CERN openlab III

Report from International Tracking Workshop

GSI/Darmstadt, 9 – 11 June 2010

Sverre Jarp, CERN openlab

Background

- ~40 participants
- Two days of training, three-day workshop
- Main interest from heavy-ion experiments
- Fixed target:
  - CBM, Panda, Hades (GSI)
- Collider:
  - ALICE (CERN), STAR (Brookhaven)
- All depend on an excellent High-Level Trigger (HLT) capability

Goals

- Develop common methods for tracking and vertexing
  - Could there be a common (and complete) toolkit ?
    - As we have it for simulation and data analysis today
- Optimize performance
  - Algorithmic work
  - Exploit modern hardware fully
    - CPUs (Vectors, Cores)
    - Accelerators:
      - GPUs (NVIDIA, AMD)
      - Intel MICA

Ivan Kisel's original work

http://www-linux.gsi.de/~ikisel/17_CPC_178_2008.pdf

- Main performance improvements came from the algorithmic work (not the hardware optimization):

Stage	Description	Speedup
	Initial scalar version	-
1	Approximation of magnetic field	50
2	Optimization of the algorithm	35
3	Vectorization (single precision)	4.5
4	Porting to Cell processor	1.5
5	Parallelisation on 16 SPEs (2 Cells)	10

Stages 1–2: 1750

Stages 3–5: 67.5

Training

- Four half days:
  - C_t (C-throughput): Hans Bapst/Intel
  - Vc (Vector Classes)/SIMD: Mattias Kretz
  - CUDA/OpenCL: David Rohr
  - Multithreading (with TBB): Mattias Kretz
- Full coverage onWeb site (also tutorials):
  - https://www.gsi.de/documents/FOLDER-9871272014070.html

Workshop agenda

- Six half-days with multiple topics:
  - Fixed-Target Experiments
  - Collider Experiments

  - Software Architectures
  - Computer Architectures

  - Reconstruction Methods

  - General discussion and Future plans

CBM (Compressed Baryonic Matter)

- Non-homogeneous magnetic field
- Detector layout:

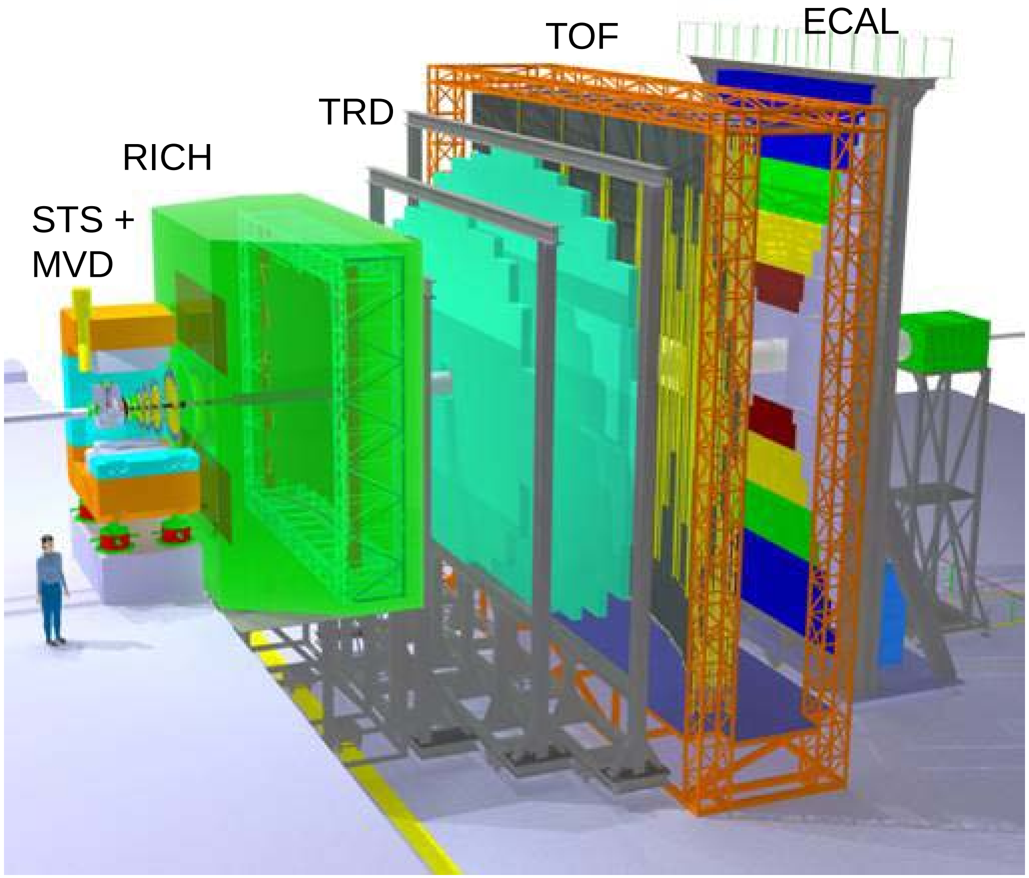

CBM (continued)

- 600 tracks at 10 MHz
- No bunch structure, measurements in 4D
- Requirements:
  - Fast and efficient tracking of all particles
  - Fast and efficient trigger algorithms
  - Fast offline event reconstruction

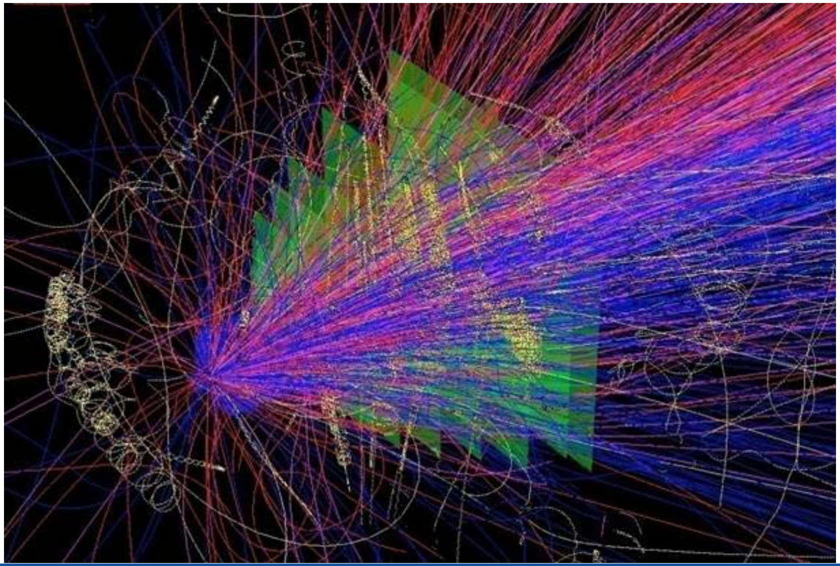

Lots of interesting talks

- A small selection:
  - R.Frühwirth/A.Strandlie: Adaptive methods for reconstruction
  - I.Kisel: Event Reconstruction in CBM
  - I.Kisel: Methods for Track Finding
  - S.Gorbunov: CA HLT Tracker in ALICE
  - M.Al-Turany: GPUs for Event Reconstruction (PANDA)
  - S.Lebedev: Ring Reconstruction in the CBM RICH Detector
  - A.Lebedev: Fast parallel tracking algorithm for the CBM Muon Detector
  - I.Kulakov: Tracking in ALICE and STAR TPC
    - CA Tracker ported (better parameterisation, bug fixing)
  - M.Bach: SIMT Kalman Filter
  - S.Jarp: Harnessing future CPU/GPU hardware
  - K.D.Örtel: Intel’s MIC Announcement
  - H.Babst: Productive Parallel Programming (with C_t)

R. Frühwirth/A.Strandlie

- Adaptive track-fitting:
  - Do preliminary pattern recognition (or none at all)
  - Submit hit/track collective to adaptive fit
  - Inspect posterior weights of hits/tracks and remove outliers
- Various implementations:
  - Elastic arms
  - Elastic tracking
  - Combinatorial Kalman filter
  - Gaussian-sum filter
  - Deterministic annealing filter (DAF)

GPUs in PANDA

- Track propagation in the PANDA experiment
  - Runge-Kutta propagator from Geant3
    - All tracks propagated in parallel
    - Texture memory used for field map optimization

#tracks	CPU (single core, no SIMD)	GPU emul. (on CPU)	Tesla C1060
100	210	160	5.0
1000	210	177	1.9

Time in microseconds/track

Adapted from presentation by M. Al-Turany

CBM RICH Ring Reconstruction

- Method used and preliminary results:
  - Ring Finding
    - Localised Hough Transform
  - Ring Selection
  - Ring Fitting
    - Ellipse-based

