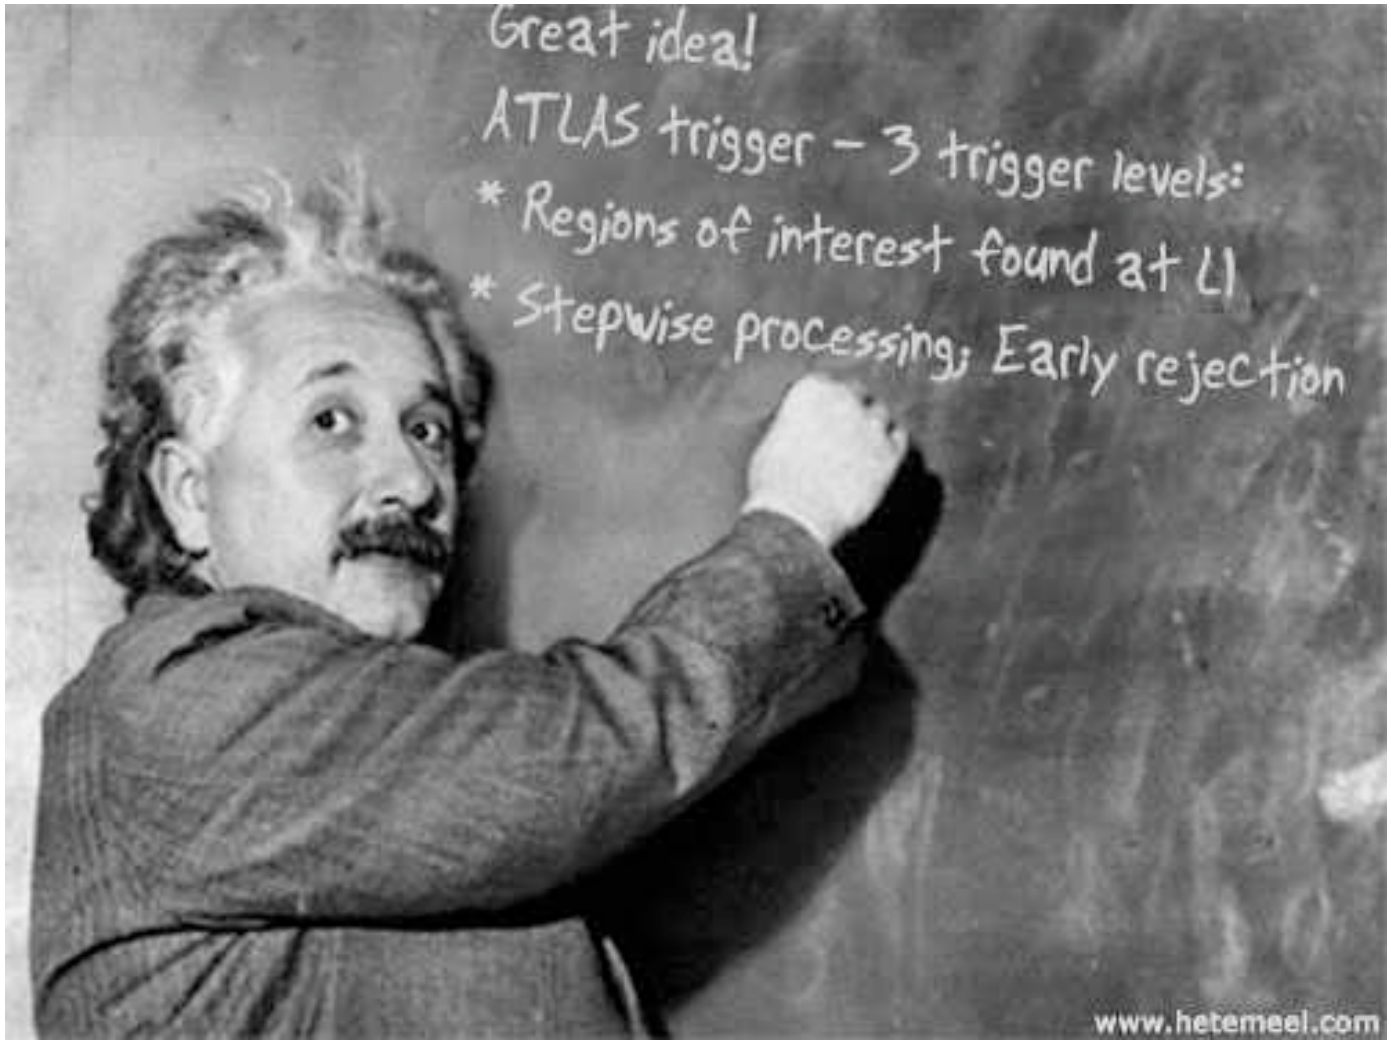

Status of trigger software and EDM

Royal Holloway University of London

Ricardo Gonçalo, RHUL – on behalf of ATLAS TDAQ

- Introduction
- Preparing for LHC data
  - A few examples among many
- EDM migration and core software
- Trigger configuration from database
- Slice activities
- Online performance
  - April technical run
  - M7 technical run

