Accelerating track reconstruction algorithms using GPUs at ATLAS

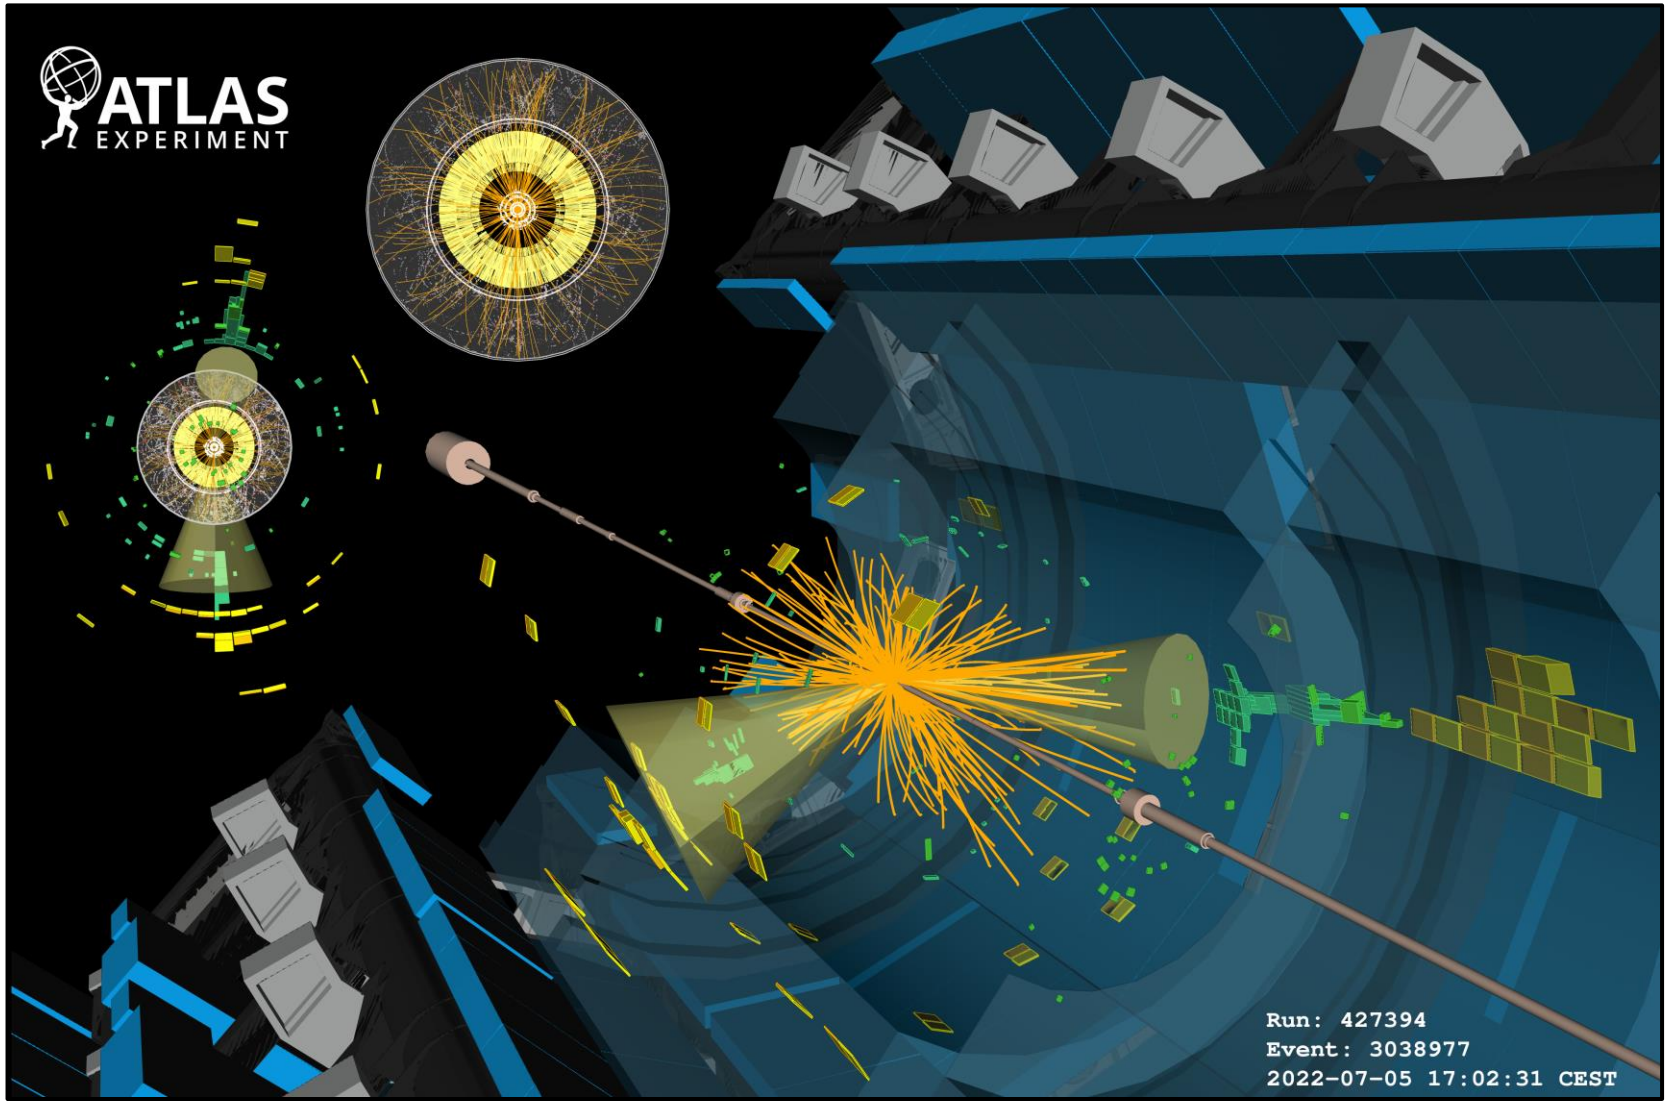

ACTS track Reconstruction

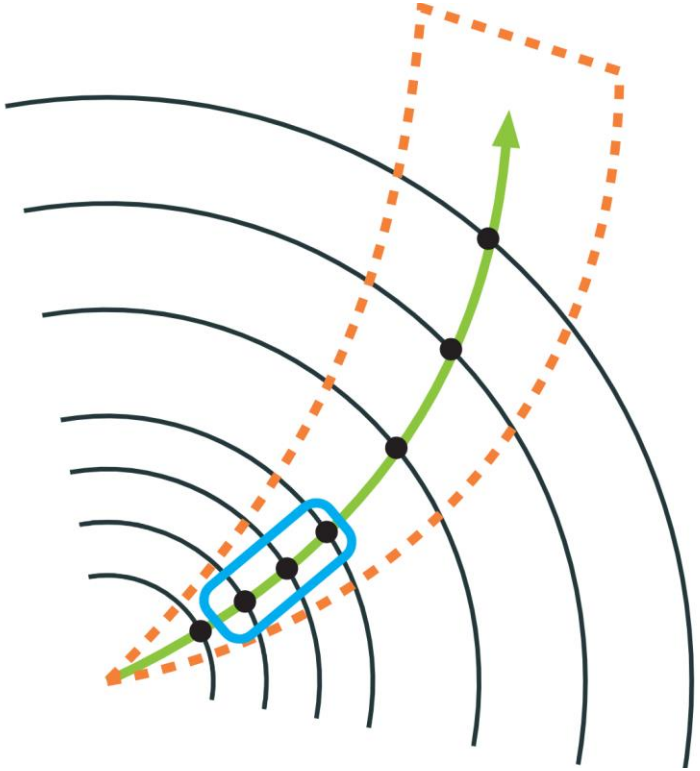

Seeding

Track finding

Track fitting

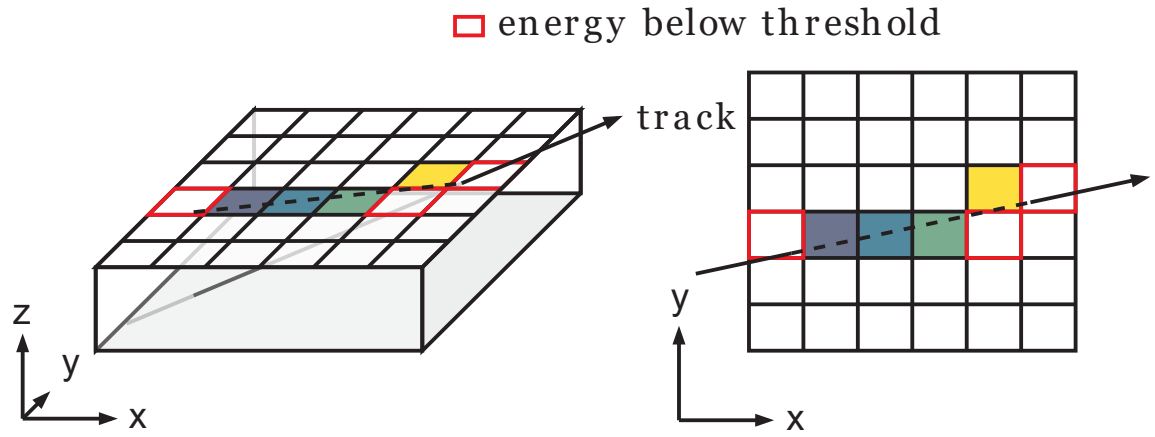

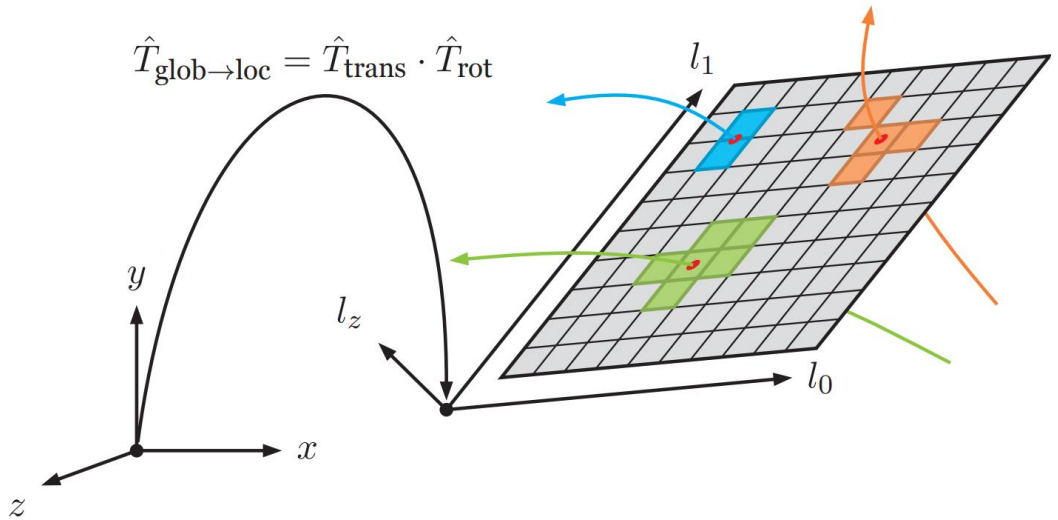

