

Trigger and reconstruction

Stewart Martin-Haugh (RAL ATLAS), Chris Jones (Cambridge LHCb)

Efficient Computing in High Energy Physics
24 March 2020



Science and
Technology
Facilities Council

Benchmarks

- ▶ Set of benchmarks for each major experiment
- ▶ Let's say ATLAS, CMS, LHCb - DUNE?
- ▶ HLT and offline reconstruction (if experiment makes a distinction)
- ▶ Some kind of common format (Docker container?)
- ▶ Agree on number of threads, concurrent events

Low-hanging CPU fruit

- ▶ Nothing too low
- ▶ Could easily run benchmarks with preloaded libraries (malloc, cmath replacements)
 - ▶ ATLAS uses tcmalloc and IMF - 10% improvement for each
 - ▶ Useful for other experiments?
- ▶ Try to restart efforts around memory profiling, heap analysis
- ▶ O2,O3 - certainly ATLAS uses O2
- ▶ Link-time optimisation: SMH tried this again after Edinburgh but still hitting bizarre compiler/linker errors
 - ▶ Contact [GCC expert?](#)

Ongoing work

- ▶ Manuel Schiller has been working on the [optimized C++ version](#) of an NN classifier
- ▶ Next steps: non-vectorized version, and more aggressive vectorization along with some benchmarks

GPU benchmarks

- ▶ ATLAS has examples for [CUDA](#) and [SYCL](#)
 - ▶ Annoying CUDA detail: nvcc only supports C++14, ATLAS/Gaudi is on C++17
- ▶ Compare with CMS CUDA example?
- ▶ Compare with LHCb Allen project?

FPGAs, tracking with ML, use of timing detectors

- ▶ May be challenging to get anything serious up and running in the timescale we have
- ▶ Should be able to summarise work in these areas across experiments