Introduction

- Trigger reduces 40MHz input rate to 200Hz recording rate
- Selects events for physics studies and detector alignment / calibration
  - First step in any physics analysis
- Assigns events to streams for easy access
- Selects (express stream) events for rapid feedback before reconstruction
- Provides real-time data quality monitoring
- Work spans physics groups, combined performance, data preparation… and trigger groups
  - Usually organised in “slices”: muons, electron, photon, jets, taus, Bphysics, Bjets, Minimum bias, cosmics

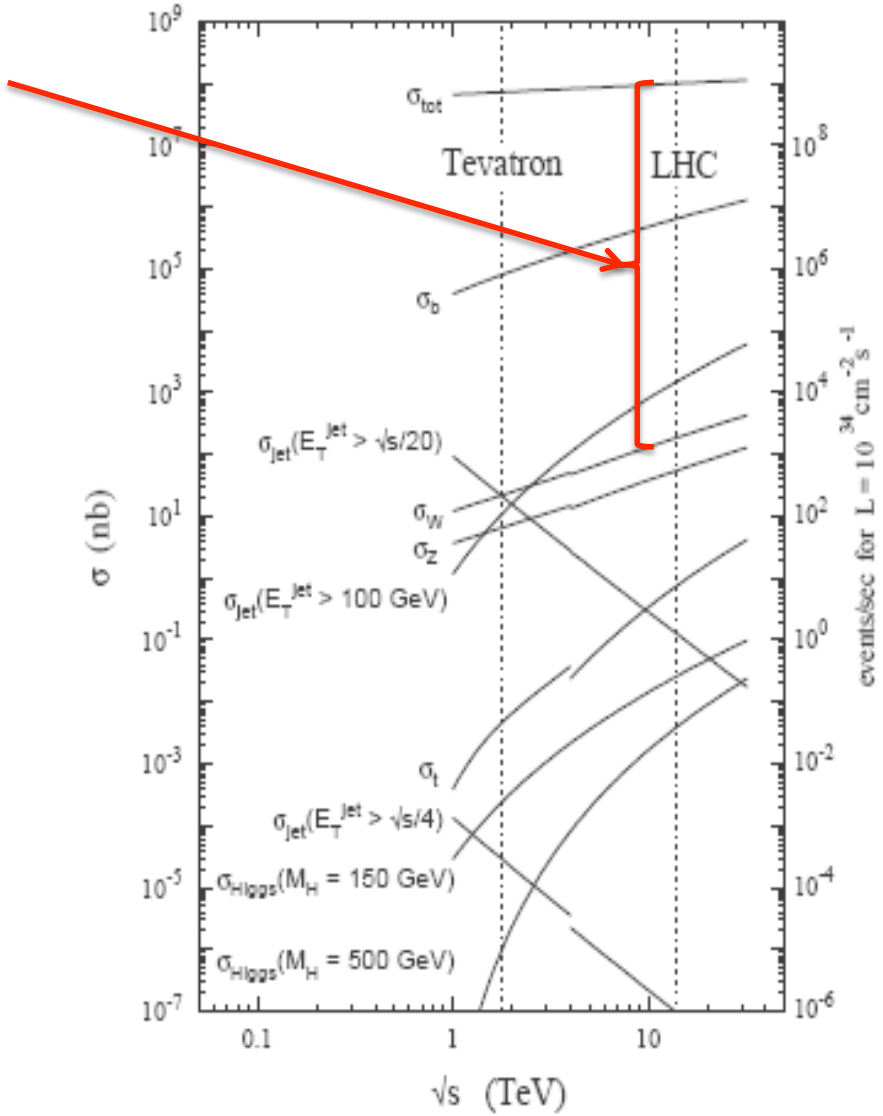

Selection method

Event rejection possible at each step

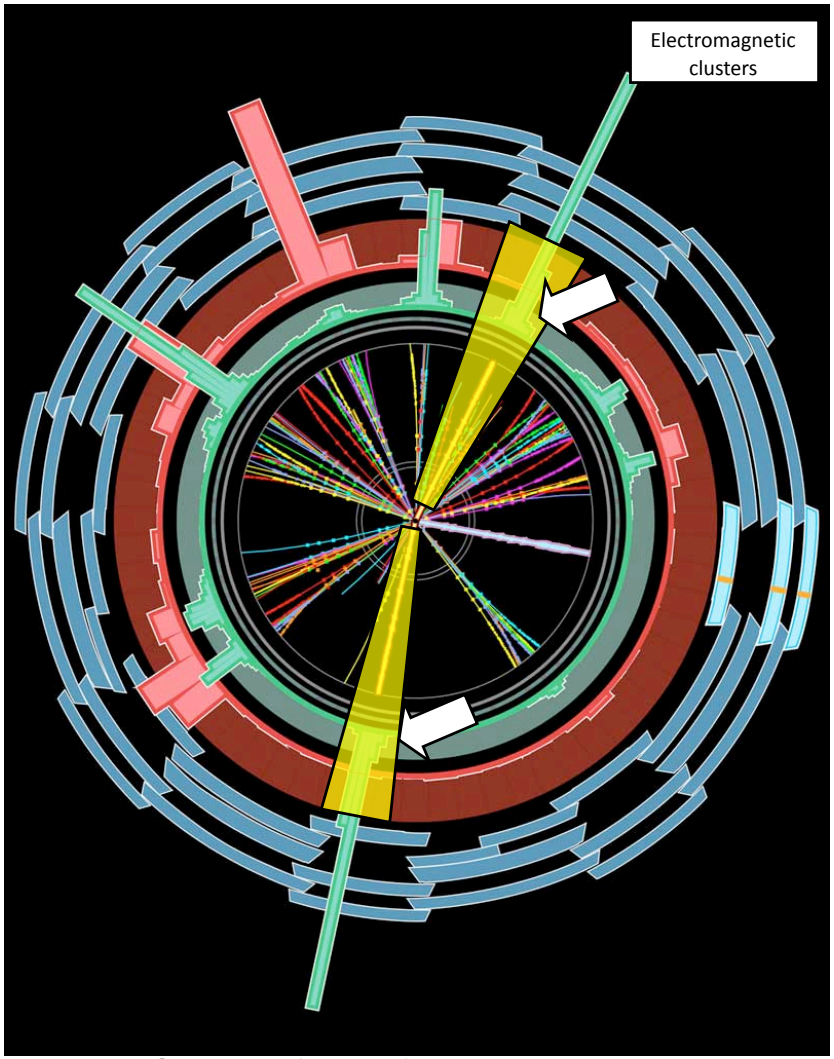

Level1 Region of Interest is found and position in EM calorimeter is passed to Level 2

Level 2 seeded by Level 1
Fast reconstruction algorithms
Reconstruction within RoI

Ev.Filter seeded by Level 2
Offline reconstruction algorithms
Refined alignment and calibration

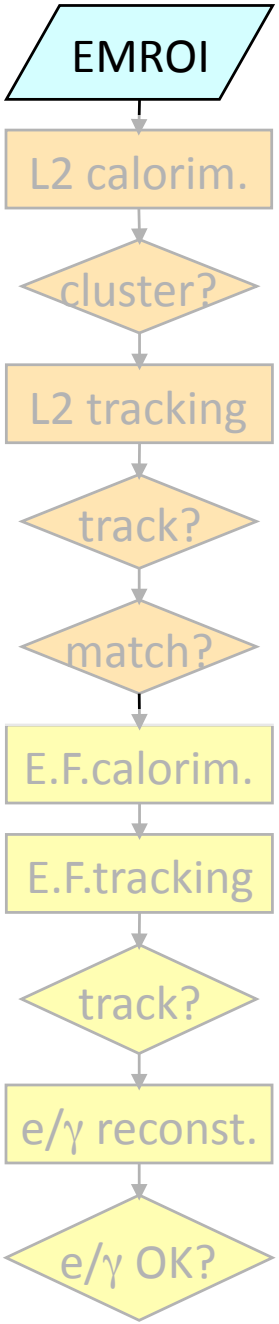

Preparing for LHC data

- A lot of effort is being put into preparing for the 2008 run:
  - Adding last missing components
    - e.g. MBTS minimum bias trigger, HLT forward jets, beam-spot finding, new developments in TRT standalone tracking
  - Improving robustness of trigger code against likely problems
    - Tested code with realistic data: residual mis-alignement, displaced beamspot, pileup, data-flow errors, beam-induced backgrounds
  - Improving flexibility to adapt to unexpected conditions
  - Establishing initial trigger menu (talk by Brigitte Vachon)
  - Improving monitoring capabilities (talk by Cristobal Padilla)
  - Testing code with real data in M weeks
  - Evaluating performance with complete menus in technical runs

- Three ad-hoc workshops since March dedicated to online running issues:
  - Trigger Robustness – 4th March
  - Trigger Data-Quality Assurance – 6th May
  - Trigger Efficiency Determination from Data – 1st July (talk by Teresa Fonseca Martin)
  - Lots of discussion and much progress made!

