Multi-Threaded Algorithms for GPGPUs in the ATLAS High Level Trigger

Patricia Conde Muíño
on behalf of the ATLAS Collaboration

PROGRAMA OPERACIONAL POTENCIAL HUMANO
INVESTIGADOR FCT
Fundação para a Ciência e a Tecnologia
MINISTÉRIO DA EDUCAÇÃO E CIÊNCIA

ATLAS TDAQ System in Run 3

- Higher luminosity → larger pile-up → larger volume of data

  Advanced algorithms needed to provide the same rejection

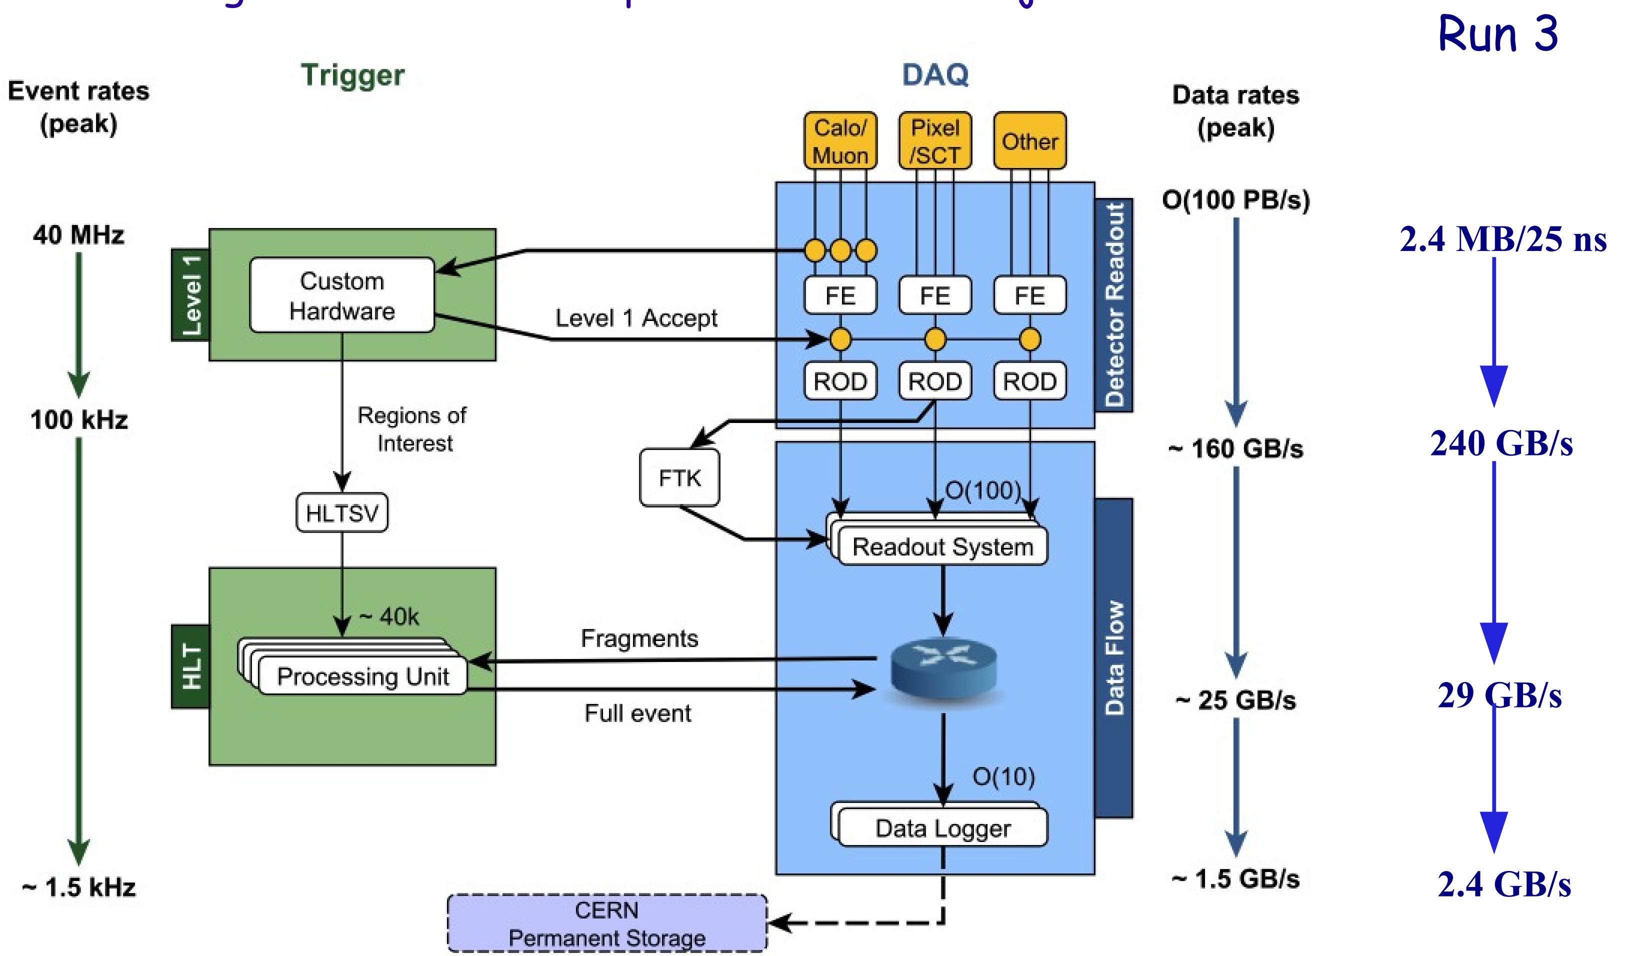

General Purpose GPUs for triggering

- HLT farm size is limited, mainly by power and cooling
- CPU time increases with pile-up
  - Dominated by combinatorial nature of the tracking algorithms
- GPGPUs: provide massive parallelization potential

ATLAS GPGPU prototype

- Evaluate the use of GPGPUs at trigger level
  - Figure of merit: processed events/s/cost
- Algorithms:
  - Calorimeter, tracking, muon and jet reconstruction
- Hardware: NVidia Tesla K80
  - 12 GB RAM
  - 2496 CUDA Cores per chip
  - 824 MHz GPU, 2505 MHz memory clock