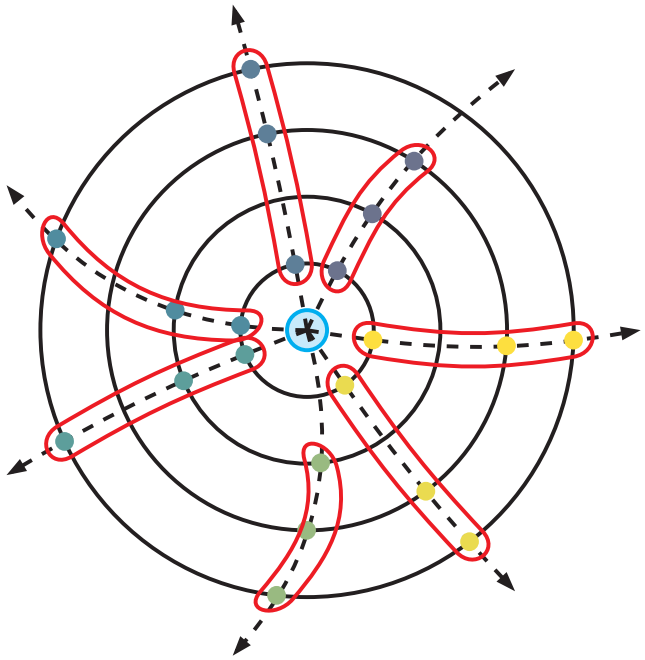

traccc project

- ACTS track reconstruction on GPU
- No physics performance loss
- Massively parallel programming
- No dynamic memory allocation in GPU threads
- Count/Find kernels rather than overallocation
- Prefix sum vector for linking a thread number to a jagged position