Algorithm Error Codes

- Trigger algorithm error codes are used to modify trigger behaviour
- Need to be flexible, to e.g. avoid harmless errors sending many events to debug stream (expect the unexpected!)
- Solution: re-map error codes until algorithms can be changed or problem fixed
- HLT::ErrorCode has Action (e.g. ABORT_CHAIN) and Reason (e.g. MISSING_ROD)
- Action and Reason can now be re-mapped for each HLT algorithm
  - E.g. (ABORT_CHAIN, MISSING_ROD) can be re-mapped into (CONTINUE, MISSING_ROD) to avoid interrupting processing of RoIs
- Original and re-mapped error codes monitored for every chain
- Time overhead is 2μs per algorithm with remapped codes

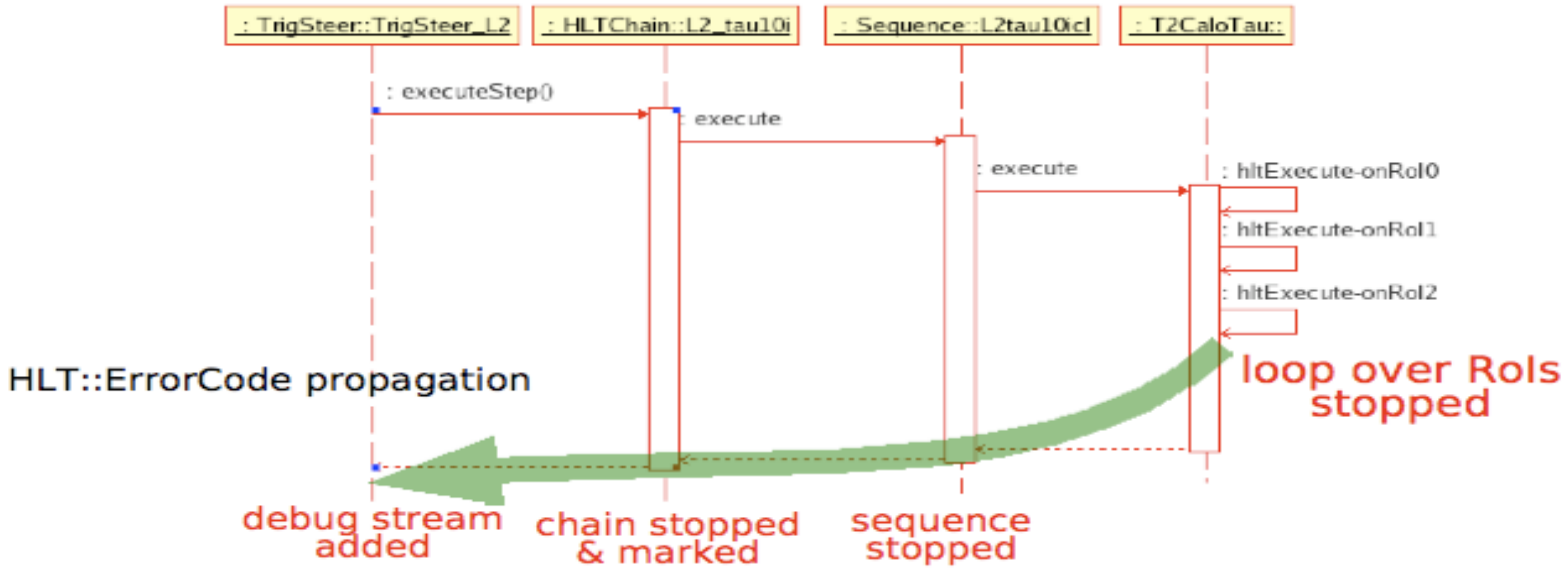

Displaced beamspot

- Lots of work has been devoted to verifying that the trigger is robust against several possible error sources

- The beam-spot displacement wrt the Atlas reference frame was found to be a possible source of inefficiency
  - Not clear what to expect from LHC - it may be a non-issue, but better be prepared

Two aspects:

- Tracking algorithm robustness at L2
  - Robust tunings exist for the most commonly used tracking algorithms
  - Their performance is being studied
- Determination of beam position for B-tagging and impact-parameter measurement
  - Work is advancing on the online determination of the beamspot position for each fill

