LHCb GPU Computing status report

THE MANYCORE AWAKENS

Trigger refurbish in 2020

LHCb is going for a full software trigger in 2020

- The detector will produce 30M events per second
- Each server will have to sustain 10.000 events per second
  - 20x more than in the current Run2

Direct consequences

- We are evaluating other hardware architectures
- Architecture-aware programming becomes more relevant
- Framework will require deep structural changes

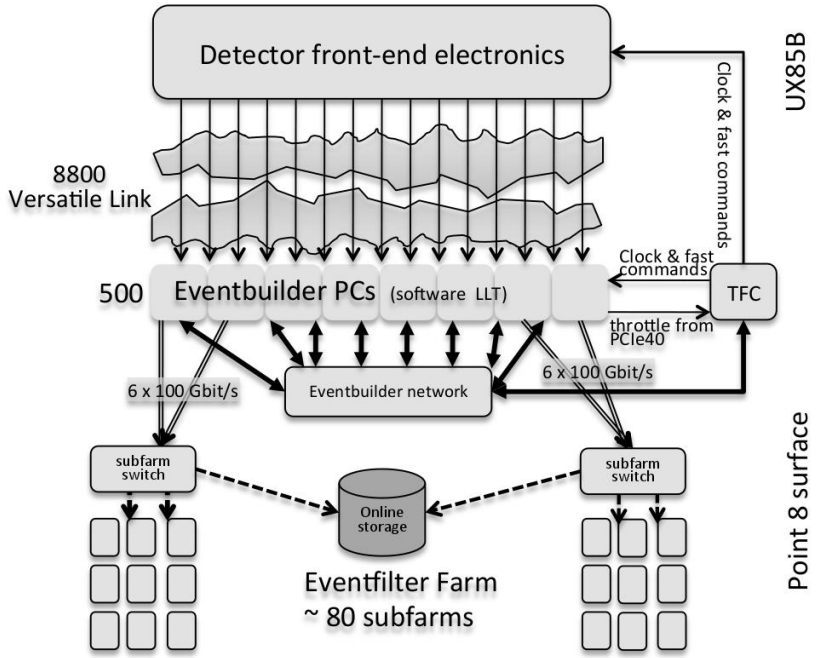

The landscape of processors

Architecture choice is not anymore a no-brainer:

- x86_64, OpenPower, ARM64, GPUs, KNL

The transistors count is not a free meal anymore

- Xeon Phi, KNL: 60+ cores
- NVIDIA: 2000+ (CUDA) cores

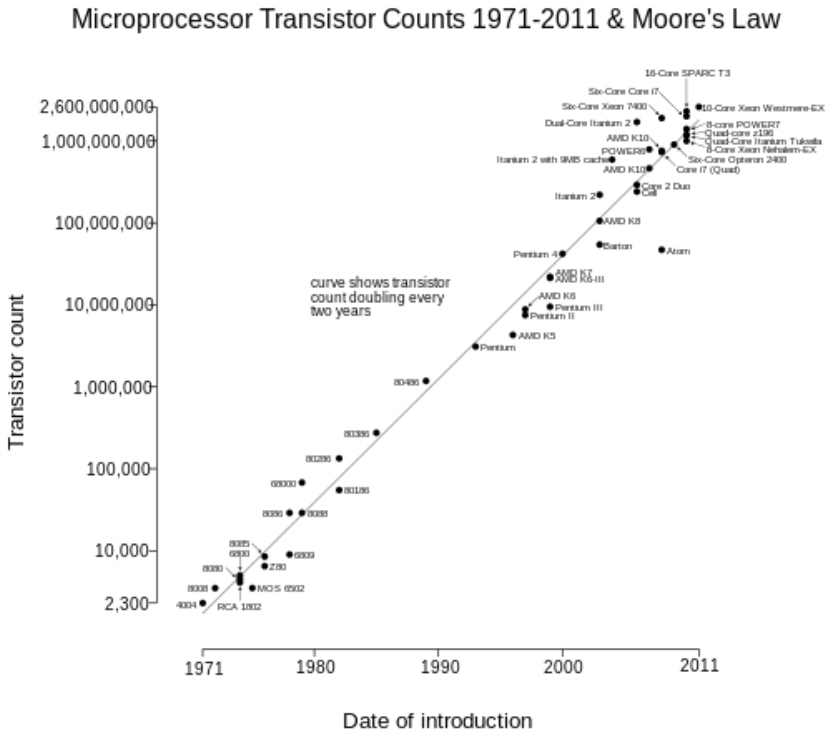

"Transistor Count and Moore's Law - 2011" by Wgsimon - Wikipedia

Our reality at present

Software is not fully harnessing current hardware

- x86-centric codebase
- Optimistic Upgrade predictions for 20XX

The usual suspects in software

- Data locality, cache
- AOS vs SOA
- Multi-threading support

This is intrinsically correlated with Hardware

- Memory hierarchy
- SIMD
- Core count

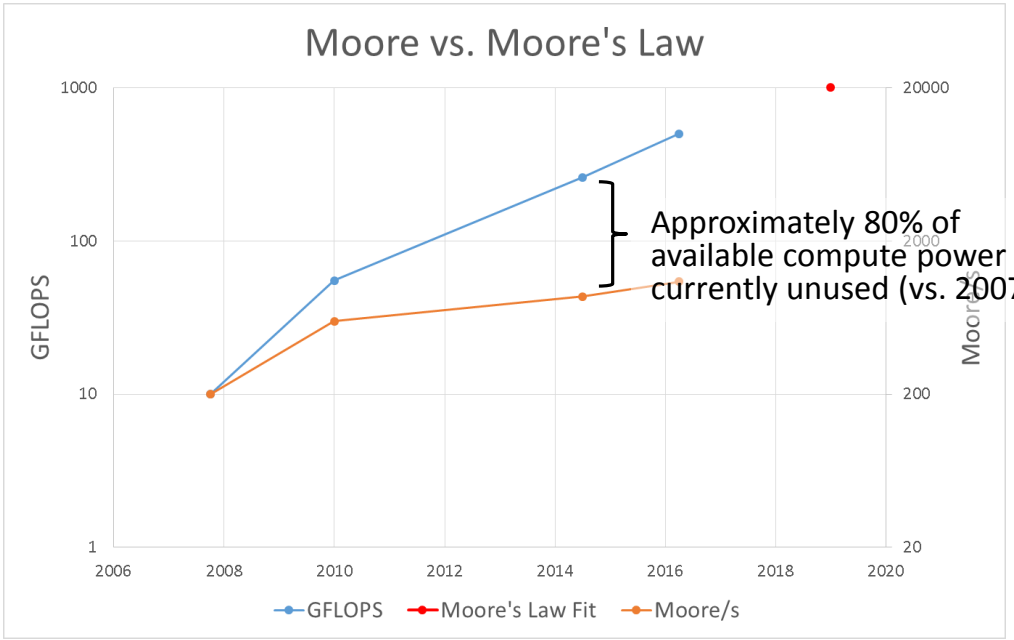

Projects related to GPU Computing at LHCb

Completed
- Velopix reconstruction
- GPU server integration
- Coprocessor manager

Ongoing
- Velopix on OpenMP
- Forward tracker on automata

Starting 2016
- lhcb-parallelization wg
  - Upgrade forward tracker
  - Kalman filter
  - PID analysis
    - RICH reconstruction
    - ProbNN

Velopix on OpenCL

- Track forwarding
  - Find track seeds and *forward* them on each sensor consecutively
- Hough transform
  - Represent all possible lines crossing a point and find intersections
- Retina algorithm
  - Recognise preexisting patterns, match hits against DB
- Other approaches
  - *Automata* approach
  - *dxdy* and segment extrapolation matching