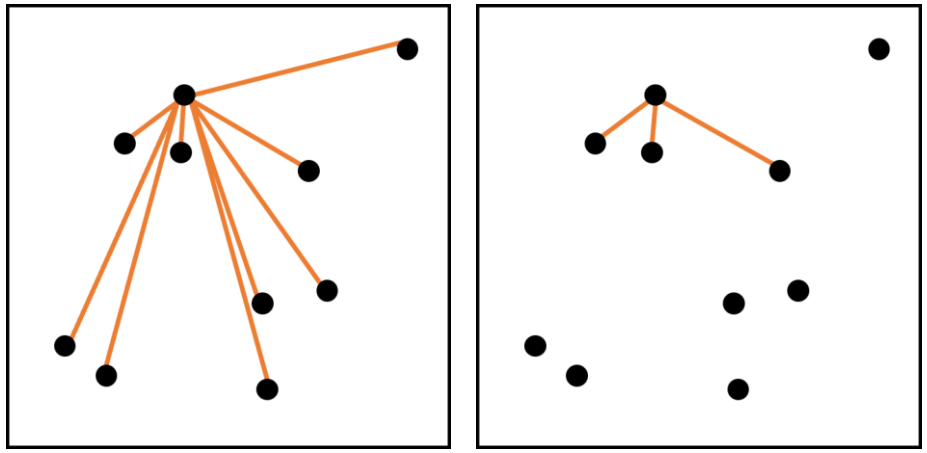

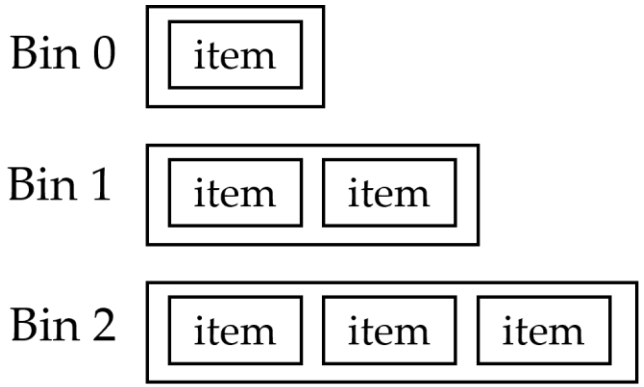

Input jagged vector

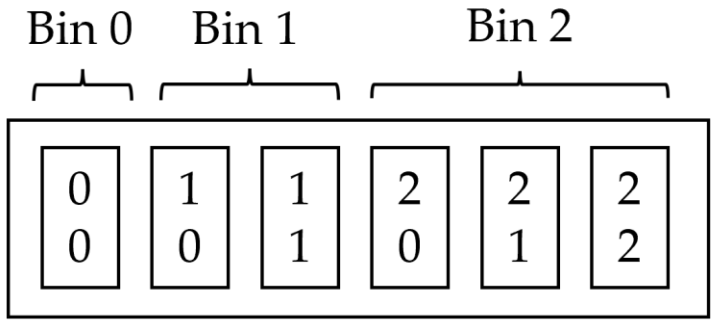

Auxiliary prefix sum vector

