

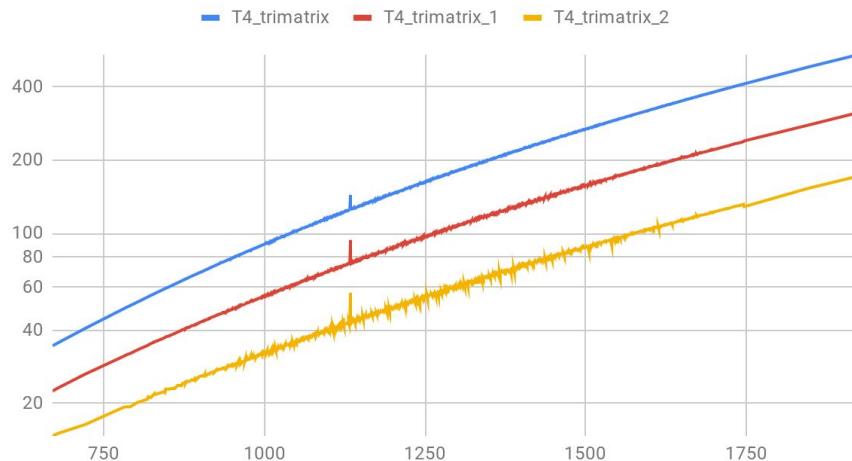
FastJet in CUDA

Andrea Bocci, Abdulla Mohamed

Day 2

- All kernels produce the correct results.
- Optimized tri matrix kernel by storing the results of double precision computations instead of recomputing them every time (**improved**).
- Did the same optimization with the grid kernel (**worse**).
 - Reason: recomputation happens in less than 2% of grid usage.

of particles, T4_trimatrix and T4_trimatrix_1



Next day

- Exploit the fact that the grid kernel does not need recomputations in each loop.

Or

- Process events in parallel if it is worth it.
 - Currently the GPU is 18.5 slower than CPU.

https://github.com/asubah/fastjet_gpu