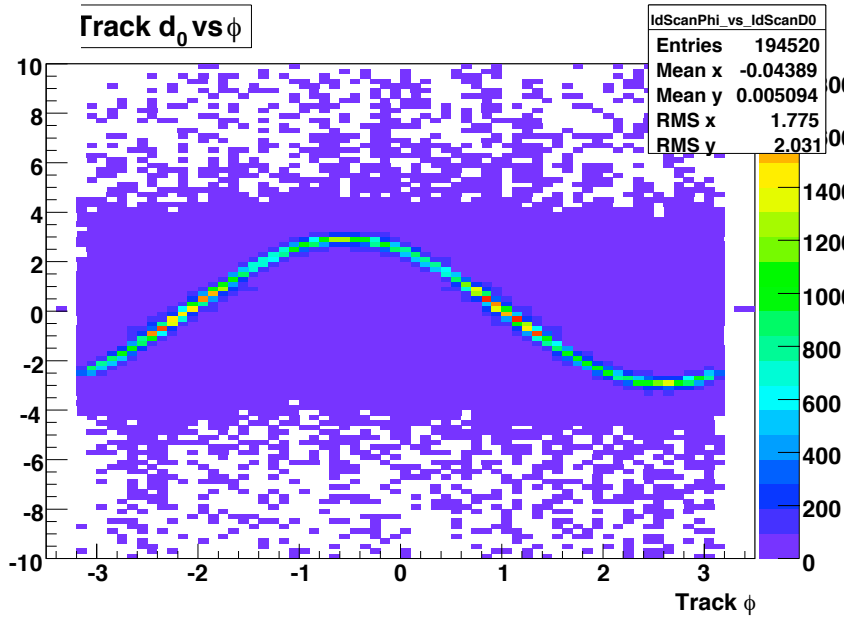

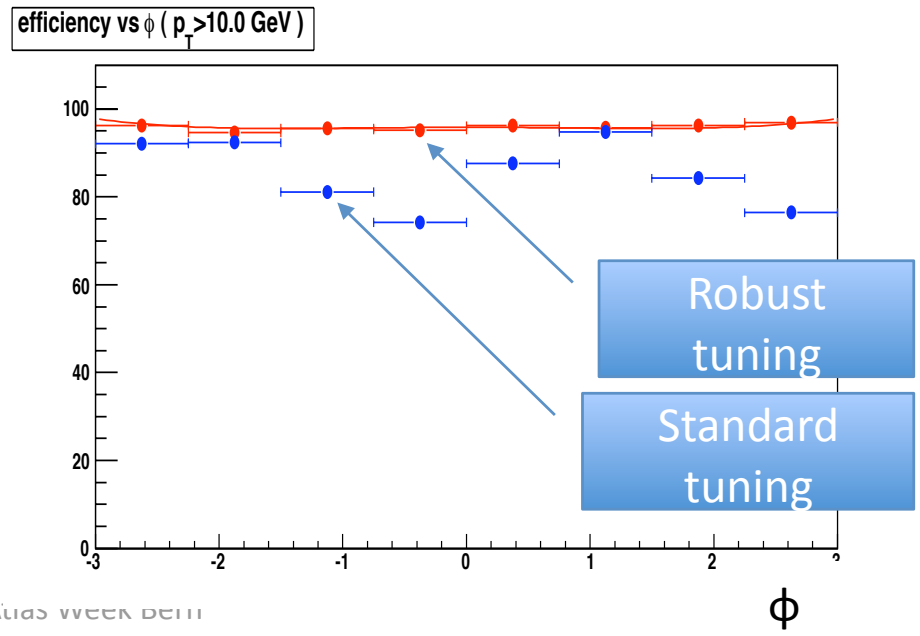

Event Data Model (EDM)

Recent changes of EDM infrastructure:

- StoreGate now used online
  - Allows use of **ElementLinks** as persistent pointers in Level 2 EDM
  - Allows for schema evolution for trigger bytestream
  - Time overhead found to be ok

- Changed to **flat trigger containers**:
  - One container **per event** for each combination of **object class** + **feature extraction algorithm** instead of one container per Region of Interest (RoI) and algorithm
  - More efficient and natural for POOL persistency
  - Allows easier access with AthenaRootAccess

- New **Serializer**: allows EDM objects to be serialized including ElementLinks

- Changes to **TrigNavigation**:
  - Handle new container structure (maps RoIs into flat containers) and new Serializer

- Migration:
  - Timing not ideal, but avoids more painful migration later with real data
  - Mostly transparent to trigger online algorithms (Feature Extraction, Hypothesis algos)
  - Migration went quite smoothly (MIG2 nightlies) and now in 14.2.10
  - ElementLinks introduced at L2 (TrigL2BPhys, TrigElectron, TrigPhoton, CombinedMuonFeature)
  - To be done: run serializer on persistent bytestream classes, for schema evolution; propagate use of ElementLinks to other L2 classes

- EDM migration allowed improvements in data size on file
- Overall size of trigger data depends strongly on data type and on menu
- Running menu for $L=10^{31}$ $cm^{-2}s^{-1}$ (no Bphysics):
  - AOD total size: 48 kB/event
  - ESD total size: 86 kB/event
- The HLTResult may be reduced further to 1-2 kB/event by slimming out navigation information - (e.g. for inclusion in DPDs)

