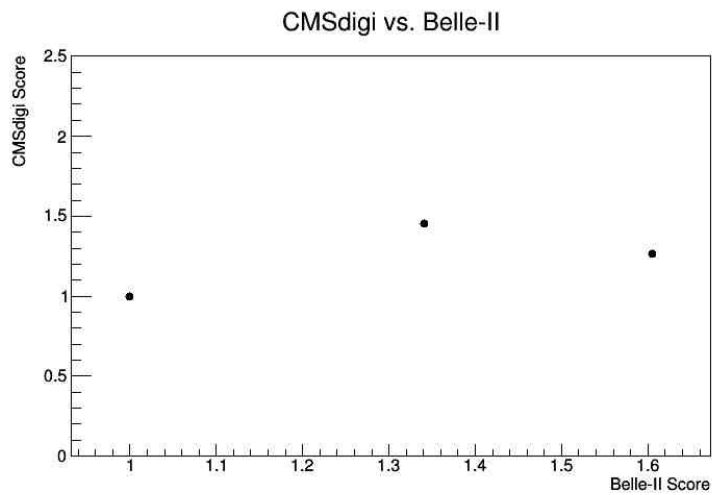
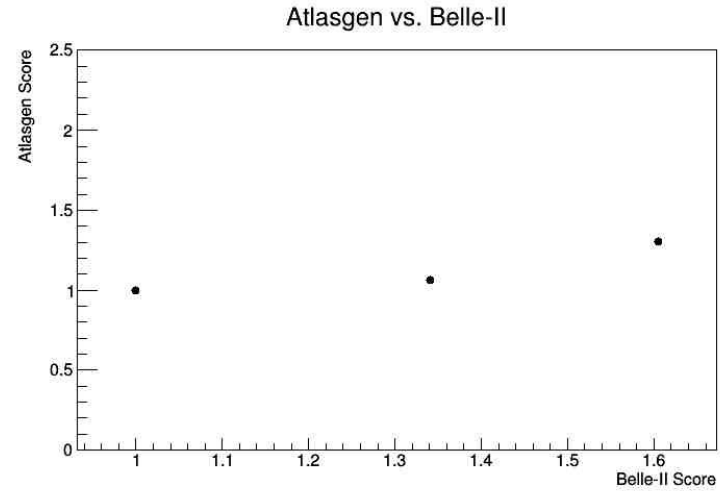
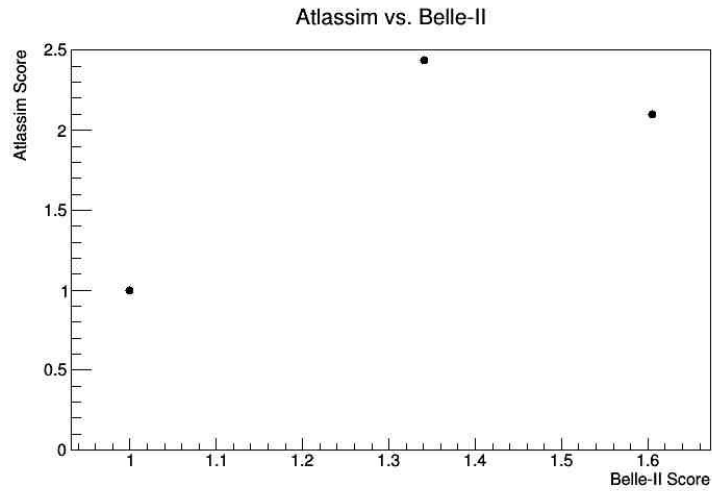
- CPU3 HV:

- 48 cores with HT on
- Two numa nodes
- 6 x 16384, 2 x 8192 MB RAM, DDR4, 2400 MT/s
- L1d cache: 32 K
- L1i cache: 32K
- L2 cache: 256 K
- L3 cache: 30720 K

Detailed Benchmark Information

- Atlas-gen: 200 events, 1 thread, 10 minutes
- Atlas-sim: 10 events, 4 threads, ~90 minutes
- LHCb-gen-sim: 5 events, 1 thread, ~25 minutes
- CMS-reco: 50 events, 4 threads, ~10 minutes
- CMS-digi: 50 events, 4 threads, ~5 minutes
- CMS-gen-sim: 20 events, 4 threads, ~15 minutes
- Belle2-gen-sim-reco: 20 events, 1 thread, ~5 minutes
- Times are for CPU1
- Threads * copies = cores of machine (all cores used by default)
- Benchmark score is total events processed per second (sum of individual scores)
- Spreadsheet with scores:
<https://docs.google.com/spreadsheets/d/1TN6xaVarEcQ6LBart505fHQ3HnRrcWerKofi5bKsA08/edit#gid=0>

Benchmark Comparison



Runs Without Dedicated HVs

Atlas Sim vs. Atlas Gen