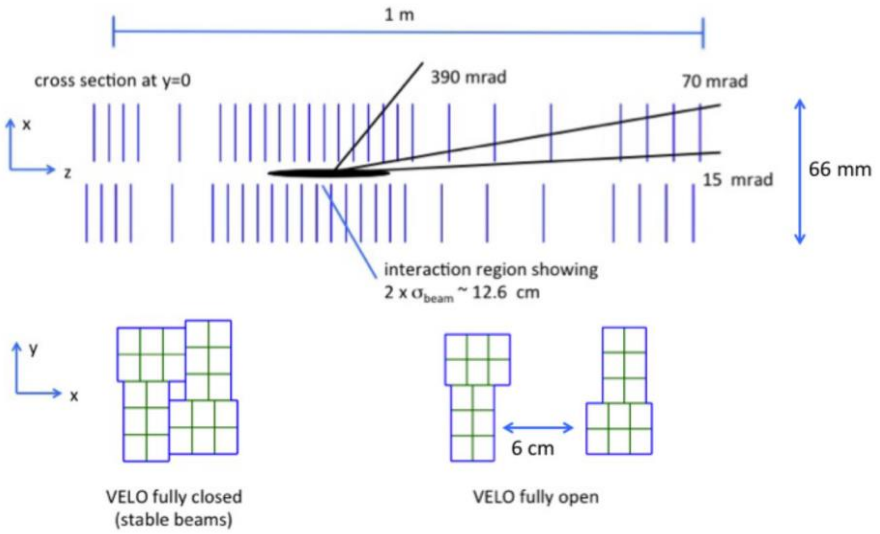

Velopix results

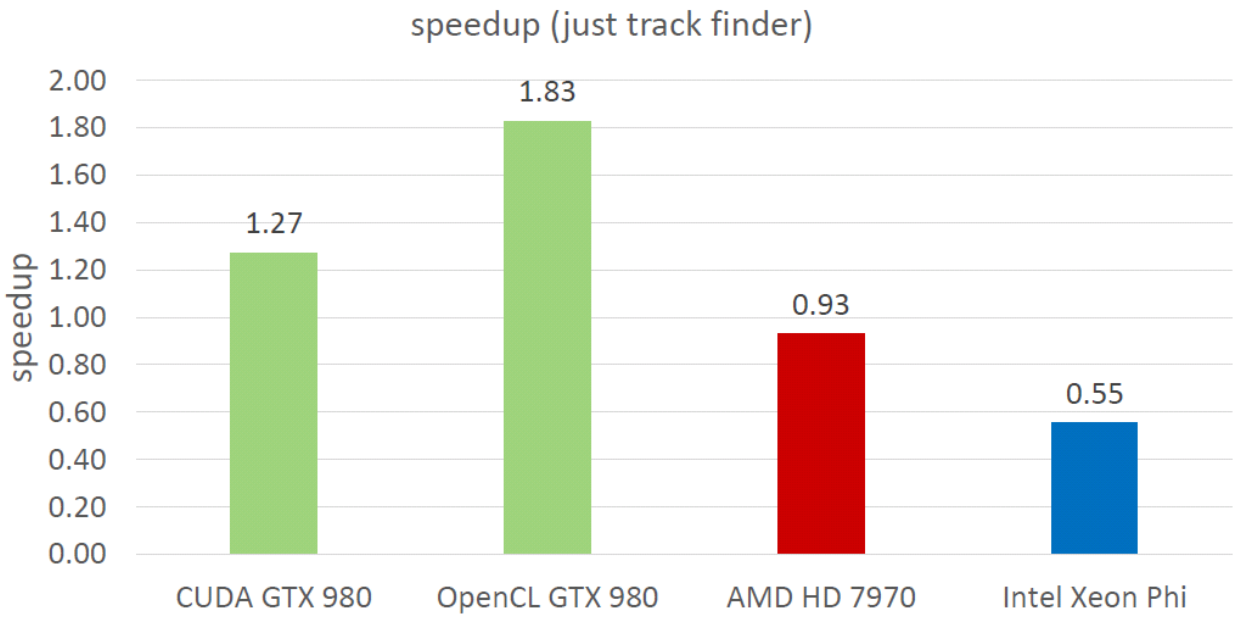

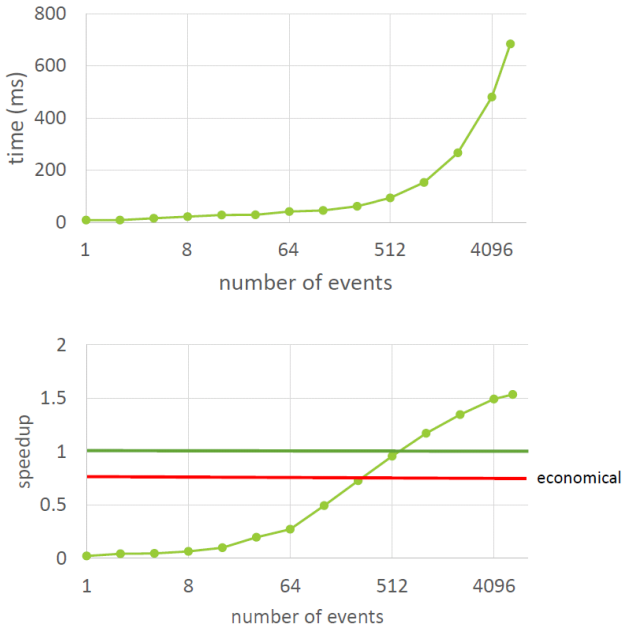

- Sequential is run on a fully loaded one Intel® Xeon® CPU E5/2630 v3 @ 2.40 GHz
- Comparison is done for throughput – *events processed per second*.
- Note the sequential version is vectorised.
- Results on the right are for the OpenCL GTX 980 implementation.

Coprocessor manager

The coprocessor manager allows accelerator hardware usage for massively parallel computation within LHCb's framework.

Some things the coprocessor manager does