13.7 HLTResult
12.1 Rec::TrackParticleContainer
9.9 TrigInDetTrackCollection
3.6 CaloClusterContainer
2.1 TrigElectronContainer
1.9 egDetailContainer
1.2 Trk::VxContainer
0.7 TrigTauClusterContainer

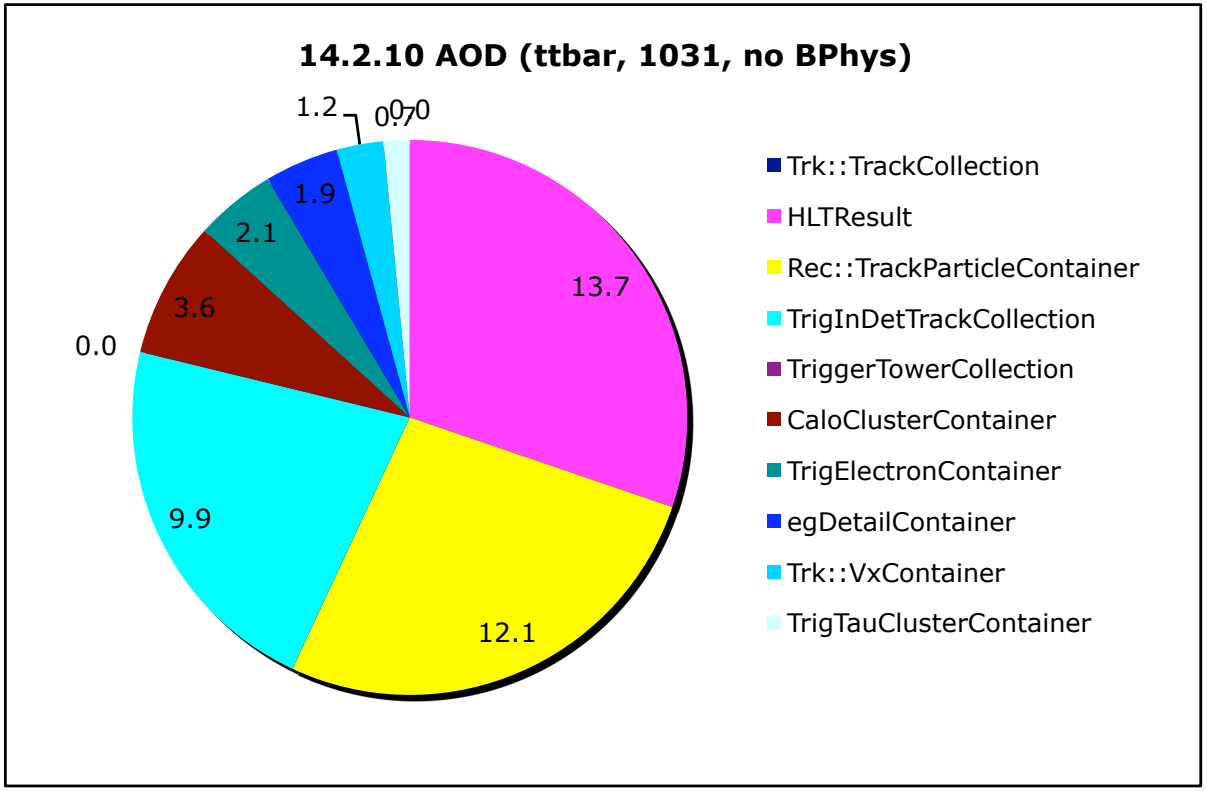

Configuration from the TriggerDB

Configuration Data Flow

Preparation

TriggerDB
All configuration data

Configures

Stores decoded Trigger Menu

Data taking

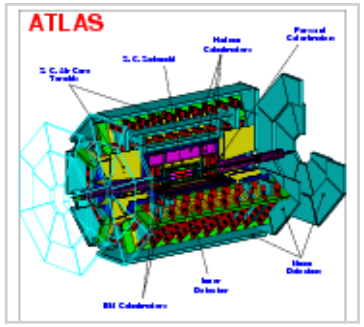

Online Conditions Database

Encoded trigger decision
(trigger result from all 3 levels)

Decoded Trigger Menu

Reconstruction/ Trigger aware analysis

Trigger Result

- passed?, passed through?, prescaled?, last successful step in trigger execution?