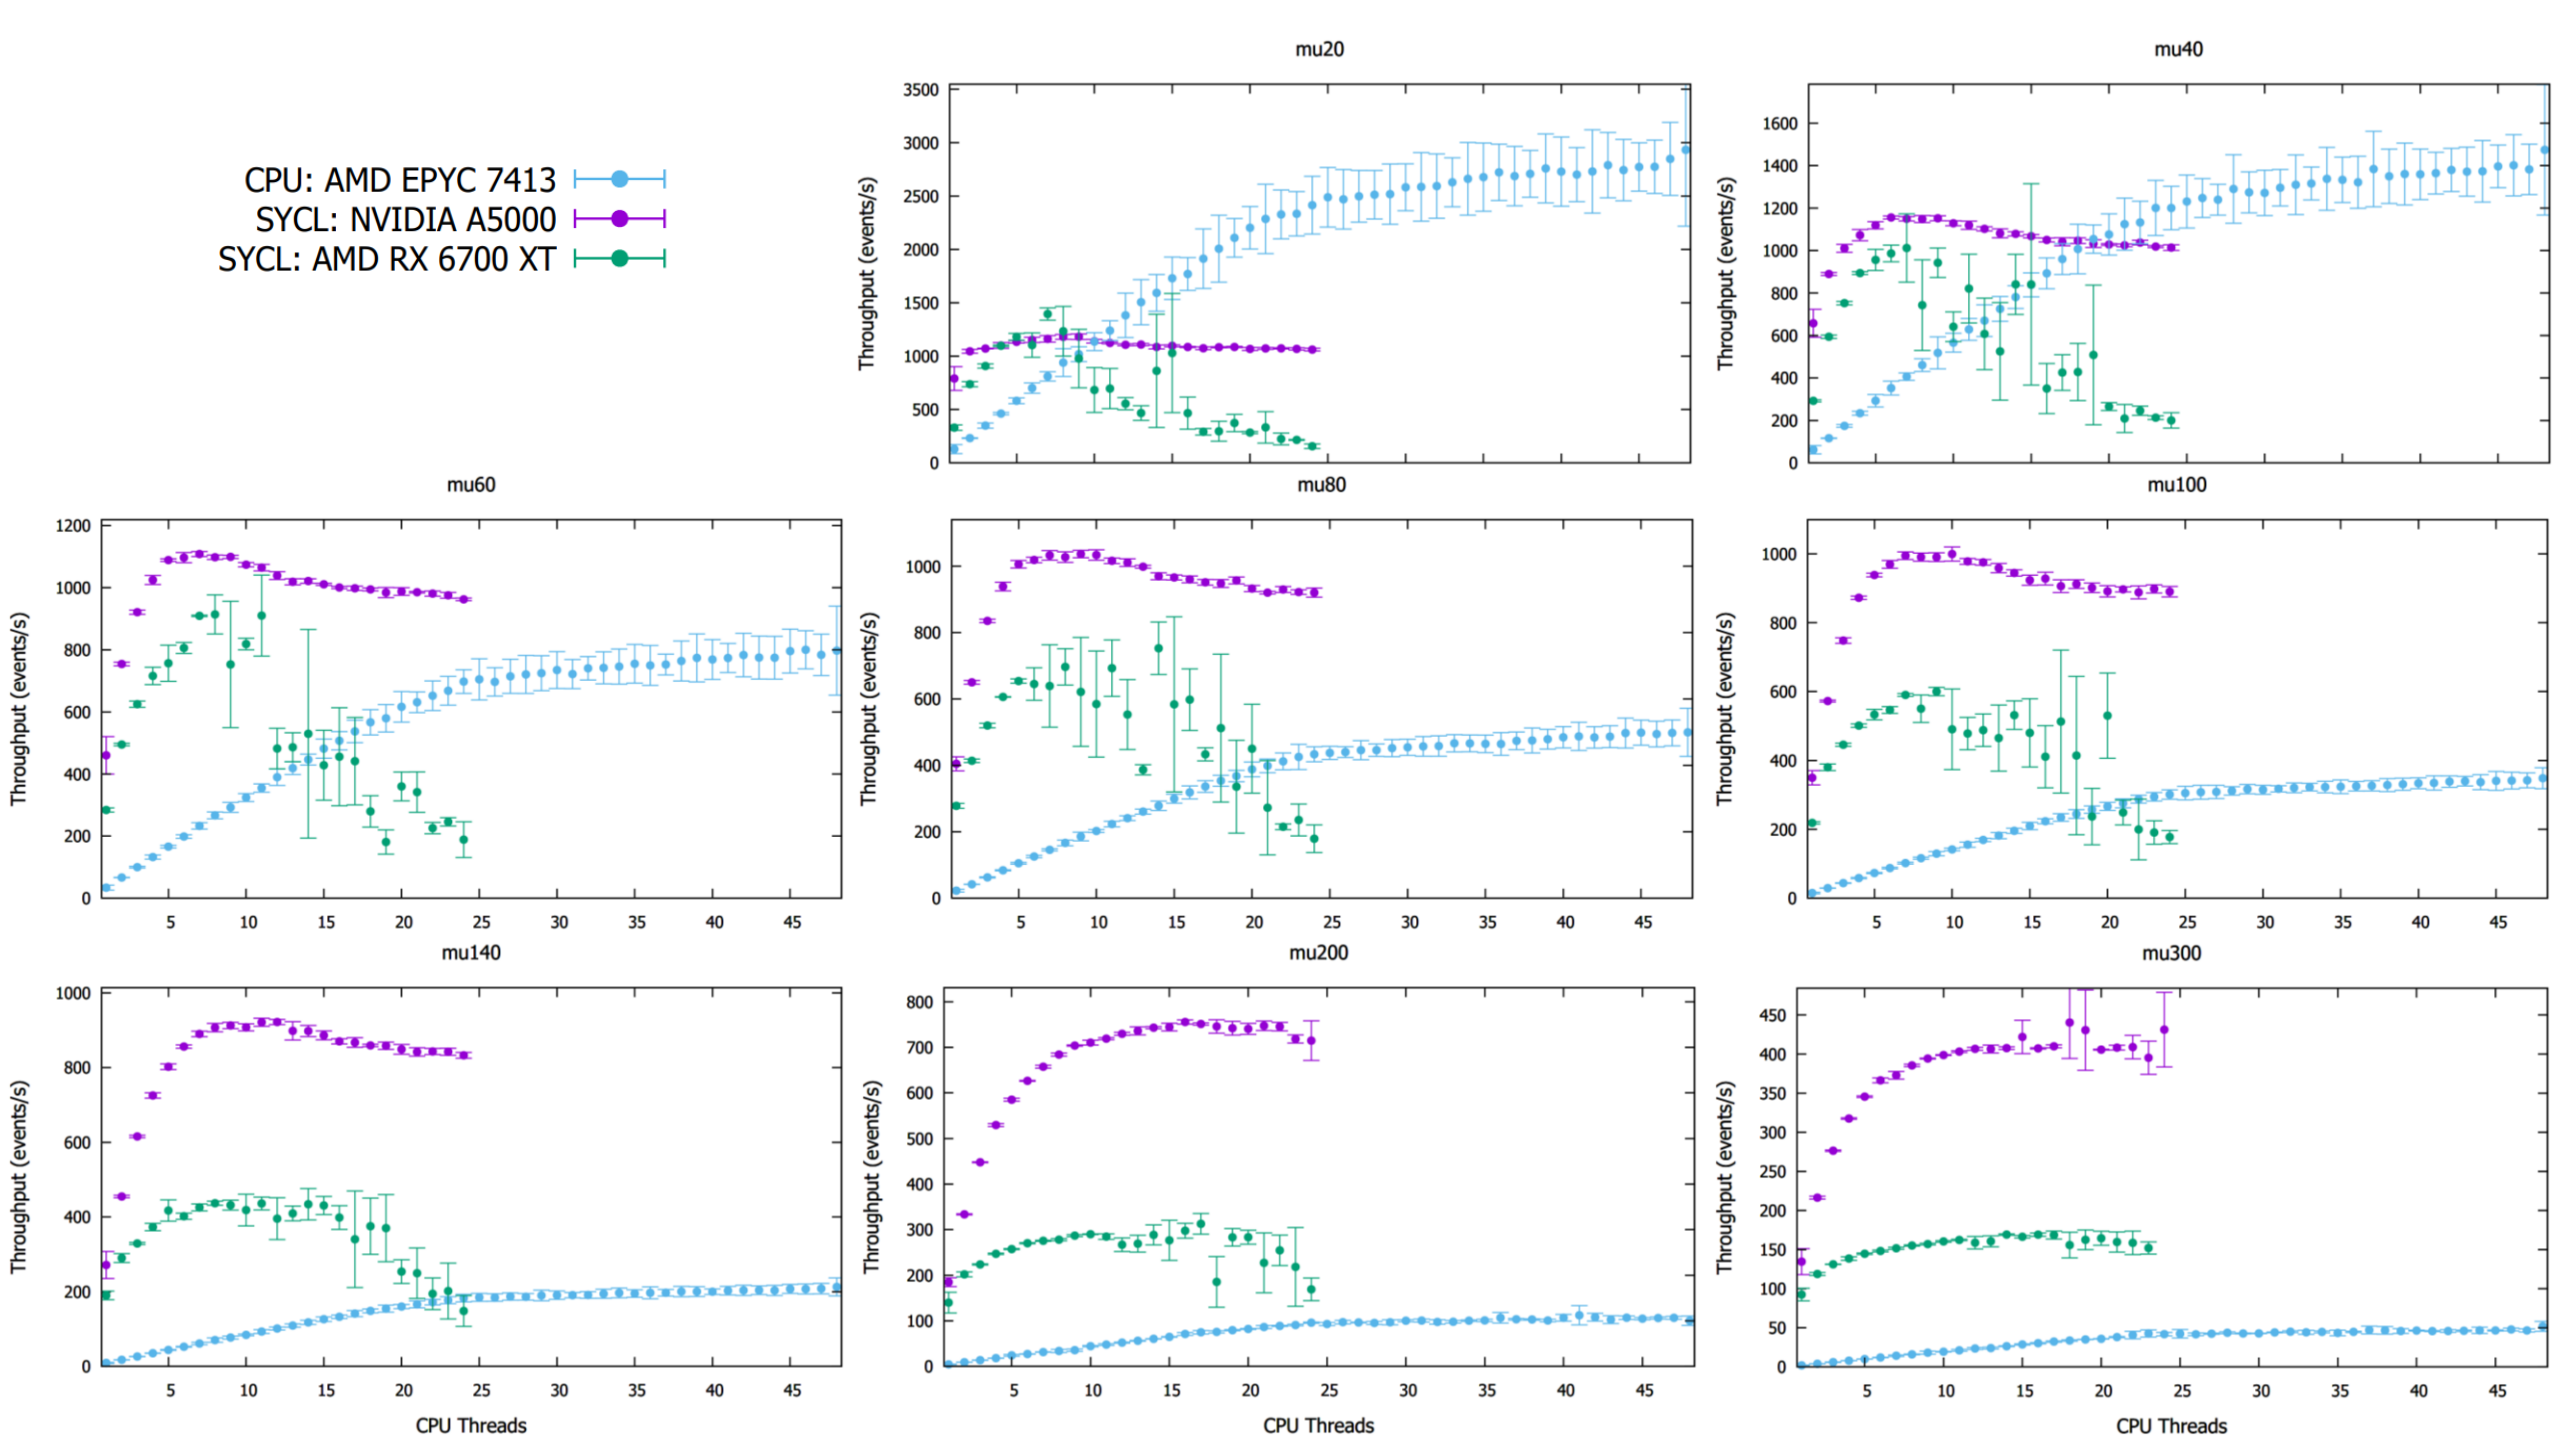
CPU: AMD EPYC 7413
SYCL: NVIDIA A5000
SYCL: AMD RX 6700 XT
mu20
mu40
mu60
mu80
mu100
mu140
mu200
mu300
Throughput (events/s)
CPU Threads