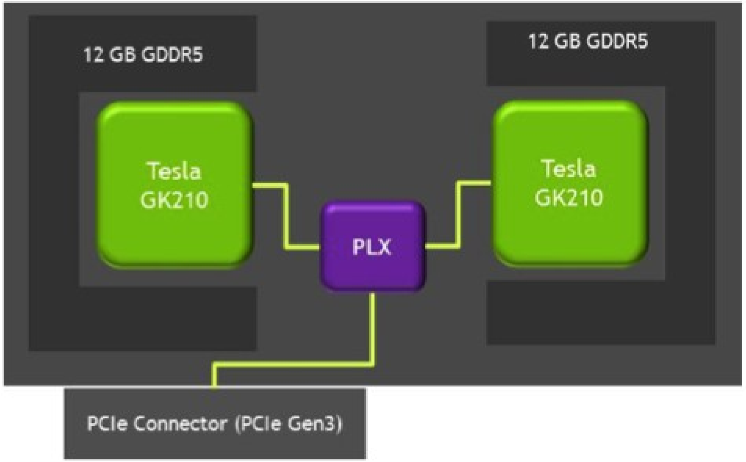

Trigger GPGPU demonstrator architecture

- Client-server architecture:

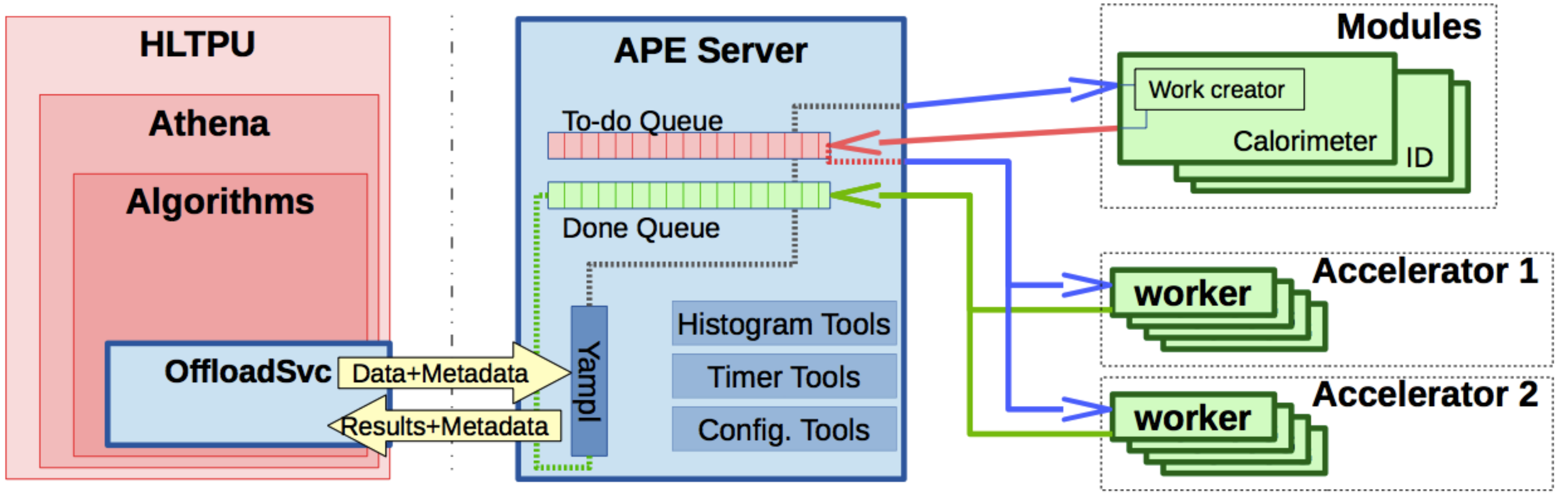

Client:

- One HLT processing unit per core
- Athena offline & online framework
  - Provides data
  - Executes chains of algorithms
  - Provides monitoring services
  - Rejects/accepts the events

Server:

- Independent from Athena
- Accelerator resource management
  - Serve many Athena processes
  - Can exploit several technologies
- Pre-allocate memory for data
- Store global/constant data