Trigger EDM

- Trigger objects for trigger selection studies

Trigger Configuration

- Trigger names (version), prescales, pass throughs

access through TrigDecisionTool

ESD
AOD
DPD
TAG

With decreasing amount of detail

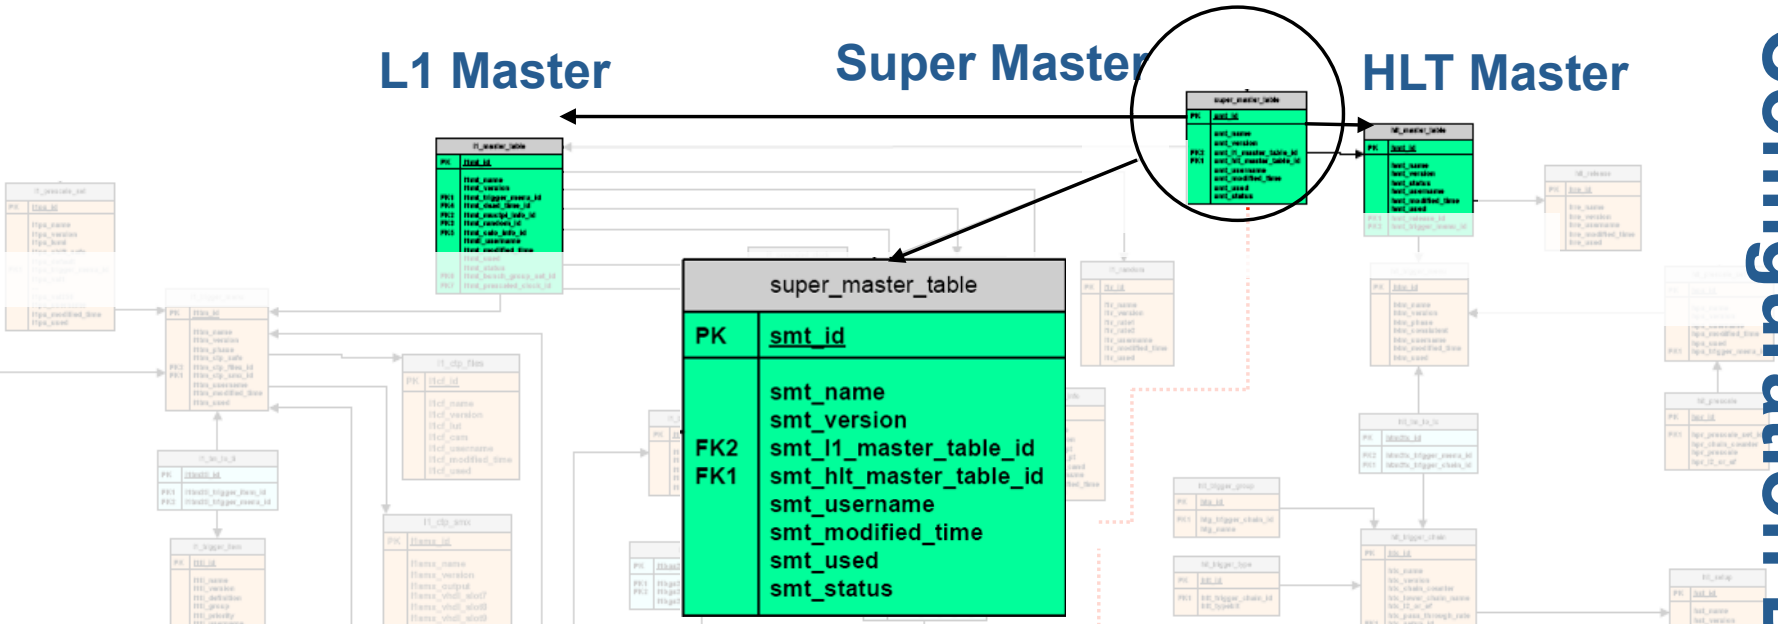

- *SuperMaster* table contains links to L1 and HLT *master* tables
- *HLTMaster* table defines complete HLT configuration
- *L1Master* table defines complete L1 configuration: the L1 menu and CTP config

⇒**the supermaster ID ("key") chooses matching L1 + HLT trigger menus**

Aim to configure the trigger **from DB also outside point1** - trigger validation on the CAF; trigger simulation; trigger development

The advantage is to access the same trigger configuration as was used for data taking
➔ easy to achieve reproducibility
➔TriggerDB to be available at Tier0/1 (Oracle), and Tier2 (SQlite)
➔First running version in 14.2.10

Trigger slice performance

- Trigger slices focusing on preparations for data taking:
  - Lots of work developed to test selection robustness
  - Migrating code to it's final form – in some cases trying out new algorithms
  - Work Offline Data Quality monitoring developed within each slice

- Much still to be done – some of which needs experience with real data:
  - What are the most common online errors
  - What are the most important sources of fakes
  - etc

Muon Slice

Latest developments:

- Migrated data preparation tools to final format
- Finished implementing EF algorithm
  - And associated EDM classes and T/P converters
- Optimisation of calorimeter and tracking isolation at L2
- Started development of offline Data Quality monitoring tools