- Automatically batches data
- Schedules algorithms
- Can be used across a network
- Can run GPU and CPU algorithms side by side
- Helps out in algorithm development

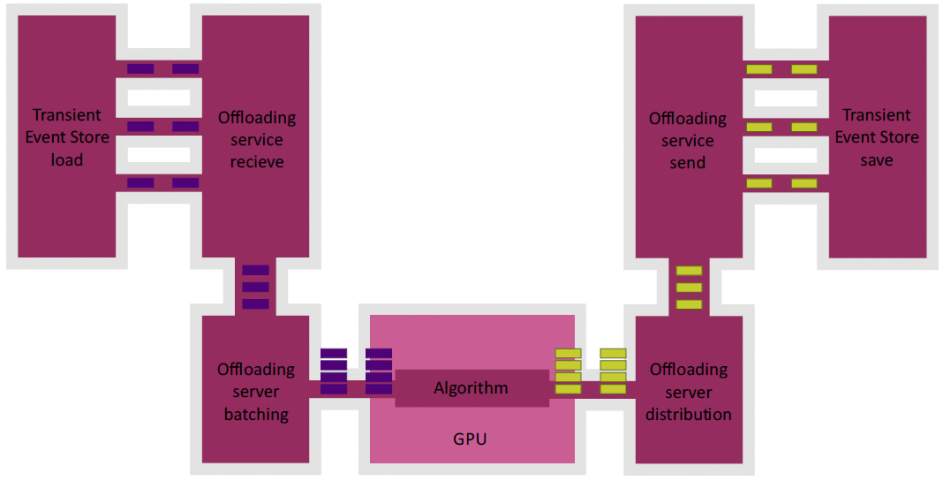

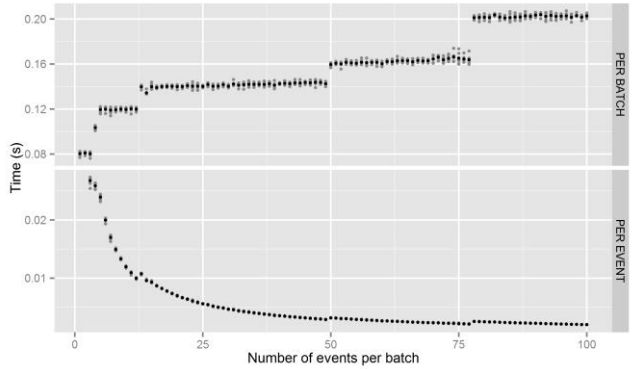

GPU server integration

- Proof of concept GPU server integration
- CPU HLT chain, reconstructing events with GPU-based tracking

Some early conclusions

- Software chain is not prepared for seamless integration, many tweaks were required
- External hardware is integrated through coprocessor manager

