FAIR Event reconstruction workshop

Final Version (3.0) of Minutes, 28.11.2012

29.10.2012 at GSI Darmstadt

Participants (alphabetically):
Mohammad Al-Turany, Denis Bertini, Volker Friese, Tetjana Galatyuk, Klaus Götzen, Norbert Herrmann, Marian Ivanov, Wolfgang Kühn, Sören Lange, Volker Lindenstruth, Yvonne Leifels, Yutie Liang, Silvia Masciocchi, Peter Malzacher, Johan Messchendorp, Walter Müller, Peter Senger, Christian Sturm, Joachim Stroth, Florian Uhlig and others

Silvia Masciocchi (Alice overview)

discussion points:

in Alice TPC cluster finding is performed online.
raw TPC hits are not stored anymore on tape, but only clusters.
next step: track finding online, and remove clusters
(reduce data to tape even more).

Alice HLT hardware so far uses nVidia GPUs and CUDA.

in the moment, online reco @ ALICE is factor ~10 faster than offline reco.
For Alice upgrade factor ~100 is required.
comment by Wolfgang Kühn: Moore law helps, but probably not factor ~100.
huge task to achieve the required speed-up.

a problem for Alice offline reco:
software needs too much memory: 5 Gbytes per 1 event.
TPC needs 1.5 Gbytes.
TRD needs even more than TPC (reason not clear yet).
calibration tables require large memory.
-> proposal to use HLT tracks as seed (in future).
might help to reduce the required memory.

calibration of TPC drift velocity is important.
comment by Volker Lindenstruth:
it already happened during Alice HLT operation,
that tracks were recontructed "outside the TPC".
problem: drift velocity is strongly time dependant.
For proper online v_drift must even be extrapolated
into the near future (during an ongoing run).

question by Volker Friese:
in an former Alice report a statement was made,
that calibration is very stable and parameters are used for ~1 year.
answer: one must distinguish
-> alighment parameters, only changed on timescale of ~1 year
-> drift velocity is changed daily or even hourly.

question by Peter Senger:
20 kHz upgrade requires new hardware?

theoretically yes, but practically not feasible to exchange
all the frontend electronics.

---

question by Peter Malzacher:
can online tracks be re-calibrated later?
(if TPC hits are removed before tape)?
in the moment, data are not stored in a way to do easy re-calibration
(only easy if hit data would be kept).
offline reco must be completely re-written
estimated schedule ~2 years. non-trivial work.
related Question by Florian Uhlig:
how many FTEs are working on it?
Volker Lindenstruth: availability of people is one of the biggest concerns.

---

question by Walter Müller:
20 GB/s taping rate is limit for Alice upgrade,
what 's the reason?
(right now 4 GB/s o.k.)
-> limit is link to computer center (12 GB/s)
assumption is to reach compression 20->12 GB/s.

---

question by Wolfgang Kühn:
space charge effects?
in the moment, effect is small ~100 um
(occupancy ~5% in Pb+Pb)
Marian Ivanov: code is existing, but not needed so far

.

Joachen Markert (Hades overview)

discussion points:

Hades is writing online DSTs,
but quality physics results require detailed calibration.
Publication within 2 weeks after data taking (as for some Alice papers)
would not be feasible in Hades.
quality for online trigger -> efficiency maximized.
quality for physics paper -> purity maximized.
approach is different.

---

double hit probability in Hades in Au+Au is quite high, per 1 wire <30%
(example from Au+Au at 1.23 GeV).

---

question by Volker Lindenstruth:
for central collisions 30% ghosts?
yes, but this is stage before (clean) track candidates.

---

one of the difficulties in Hades:
vertex resolution (as no vertex detector installed)
question by Wolfgang Kühn:
CPU time comparison between 4th-order-Runge-Kutta (RK4) and Kalman ?
hard to say. one method requires requires segment fitter pre-step, other not.
coordiate transformations require time.
theoretically (according to number of algorithm steps) should be 1:1.
~10 events per 1 second
For RK4, ~30% of CPU time is spent inside the algorithm.

Comment by Volker Lindenstruth:
how long would by all hades data re-processing at GSI farm?

-> only 1 week. not considered a problem.

---

problem is extrapolation through the kick B field
(because there is no detector).
if Kalman input parameters are wrong,
extrapolation to outer 1st plane if off by 50 cm

---

what needs to be improved for Hades for SIS100 ?
so far no timestamps in data words.
but: priority is present, not future
(i.e. to do physics experiments with existing setup).

Klaus Götzen (Panda overview)

Discussion points:

---

comment by Denis Bertini:
seems a very complex "trigger" scheme.
"trigger" is given by algorithm.

---

question by Peter Senger and Walter Müller:
is there a simple physics case for day zero ?
a.) first candidate for scalar glueball -> 3 pi0
(so final state is 6 photons, no charged tracking required)
b.) J/Psi trigger is easy
(so any charmonium state which decays into J/psi is "easy")

10^7 is theoretical reduction factor for rare channels.
what is practical reduction factor ?
-> 1/1000 is planned.