Measurements at peak performance

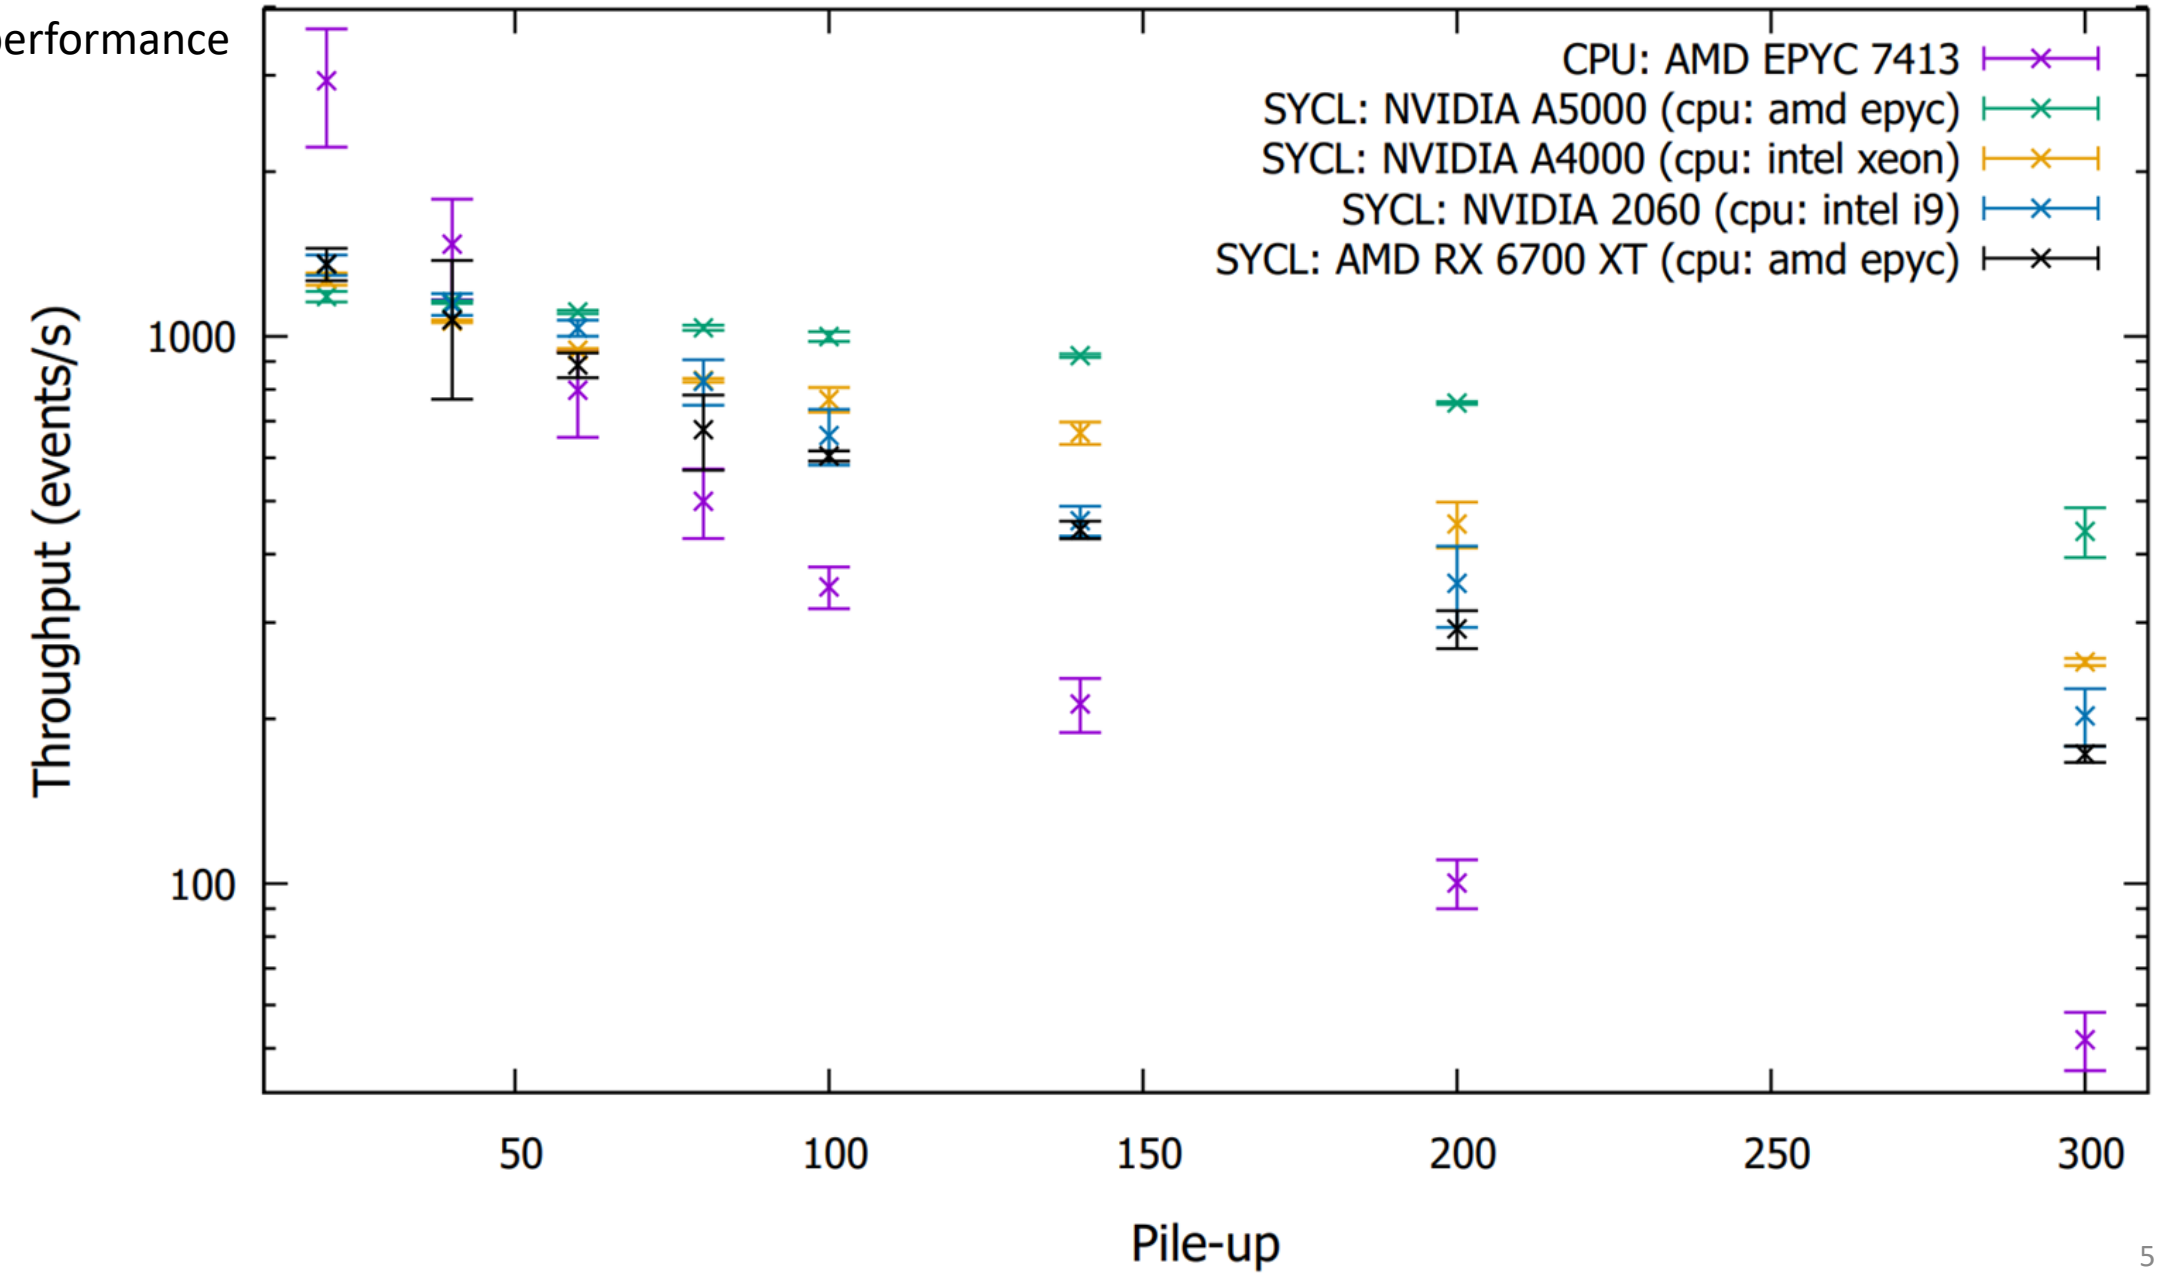

Backup

Current status - functionality

- Different backends
  - ➢focus on SYCL/CUDA
- Large part of track reconstruction running entirely on the GPU
- Still being extended

