

IRIS-HEP RETREAT

LIGHTNING TALK

Bei Wang
Software Engineer
Princeton University

Sep 12, 2019

About Me

- **Bei Wang**, Ph.D. in computational science and engineering, **software engineer** at Princeton university
- Has been working on R&D in parallel computing applications for over 8 years in plasma physics, fluid dynamics and high energy physics
 - Co-PI of **Intel Parallel Computing Center (IPCC)** at Princeton
 - Code development and optimization on modern computing architectures

Project: Particle Tracking

- **Motivation:**
 - MkFit is a parallelized KF-based code that does particle tracking and reconstruction on multicore CPUs and manycore GPUs
 - MkFit needs all hits as an input
 - Currently, all hits needed by MkFit are obtained from a CPU code that performs strip unpacking and clustering **serially**
- **Goal:**
 - Investigating strip unpacking and clustering on GPU
 - Could also use the GPU to produce hits in the MkFit data format

Current Status

- Developed a parallel version of the clustering algorithm
 - Git repo: https://github.com/beiwang2003/strip_clustering_gpu
 - Data are stored as Structure-of-Array (SoA) format
 - SIMD friendly implementation
 - Scale on multicore CPU using OpenMP

Cluster Criteria

- The ADC value of each strip must be $> 2 \cdot \text{noise}$
- **At least one strip has ADC value $> 3 \cdot \text{noise}$**
- The quadrature sum of ADC values of the cluster must be $\geq 5 \cdot$ the quadrature sum of the noise
- The strips are consecutive or have gap of at most N strips (N is a input parameter)

New Parallel Algorithm

1. Find out all seed strips
 - we call strip with $ADC > 3 * noise$ as seed strip
2. Keep only one seed strip (the first one now) if there are multiple seed strips with consecutive strip id
 - **issue**: we might still have duplicated clusters generated from non-consecutive strips
3. For each non-consecutive seed strip, find the left and right boundary of the candidate cluster
4. For each candidate cluster, check if it is a true cluster. If so, do some adjustment for its adc values. A true cluster satisfies:
 - Squared noise is large enough
 - chargePerCM is large enough

Future Work

- Developing the first GPU version of the code (in CUDA)
- Current clustering algorithm assumes that the input data for a single event are stored in SoA format. More work is needed if the output from unpacking is in a different format
- Benchmark with the CPU version