---

question by Silvia Masciocchi:
deadtime?
time distribution system SODA
- each frontend is connected to SODA
- each piece of data has a timestamp
if a detector has deadtime, then simply a data package are missing.
however, based on data of the other detectors
for a given timestamp, it can be concluded that data are mssing.

Sören Lange (Panda FPGA, part 1)
Yutie Liang (Panda FPGA, part 2)

combined discussion for both Panda FPGA talks:

---

resolution of online tracking is crucial.
yes, but Panda measures exclusive channels,
so constraints can be used.
Example: p pbar -> J/Psi pi+ pi- (shown in the talk of Klaus Götzen)
instead of calculating [e+e-] invariant mass for J/Psi
(using the high momentum leptons with limited tracking resolution)
we can see J/psi peak in [pi+pi-] missing mass,
which has factor ~10 better mass resoultion.

Pile-up filter ?
If we consider one 220 ns readout frame of the Panda STT,
we get pile-up hits e.g. from later collisions.

However, these pile-up hits belong only to early drift times
of the next STT frame.
-> these are tracks with missing hits
(only the short drift times contribute to hits,
and hits with longer drifttimes are missing).
-> incomplete tracks, which could be cleaned up
in a 1st step of the online system.

Mohammad Al-Turany (Panda GPU)

discussion points:

Mohammad emphasized that the effort is mainly
trying to solve the parallelization concepts
(i.e. how to parallize a given problem)
for a specific hardware platform
and not so much the coding in a special language
(i.e. CUDA or OpenCL)
in other words: finding out the syntax of a given function
is a minor effort, and function concepts in CUDA and OpenCL are similar.

comment by Volker Lindenstruth:
CUDA was used in Alice only because at the time of writing OpenCL
was not available yet. The transformation to OpenCL is a decision for
the upcoming long shutdown but changing CUDA code into OpenCL code for Alice
HLT turned out to be much more work than anticipated.

Volker Lindenstruth pointed out that the presentation of Mohammad is his personal opinion and does not represent the official IT standpoint.

Volker Friese (CBM overview)

discussion points:

MVD is slowest detector in CBM.
(20 us readout time of MIMOSA pixel).
-> will not be used for charmonium data taking
(MVD will be removed from setup)

Ivan Kisel (Alice and CBM GPU)

discussion points:

efficiency for secondaries is quite low -> ~50% -> why?
track are in outside detector region, not enough hits`on track

comment by Sören Lange:
HERA-B had strongly inhomogeneous B field
parabola track model is only valid for pure dipole field.
-> might be the reason why Hough transform failed.

cellular automaton has input (x,y,z) spacepoints.
application for Panda STT (drift circles) would require some thinking.

CA is "SIMD" (single-instruction-multiple-data)

fits well to GPU.

CPU time of 47 ms/event contains matrix operations in Kalman filter

question by Mohammad Al-Turany:
is the code available? stand-alone or in FairRoot ?
MUCH and TRD tracking is in the repository.
other packages are available by request as stand-alone versions.

Final discussion about conclusions

(where do we go from here?)

1. CA as a generic package is a nice approach.
investigation of applicability to Panda drift circles is appreciated.

2. for all experiments:
hardware can not be fixed now.
market develops fast and
hardware in 5 years will be different from now.
investigation must continue.

3. Investigation of Hough transform approaches
on both FPGA and GPU should continue.
lots of work went into it.
Panda participants would like to ask GSI IT to continue development
of GPU based Panda tracking algorithm to be able to compare to the
FPGA implementation, in particular with the goal to also include
the newer developments in Giessen (drift circles, t_0 determination etc.)
Volker Lindenstruth points out that IT has very limited resources
and has to plan. Therefore it is requrested to define a work plan
with milestones for such projects.

4. In order to assemble a demonstrator system it would be highly desirable to have
the event server system (proposed in the talk by Mohammad Al-Turany),
which can provide realistic simulated data to different online computing nodes by:

A. High speed optical links (PEXARIA2) to prototype DAQ and event filtering boards ("ATCA compute nodes").
B. High speed optical links (PEXARIA2) to GPU directly through GPU-Direct protocol
C. Network to CPU/GPU Farms

Comment by Norbert Herrmann:
such a demonstrator system would also be important for CBM.
Volker Lindestruth points out that these efforts are important but point A is
considered as DAQ effort and are not part of the IT mandate. This subject belongs
more into the EE department.

5. CBM needs time-based simulation for 20 us MAPS readout frames.
However, approach is different (comment by Walter Müller):
Panda has time and space data mixed,
CBM can add time to the track model.
one should keep in mind, that in CBM STS
all particles are almost speed-of-light
(no drift time variations).

6. next workshop:
in ~3 months.
date will be determined by doodle.
Peter Malzacher asks if for the next meeting the GSI IT could be invited too.
Volker Lindenstruth pointed out that the meeting was open.

---

minutes taken by Sören Lange
This is 3rd version, 28.11.2012