GPU Inner Detector algorithms

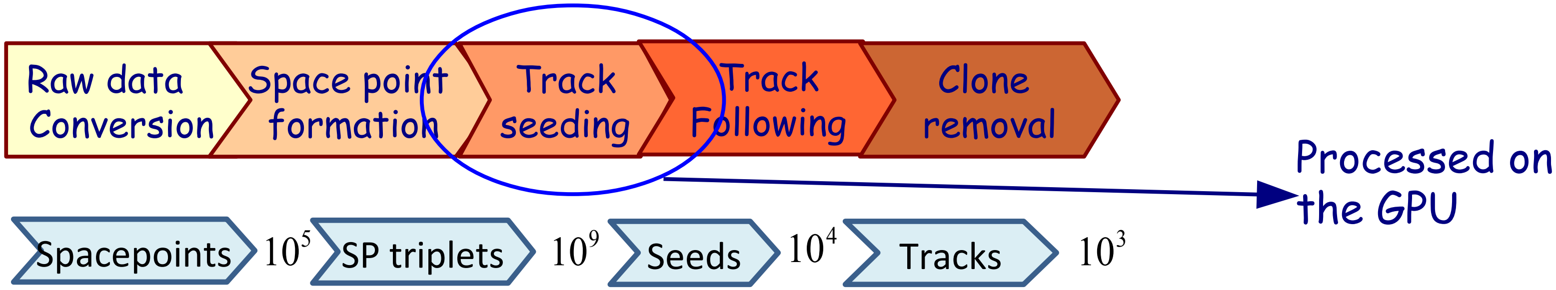

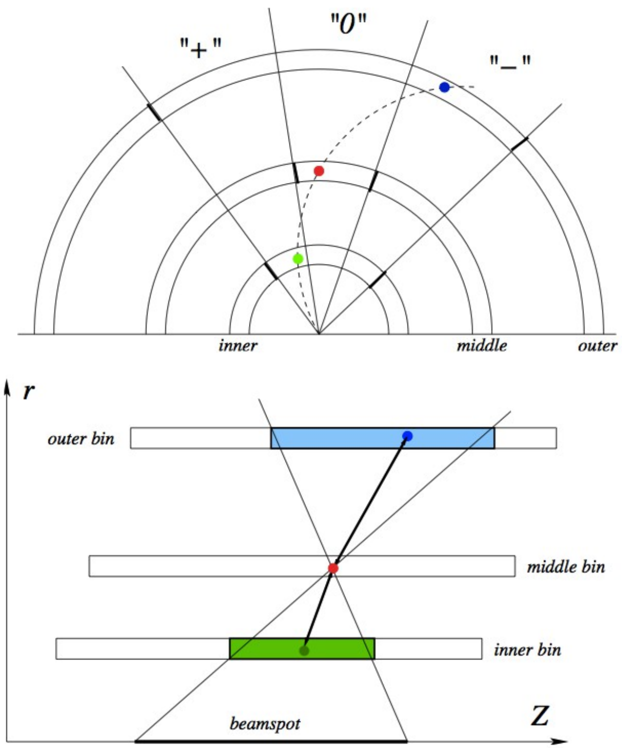

- Spacepoints (SP) arranged in wedge-shaped slices
- For each SP in middle layer
  - Search for inner/outerSP in same and neighbouring slices
  - Impose z-limits from beam spot
- GPU kernels
  - Calculate all doublets with middle-SPs
  - Store doublets in global SoA storage
  - Form triplets
    - Apply kinematic and quality cuts

GPGPU ID Algorithm Performance

- GPGPU algorithm provides same efficiency and resolution as CPU one
  (very different GPU and CPU algorithms)

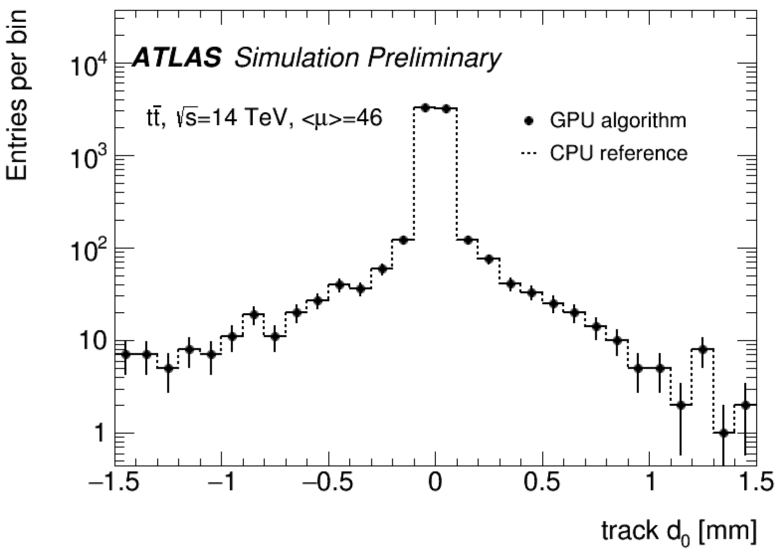

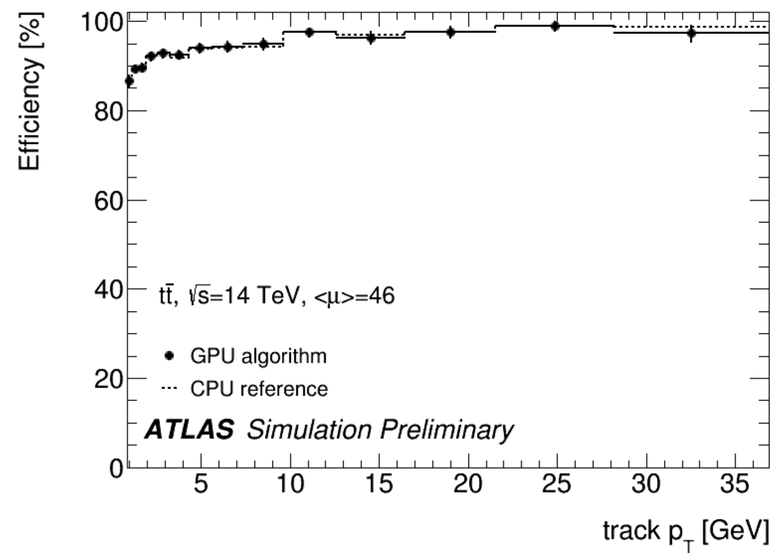

Link to GPU trigger public plots

GPGPU ID Algorithm Timing Performance

- Total algorithm execution time
  - Measured in tt-bar events (very busy) with 46 interactions/bunch crossing
  - Reduced by a factor ~5
- Very small data transfer overhead