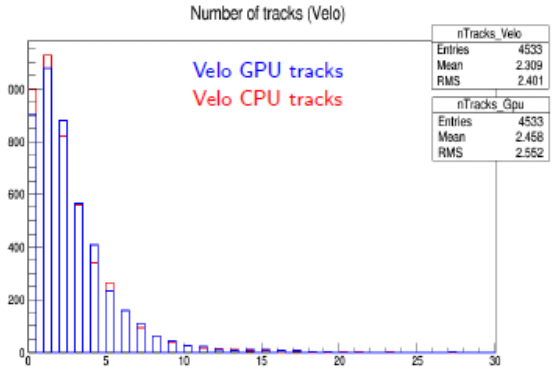

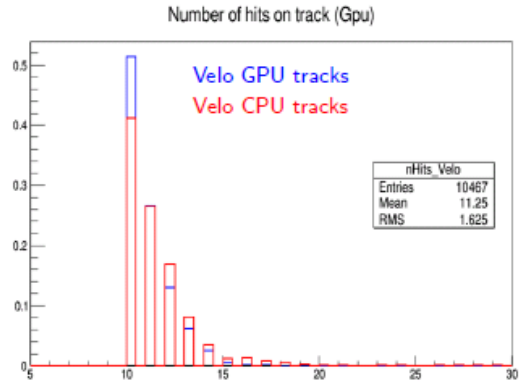

Forward tracking automata implementation

Adapted CA algorithm:

- Automata track finding procedure: tracklet generation – neighbour finding – evolution and track formation
- Do not build “long tracklets” – This reduces combinatorial
- Build tracklets between x-layers
- Add stereo hits to x-tracklets and form a “stub”
- Fit and kill bad tracklets with a cut on chi2
- Evolve and select candidate tracks

Reconstruction efficiencies and ghost rate are reasonable.

Next steps:

- Improve efficiencies and CPU timing
- Study performance on manycore architectures

Parallelization is on the rise

The LHCb community is becoming more aware of the challenges software poses for our Upgrade

- By end of 2016 we want an evaluation of most relevant software algorithms and hardware architectures
  - Vectorisation, OpenMP, CUDA, OpenCL, [insert programming model]
  - x86_64, NVIDIA, KNL, OpenPower, ARM64
- Based on this study, we will determine what to go for and factor in other as relevant considerations
  - **Scalability**, readability, integration, maintainability
  - Hardware failure rate, **training**, manpower
- Good quality software should not *only* do Physics!

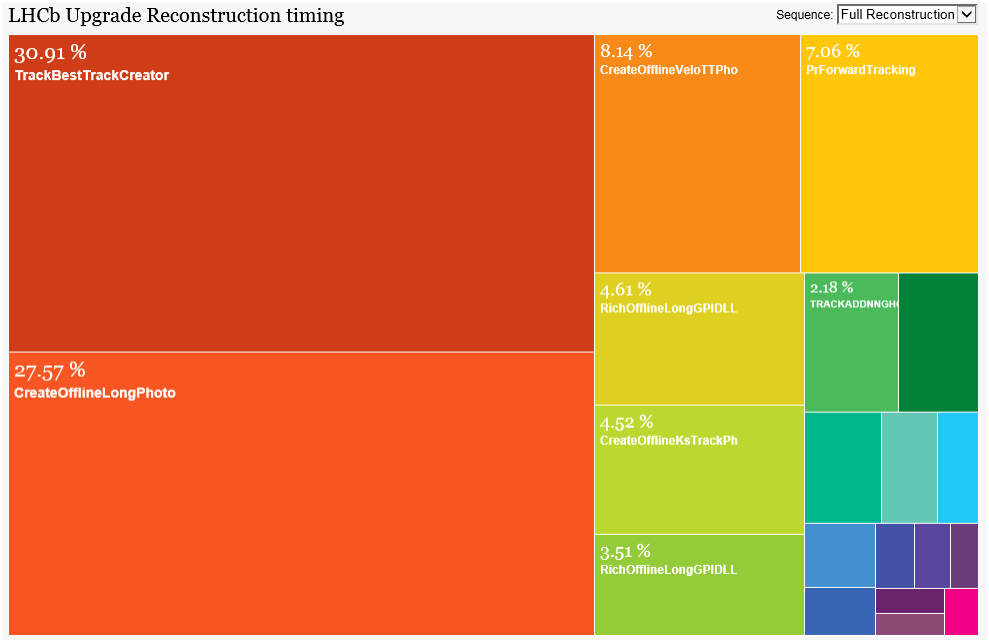

Resources here and there

- LHCb Upgrade dokuwiki
  - https://lbdokuwiki.cern.ch/upgrade:main_page

- 6th LHCb Computing workshop
  - http://indico.cern.ch/event/337568/session/5/?slotId=0#20151116