Category	Algorithms	CPU	CUDA	SYCL	Futhark
Clusterization	CCL	✅	✅	✅	✅
	Measurement creation	✅	✅	✅	✅
	Spacepoint formation	✅	✅	✅	○
Seeding	Spacepoint binning	✅	✅	✅	○
	Seed finding	✅	✅	✅	○
	Track param estimation	✅	✅	✅	○
	Combinatorial KF	○	○	○	○
Track fitting	KF	🟡	🟡	○	○

✅: exists, 🟡: work started, ○: work not started yet

CPU vs GPU algorithm

- Running generic algorithm with input jagged vector and output jagged vector

on a CPU:

- Do calculations on input
- Append new members to result

on a GPU:

- 1st kernel – Create prefix sum vector for input jagged vector
- 2nd kernel – Count number of members needed for result
- 3rd kernel – Create prefix sum vector for counted items
- 4th kernel – Fill result jagged vector

+ CPU memory allocation between kernels

- Code itself can be very different to boost GPU resource usage
- GPU introduces a lot of complexity
- Throughput gain from massive parallelism

New clustering

- From module to cell parallelism
- No jagged data in clustering
- Organise data in equally sized partitions (on CPU)

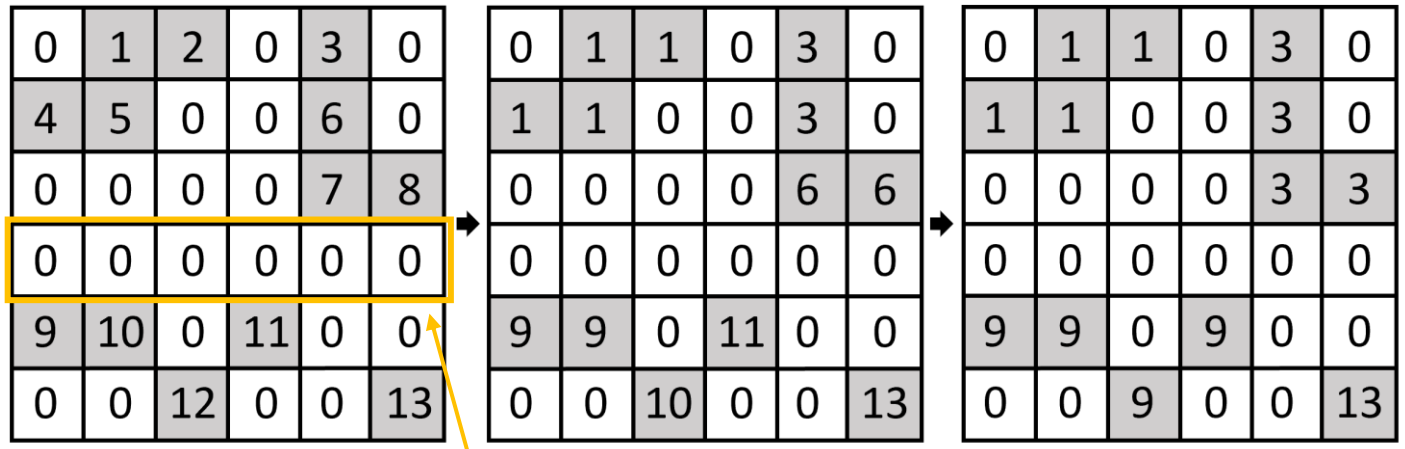

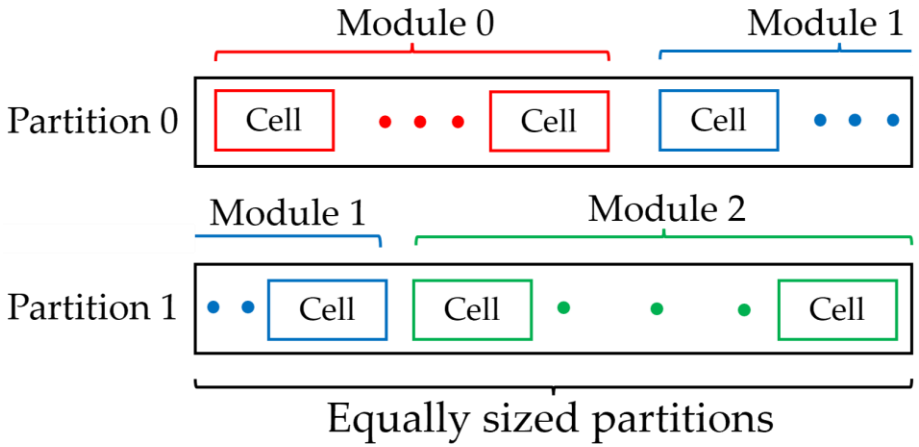

For the coding details, see [the open PR in GitHub]