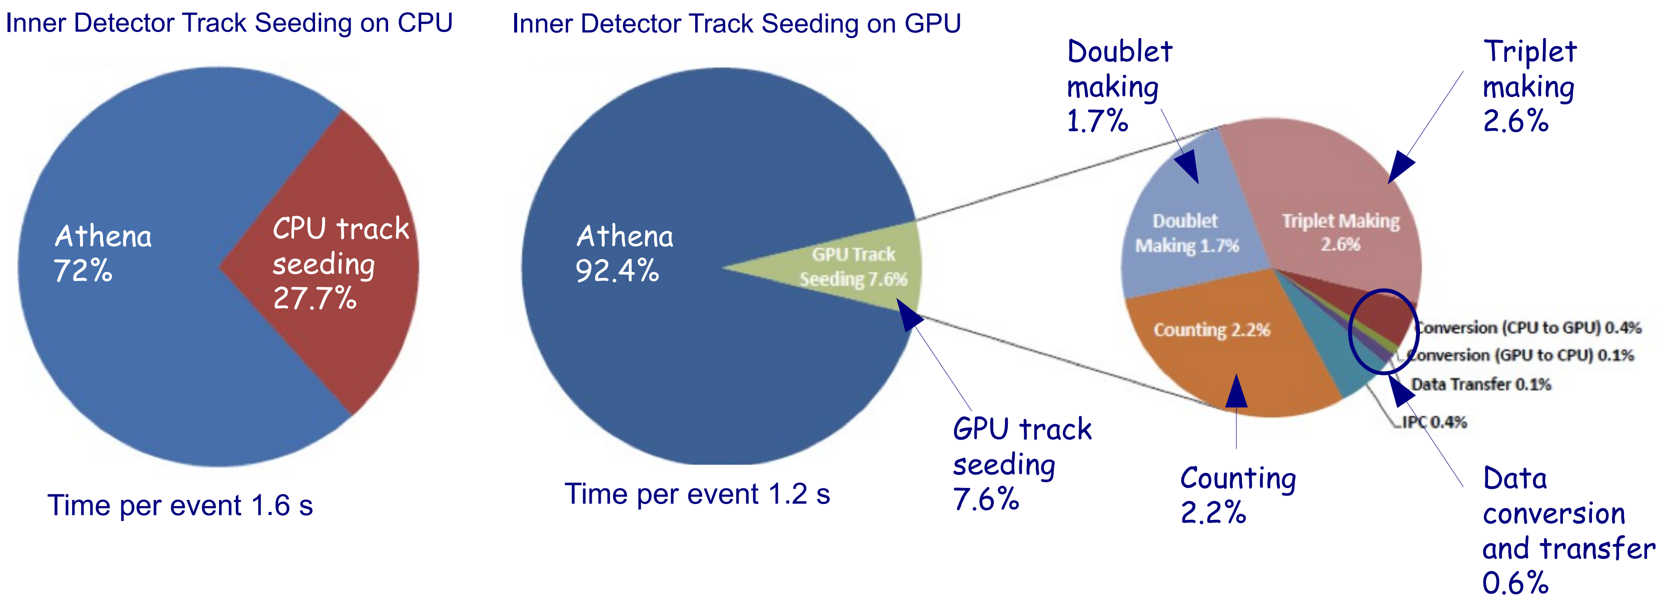

GPGPU Calorimeter Clustering Algorithm

- TopoCluster reconstruction on CPU (~8% of total time)
  - Group cells in 3-dimensions according to their signal/noise ratio

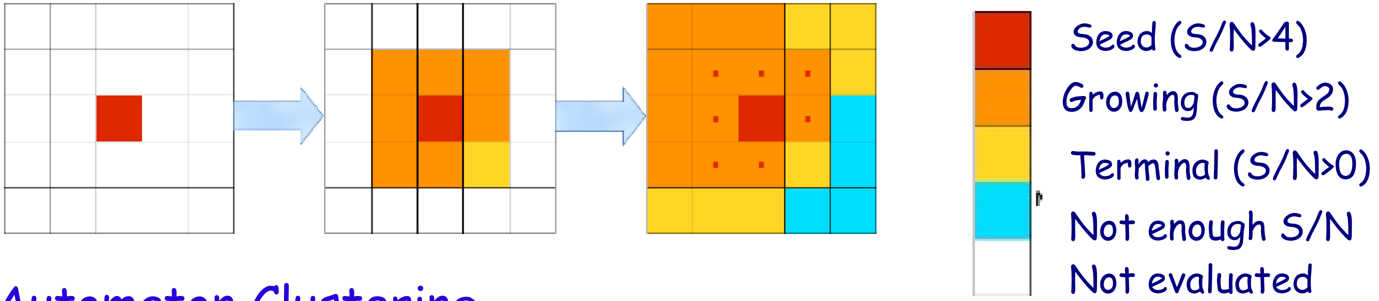

- TAC: Topo-Automaton Clustering
  - Use a cellular automaton for the GPU (maximize parallelism)
  - Propagate tag on a grid of elements (cell pair)
    - Cells get the largest tag on each iteration
  - Process all cells pairs until no tag changes

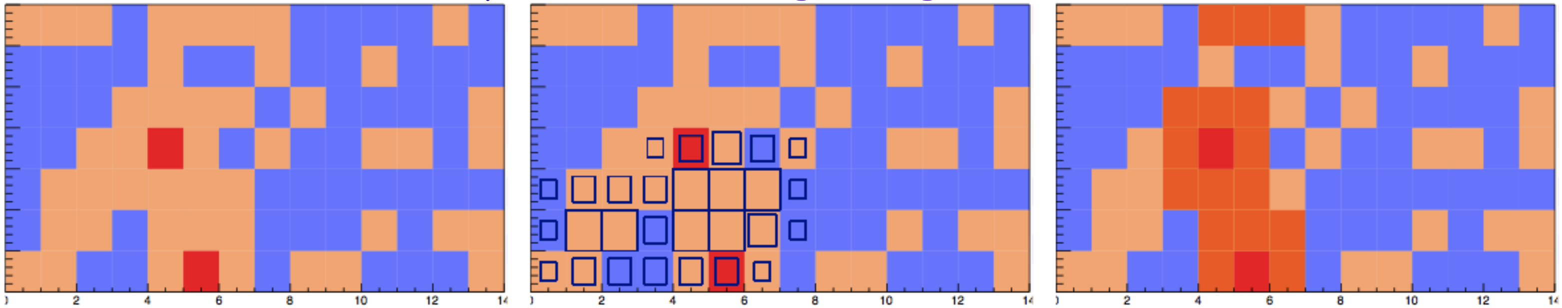

GPGPU Calorimeter Clustering Performance

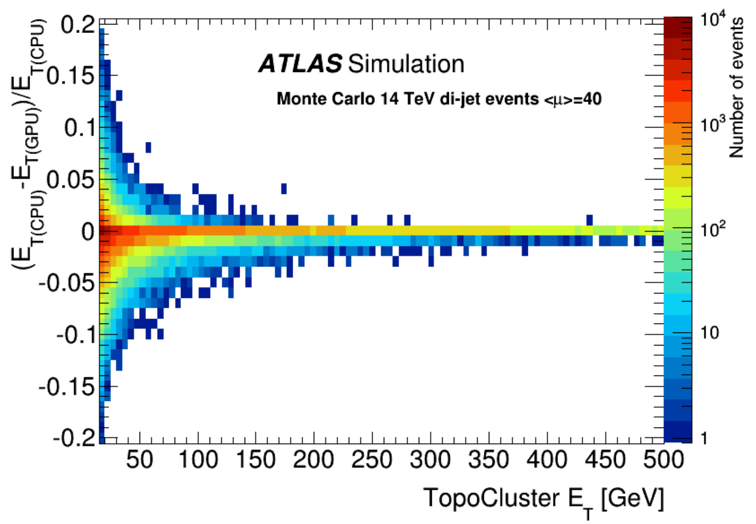

- Energy difference <5% for most clusters
- Cluster growing time reduction factor:

Sample	Pile-up	Reduction factor
tt-bar	138	2
tt-bar	46	2
di-jets	40	1.3

- Potential larger gain with parallelization of next clustering steps (splitting)

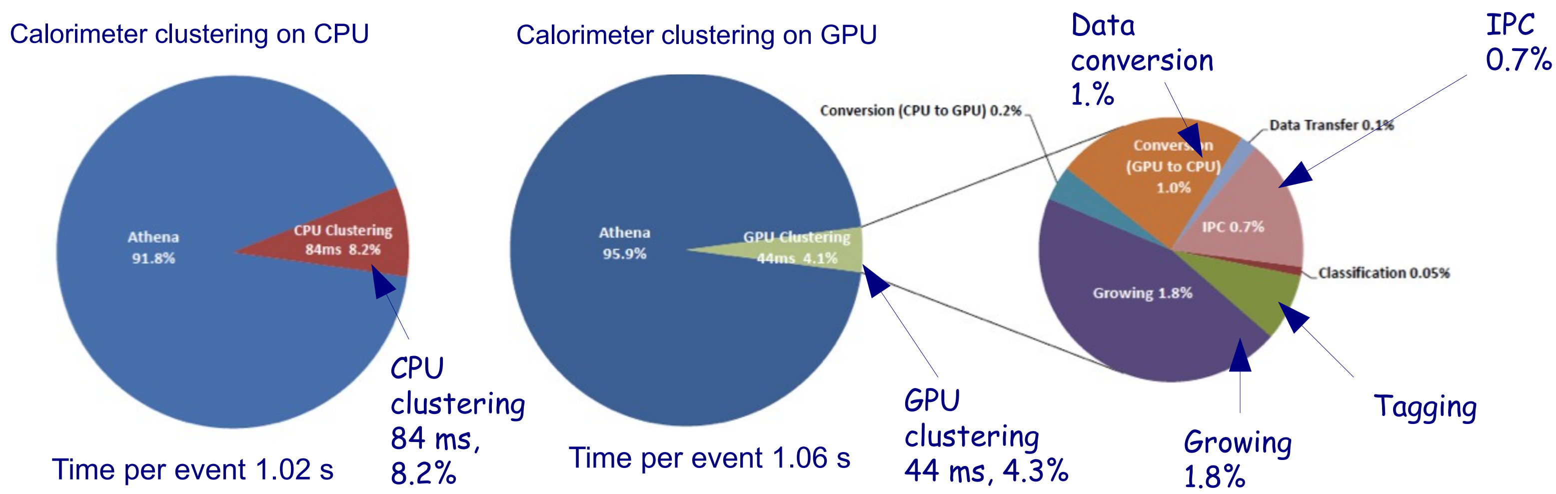

Muon reconstruction algorithms

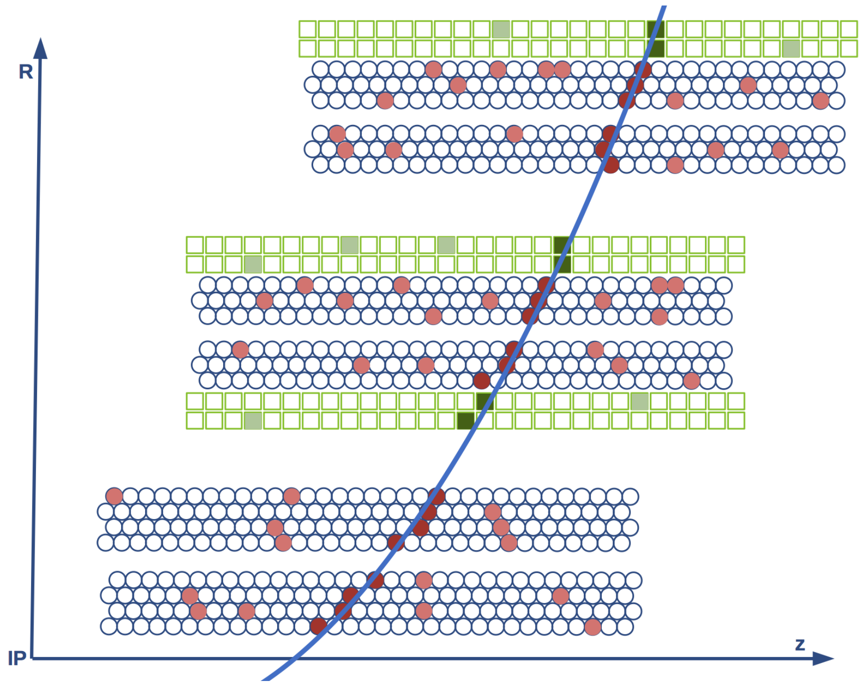

- Uses Hough Transform (HT) to convert track finding to maxima finding
  - Straight HT is used for xy-plane, curved HT in rz-plane
- GPU algorithm
  - filter hits and fill Hough-space matrices
  - Sort maxima above certain threshold
  - Associate hits in spectrometer
- 3D segments constructed by combining 2D segments in CPU

