



Parallelizing Unpacking and Clustering



Motivation

- CMS is evaluating the use of GPUs in the HLT for Run3
- Including *mkFit* in HLT workflow for Run3

Parallel Unpacker Implementation

- Unpacking of raw data from CMS Silicon Strip with SoA format
- Parallelizing “three threshold” clustering algorithm
- Multi-event processing
 - OpenMP with nested parallelism in CPU
 - CUDA streams in GPU

Performance Results

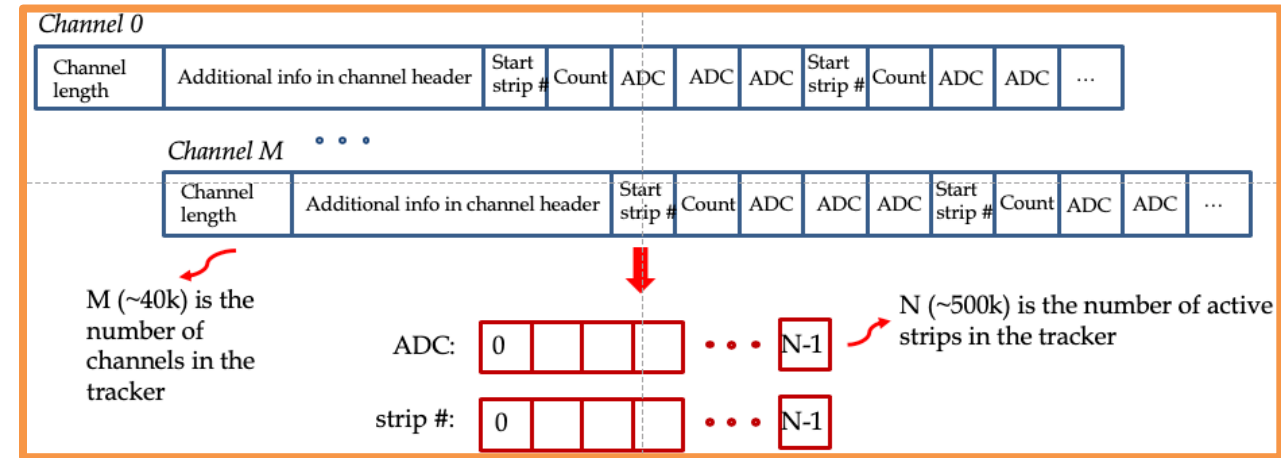
- CPU (2x14 cores Broadwell) throughput: 400-560 (events/s)
- GPU (P100) throughput: 546-649 (events/s)

Next Steps

- Integrate unpacker into CMSSW
- *mkFit* hit producer: convert clusters to global coordinates
- Investigate OpenCL implementation for FPGA

Code available at: https://github.com/beiwang2003/strip_clustering_gpu

Unpacking



“Three threshold” clustering