New clustering - performance

CPU: AMD EPYC 7413
GPU: NVIDIA RTX A5000

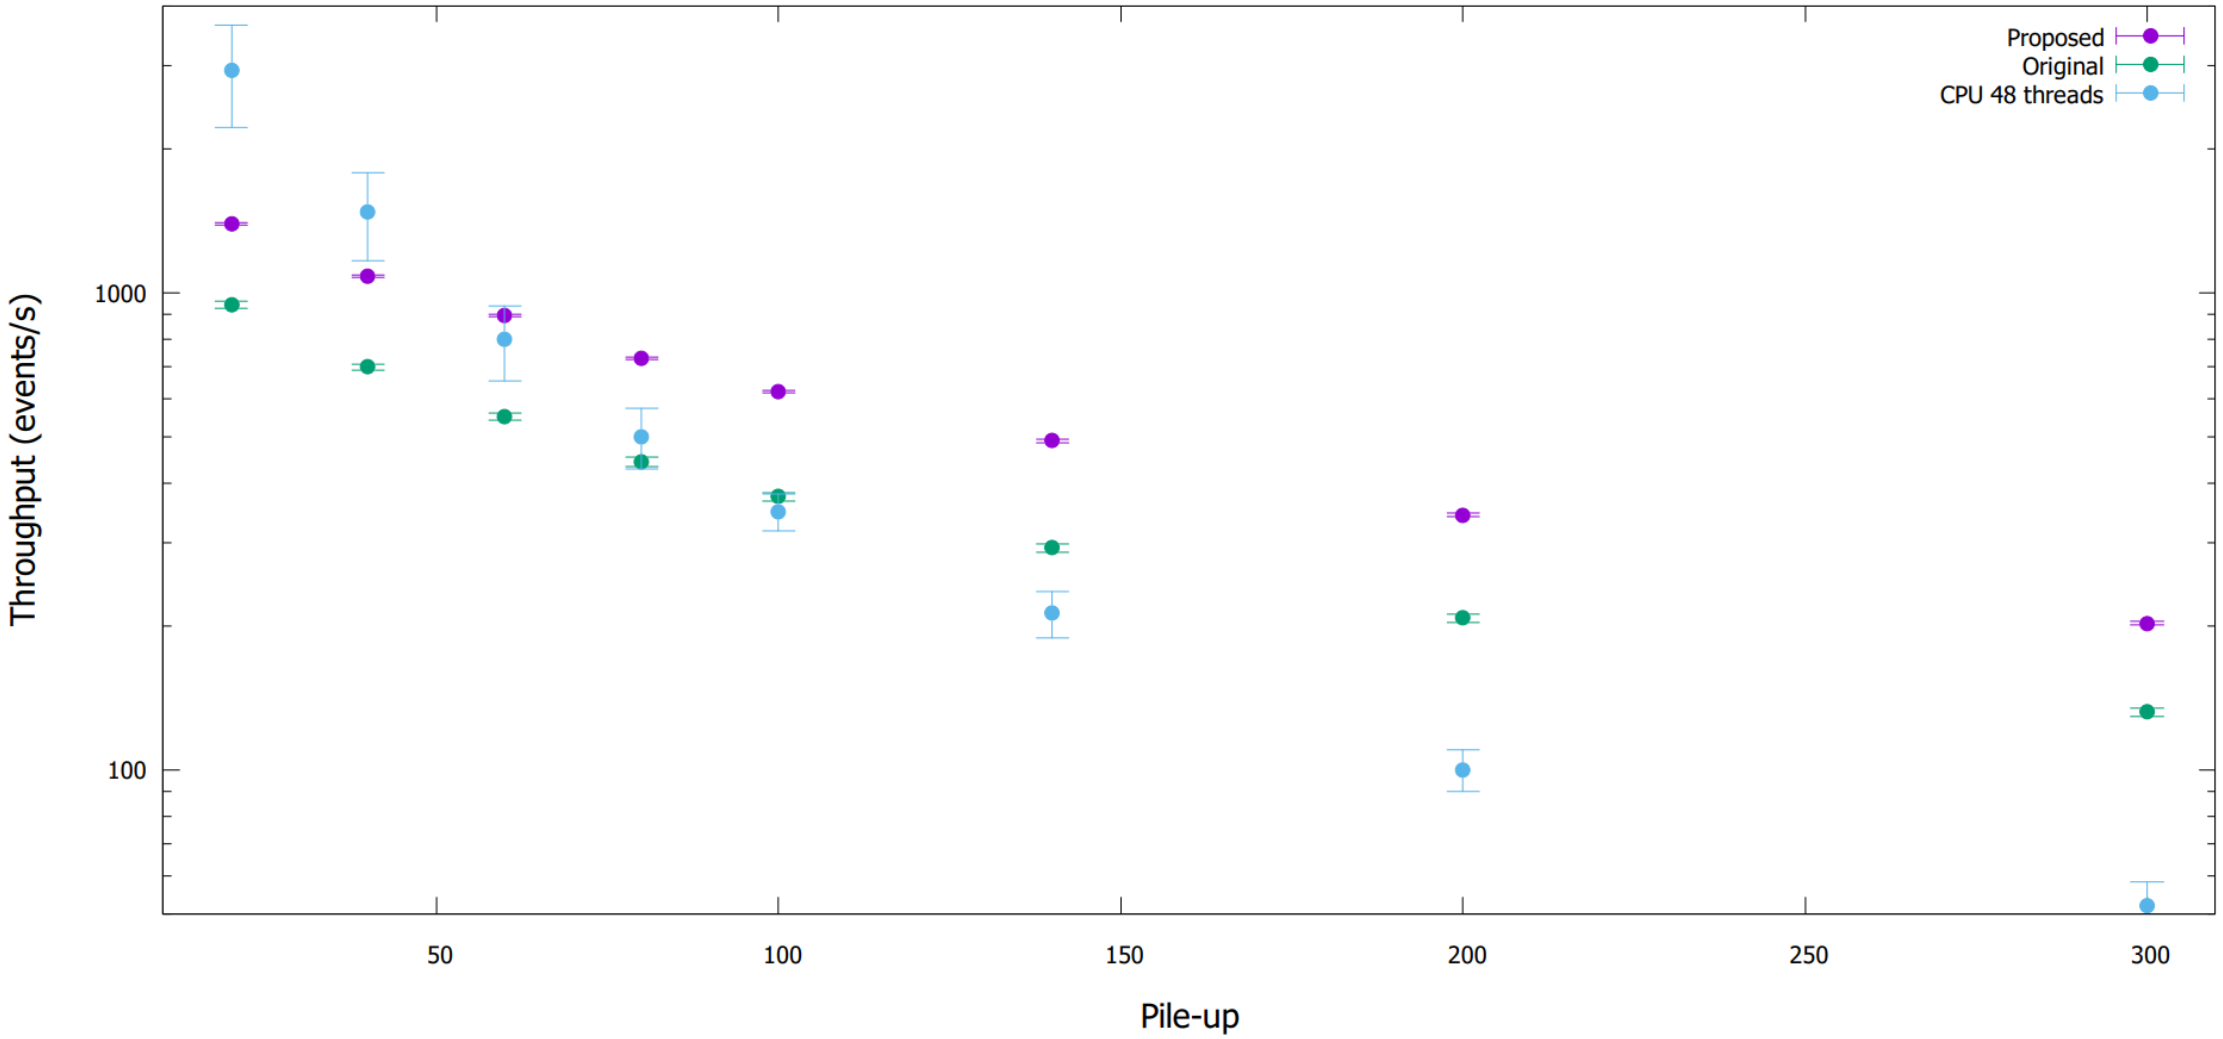

- WIP: 50% – 70% throughput improvement for full application

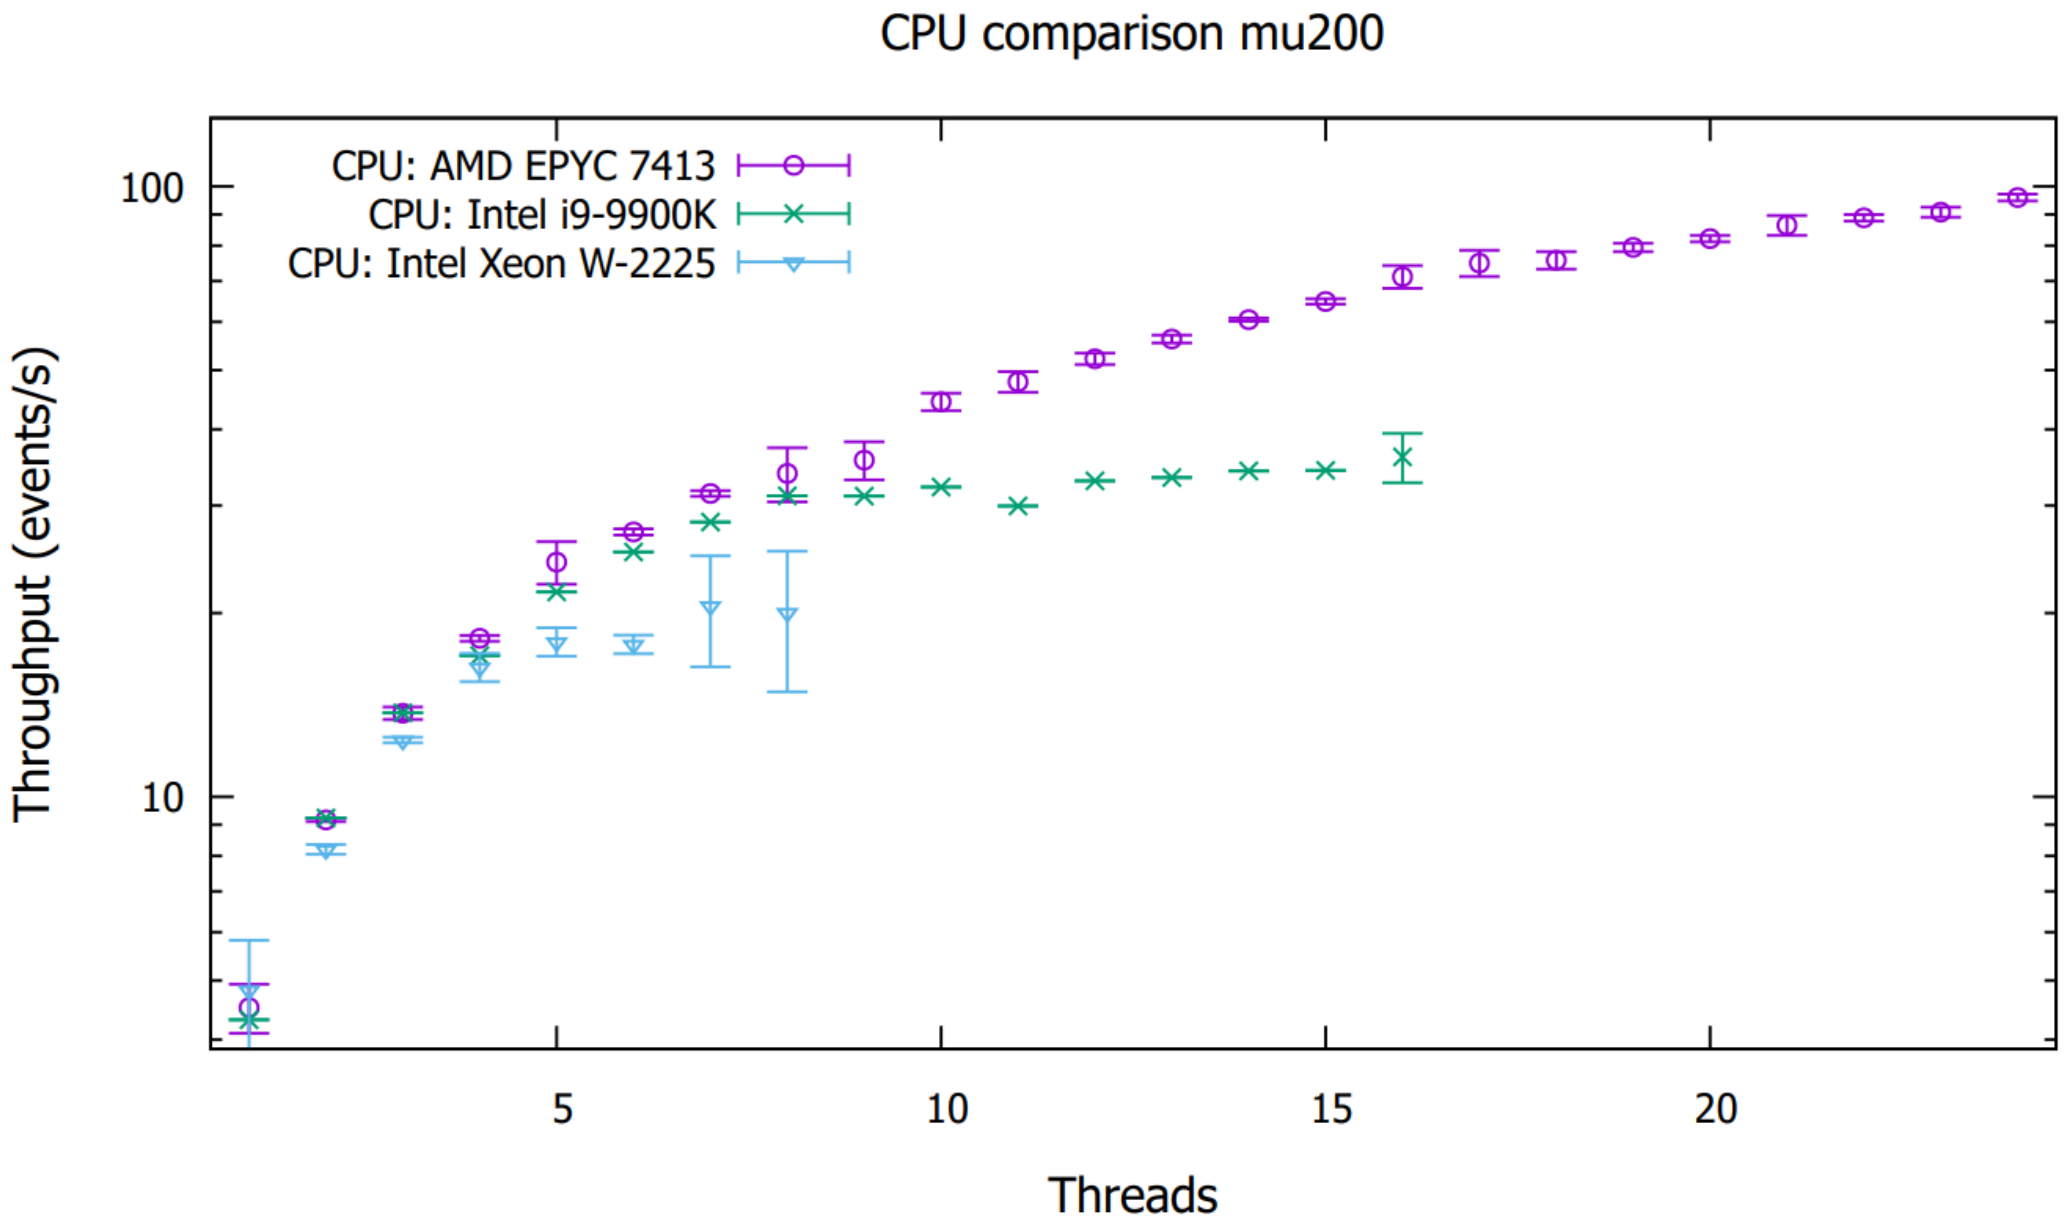
CPU comparison mu200
CPU: AMD EPYC 7413
CPU: Intel i9-9900K
CPU: Intel Xeon W-2225
Throughput (events/s)
100
10
5
10
15
20
Threads

SYCL: NVIDIA A5000 (cpu: amd epyc)
at different pile-up events