Performance of the GPU Prototype

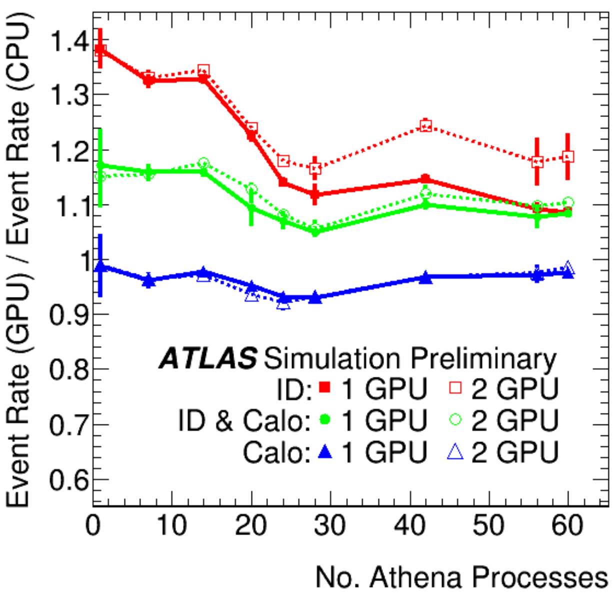

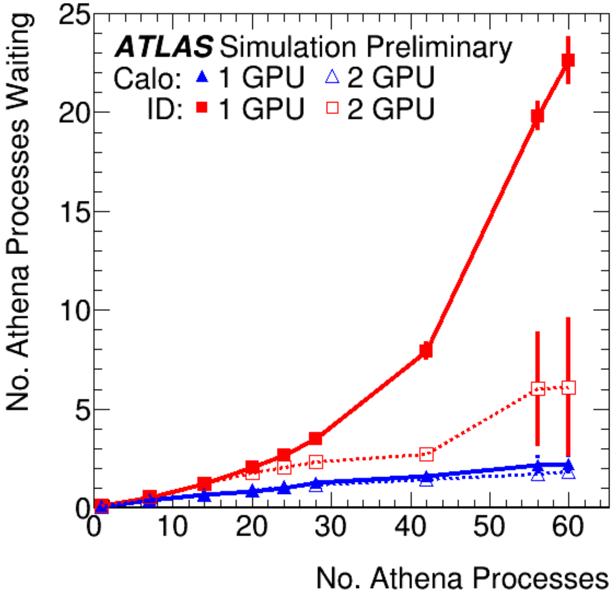

CPU: Intel(R) Xeon(R) E5-2695 v3 14-core

GPU: NVidia K80

Data sample: simulated tt events (46 collisions/bunch crossing)

- Gain in throughput: 20-40% depending on number of processes running
- 1 GPU can serve efficiently up to 14 processes
- Gains will increase when more code is offloaded (ex. Bytestream conversion, track following, cluster splitting, …)

Jet reconstruction algorithms already implemented on GPU – performance measurements underway

Summary and conclusions

- The LHC Upgrade will impose stringent requirements on the ATLAS trigger system
  - Need advanced algorithms, capable of higher rejection with same efficiency
- ATLAS is studying the use of GPGPUs for triggering
  - Require re-implementation of the algorithms to maximize parallelization
- First evaluation of calorimeter and tracking reconstruction
  - Achieved the same physics performance in tracking & cluster reconstruction
  - Total execution time reduced by a maximum of
    - A factor of 5 for tracking
    - A factor of 2 for cluster formation
  - Lesson: data structures suitable for CPU & GPU would reduce overheads
- Gain in number of processed events/s:
  - Between 20-40%, depending on number of processes accessing the GPU
  - Larger gain expected when more code is offloaded to the GPU

Acknowledgements

➢ OE, FCT-Portugal, CERN/FIS-NUC/0005/2015

Backup

GPGPU Calorimeter Clustering Performance

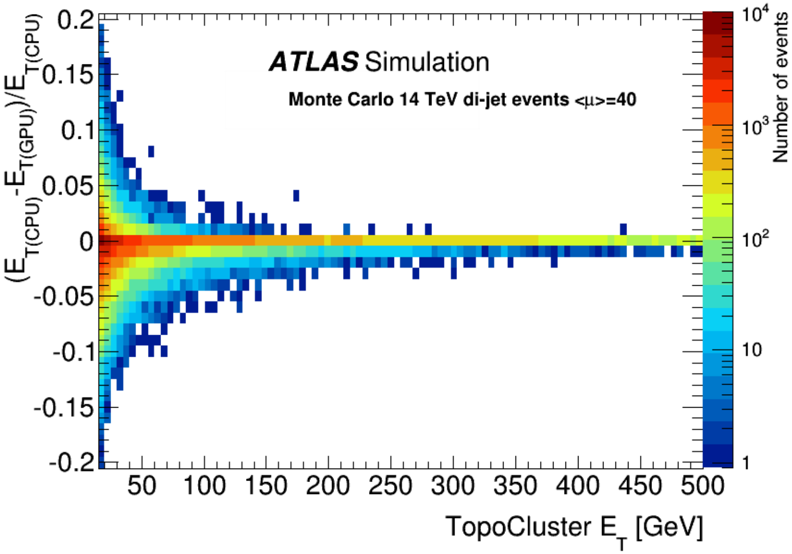

Clusters E_T spectrum:

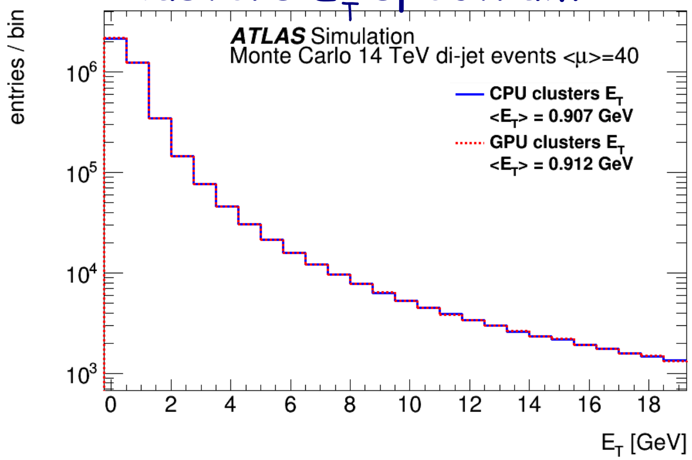

Number of jets/event

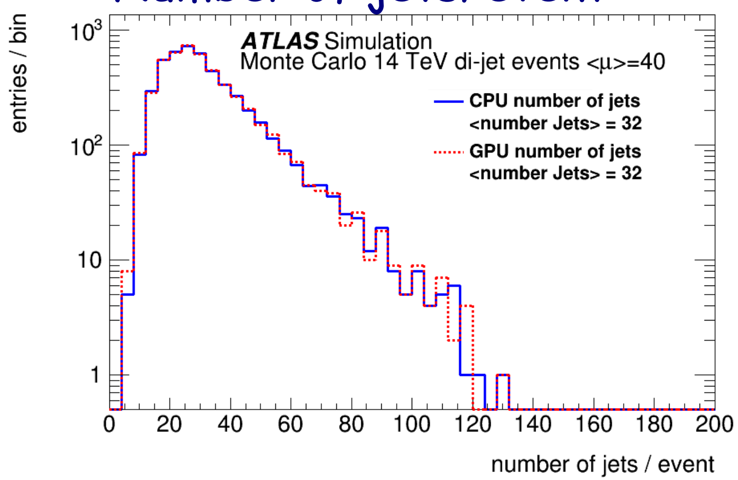

- Energy difference <5% for most clusters
- No significant effects on clusters E_T or jet reconstruction