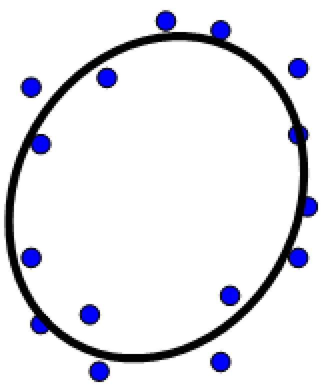

	Time, ms	Speedup
Initial	357	-
Optimization	5.8	62
Parallelization	3	2
Final	-	119

2x Intel Xeon X5550 processors at 2.67GHz (8 cores)

Parallel algorithm are under investigation

Adapted from presentation by S.Lebedev

CA HLT in ALICE

- Based on MC data:

MC, 14 TeV pp events:

HLT Tracker :	Offline Tracker :
Time = **19.6** ms	Time = **66.0** ms
Eff = 99.86%	Eff = 99.94%
Ghost = 0.19%	Ghost = 0.21%
Clone = 9.06%	Clone = 9.30%

MC, 5 TeV Central PbPb events:

HLT Tracker :	Offline Tracker :
Time = **17.6** s	Time = **160.1** s
Eff = 98.15%	Eff = 95.84%
Ghost = 1.66%	Ghost = 1.40%
Clone = 13.22%	Clone = 12.13%

~ linear time dependence:

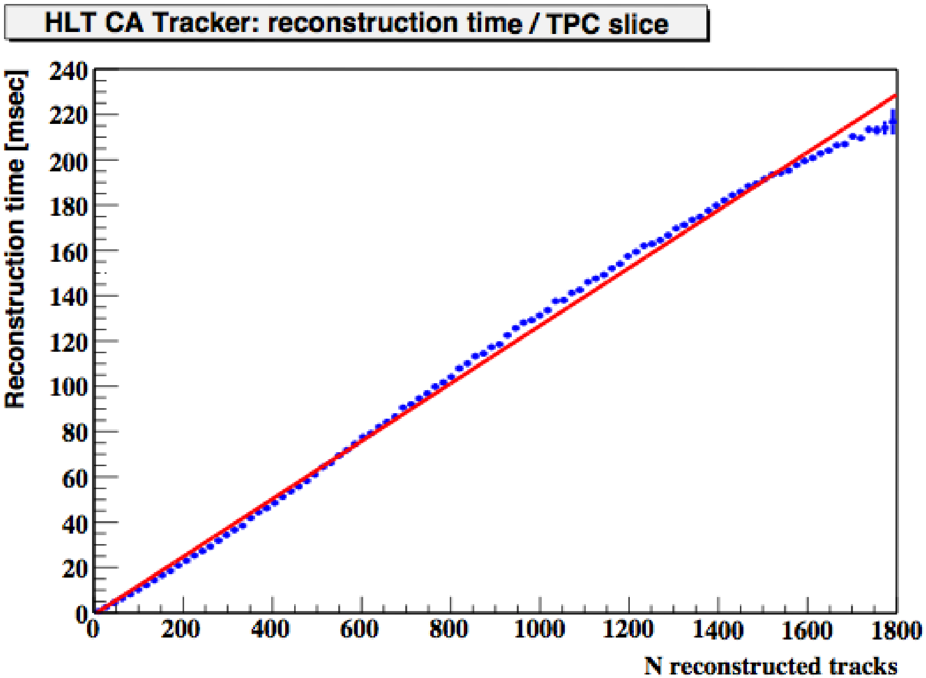

Adapted from presentation by S.Gorbunov

ALICE (1)

- TPC tracking algorithm:

Category of Task	Name of Task	Description on Task
	(Initialization)	
Combinatorial Part (Cellular Automaton)	I: Neighbors Finding	Construct Seeds (Track Candidates)
	II: Evolution	
Kalman Filter Part	III: Tracklet Construction	Fit Seed Extrapolate Tracklet Find New Clusters
	IV: Tracklet Selection	Select good Tracklets
	(Tracklet Output)	

Adapted from presentation by David Rohr

ALICE (2)

- Speed comparison CPU versus GPU (for heavy ion event)

