Trigger and reconstruction

Stewart Martin-Haugh (RAL)

Efficient Computing in High Energy Physics
18 February 2020

Science and
Technology
Facilities Council

Reconstruction

- Scope for novel reconstruction algorithms?
- Combining ML with traditional reconstruction is interesting
- What standalone reconstruction kits can we use?

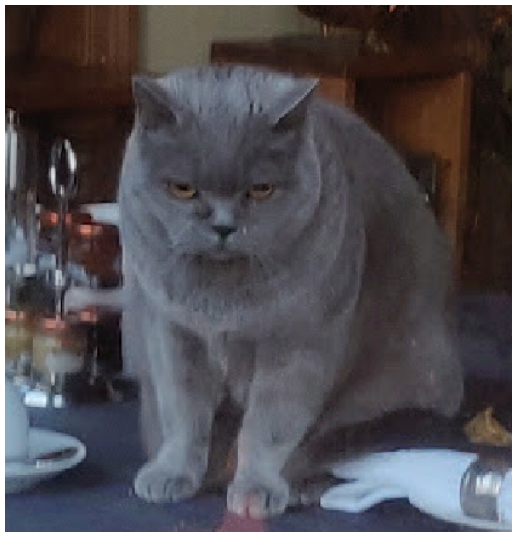

GPUs

- Event reconstruction is embarrasingly parallel
  - Two CPUs = 2x throughput
  - CPU + GPU = ?x throughput
- For HLT farms or grid site procurement, want $\approx$ 10x speedup
  - Can we establish a cost metric/model
- For some existing resources (USHPC), $\approx$ 1.5x speedup may be acceptable (better than not using it)
- Gain experience even if we're not using it optimally

Coarsening reconstruction

- Store less data (cf TURBOTLADS[1])
- Lossy compression of raw data (e.g. with autoencoders)
- Numerical precision
  - Double, FP32, FP16,FP8,FP4
  - Fixed point?

[1] Turbo stream, Trigger level analysis, Data scouting

What can we do in 6 months?

- Run ML directly on raw data for anomaly detection
- A simple algorithm implemented in CUDA, SYCL, Kokkos etc as a testbed
- Establish who in the UK is working on non-x86 architectures
- Make contact with other communities working on fast pattern recognition/decision making