GPGPU Calorimeter Clustering Timing

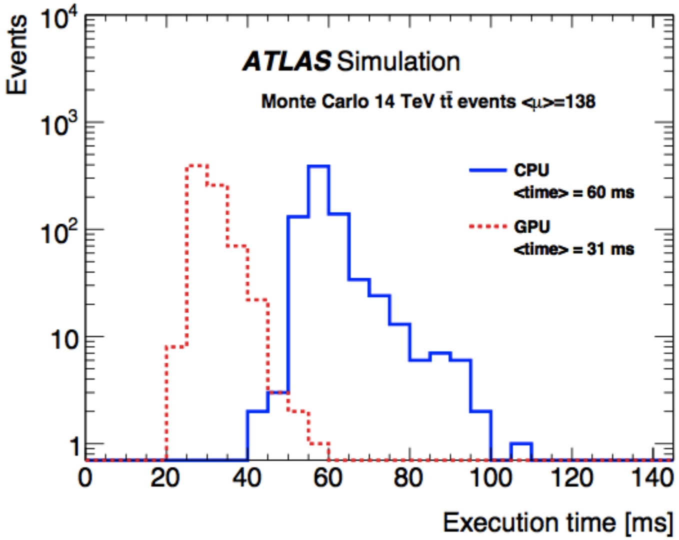

- Cluster growing time reduction factor:
  - Factor ~2 for tt events with 138 interactions/bunch crossing
  - Reduction of 30% for di-jet events with 40 interactions/bunch crossing

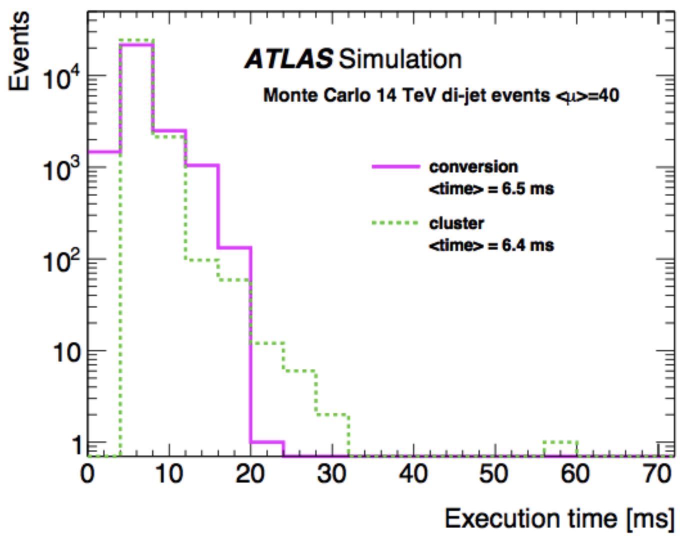

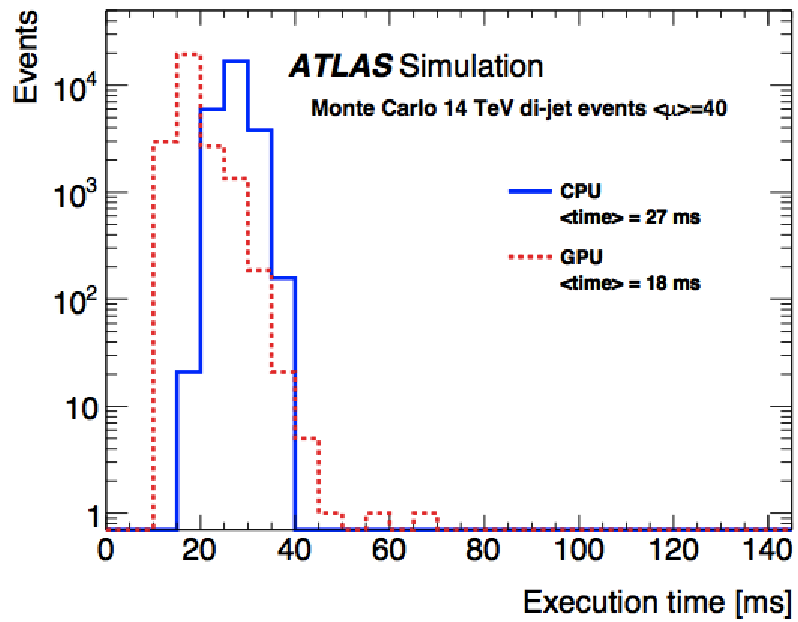

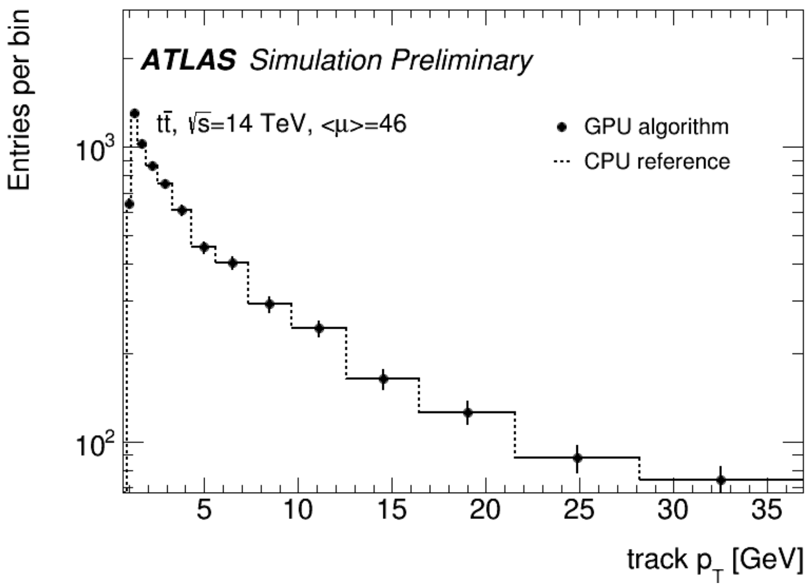

Entries per bin
ATLAS Simulation Preliminary
$t\bar{t}$, $\sqrt{s}$=14 TeV, <μ>=46
GPU algorithm
CPU reference
10^3
10^2
5
10
15
20
25
30
35
track p_T [GeV]