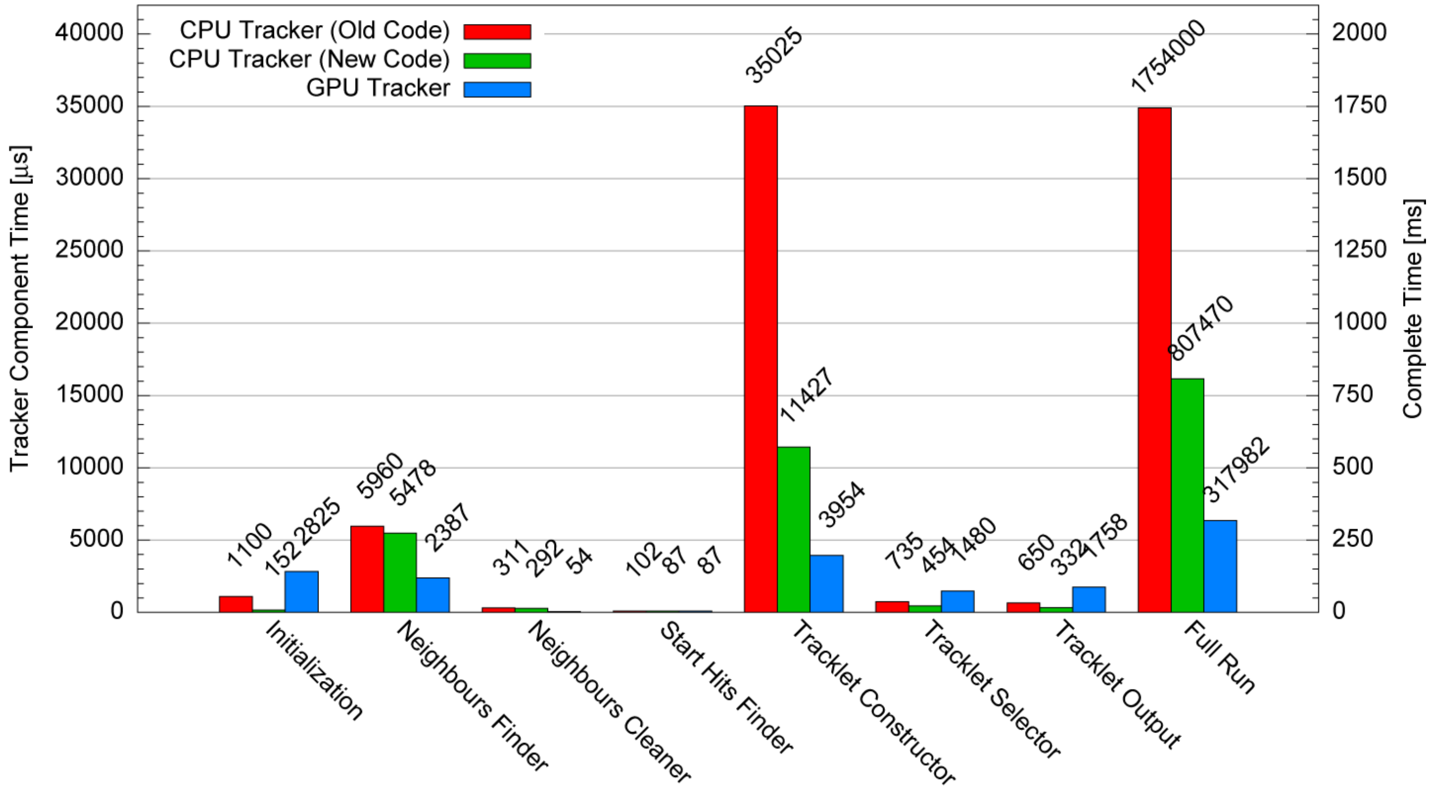

Twofold performance increase on CPU
2.5-fold performance increase compared to CPU

Adapted from presentation by David Rohr

Workshop conclusions

- Agreement on a common approach
  - Host experiments: ALICE, CBM, PANDA, STAR
  - Extend functionality of SIMD Kalman Filter (DAF, etc.)
  - CA Track Finder ported to C_t in parallel to Vc-based version
  - Vertexing package (KFParticle, used in ALICE, CBM, STAR) will be extended
- Next meeting in a year’s time: At CERN ?
- CERN openlab’s interest:
  - Realistic benchmarks for multiple platforms
    - Written with scalability in mind