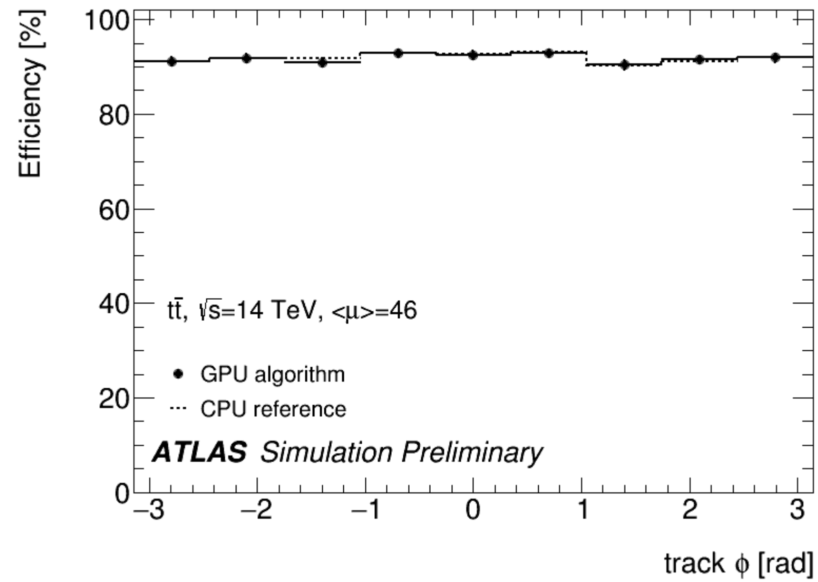

Efficiency [%]
100
80
60
40
20
0
$t\bar{t}$, $\sqrt{s}$=14 TeV, <μ>=46
GPU algorithm
CPU reference
ATLAS Simulation Preliminary
−3
−2
−1
0
1
2
3
track φ [rad]

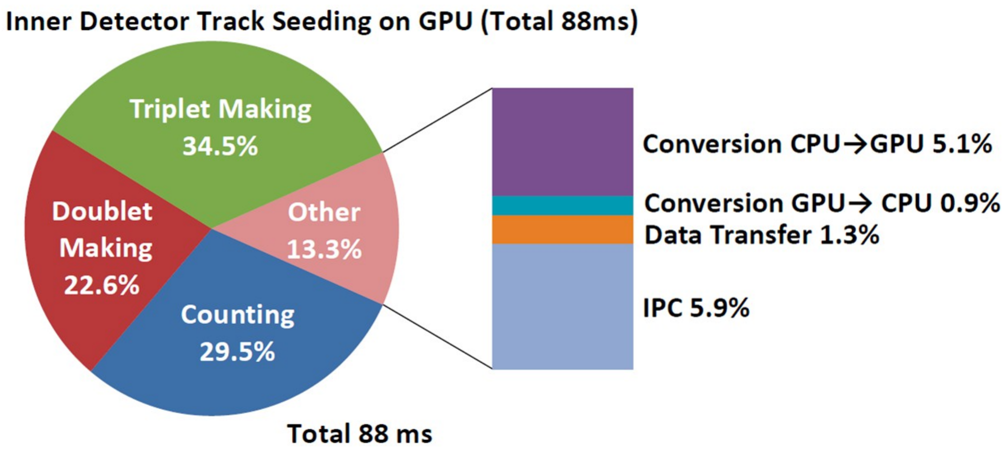
Inner Detector Track Seeding on GPU (Total 88ms)
Triplet Making 34.5%
Doublet Making 22.6%
Other 13.3%
Counting 29.5%
Conversion CPU→GPU 5.1%
Conversion GPU→ CPU 0.9%
Data Transfer 1.3%
IPC 5.9%
Total 88 ms

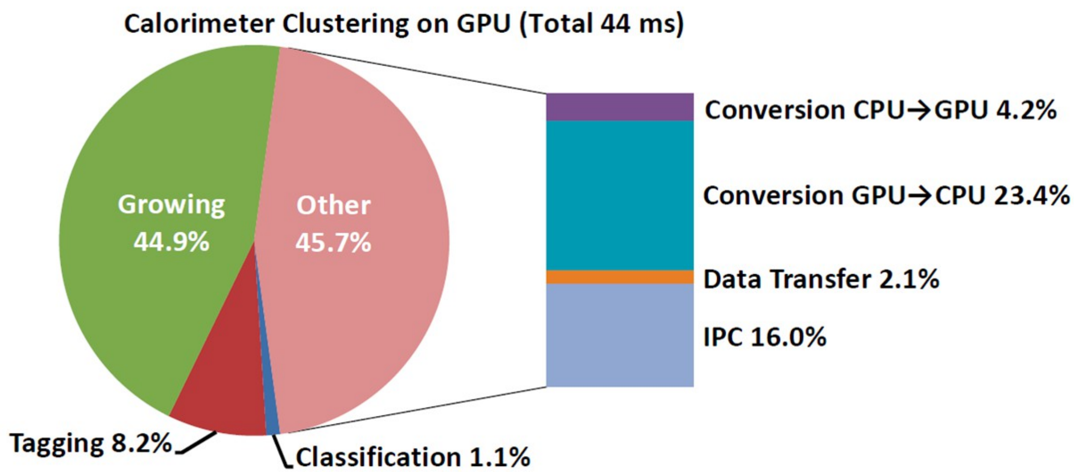
Calorimeter Clustering on GPU (Total 44 ms)
Growing 44.9%
Other 45.7%
Tagging 8.2%
Classification 1.1%
Conversion CPU→GPU 4.2%
Conversion GPU→CPU 23.4%
Data Transfer 2.1%
IPC 16.0%