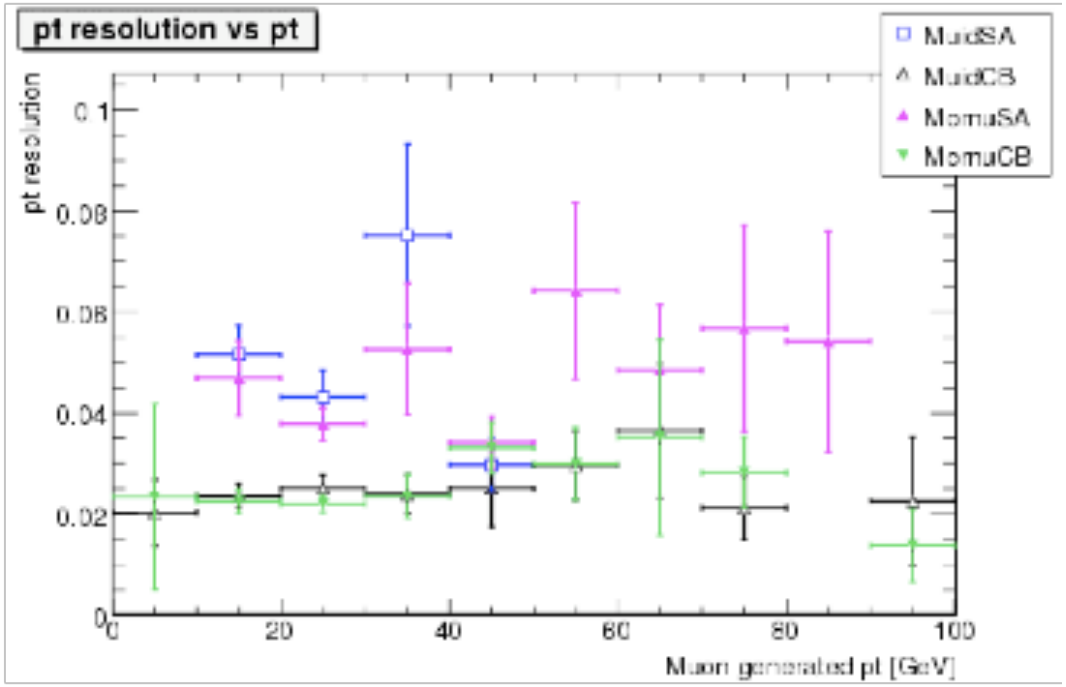

e/gamma slice

Latest developments:

- Motivated improvements in L2 calorimeter reconstruction:
  - L2 calorimeter reconstruction tuned to improve search for nearby clusters (relevant for triggering on J/Ψ)
  - Cell energy correction (longitudinal weights)
  - Energy correction for electrons in the calorimeter crack
- EF using individual cuts instead of isEM to ease early data studies
- Work progressing well for Data Quality monitoring
- Several studies on code robustness
- Studying performance with standalone TRT tracking and with backtracking

	Loose cuts	**Tight cuts**
ForwardTracking (input 9650 events)	8983 (93%)	8094 (83.8%)
TRTOnly (input 11500 events)	8747 (76.1%)	0 (0%)

Tau Slice

- Latest developments:
  - By default, run tauRec at EF (wrapped) – currently testing TrigTauRecMerged performance and, in parallel, use of TopoClusters at Event Filter
    - Some difficulties caused by event filter tracking sequence
    - Differences in TopoCluster building between trigger and offline
    - Degradation in resolution of 3-prong tau candidates
  - Improvements in L2 tracking to improve fake tau rejection

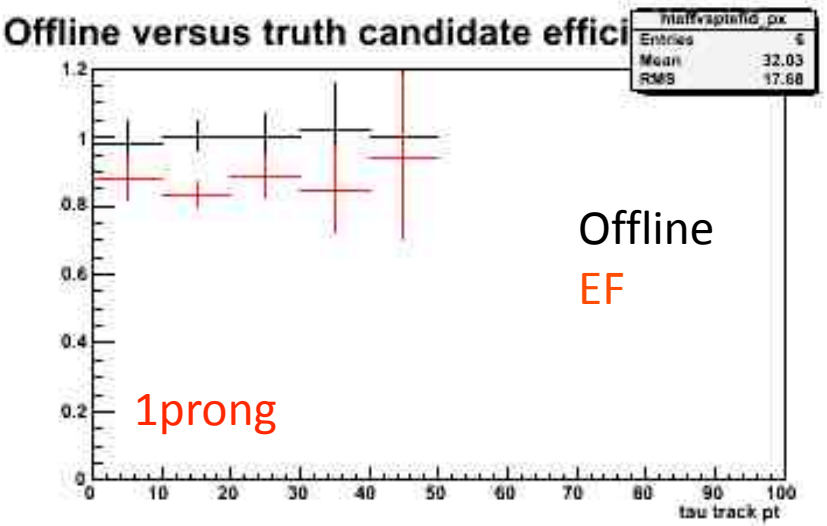

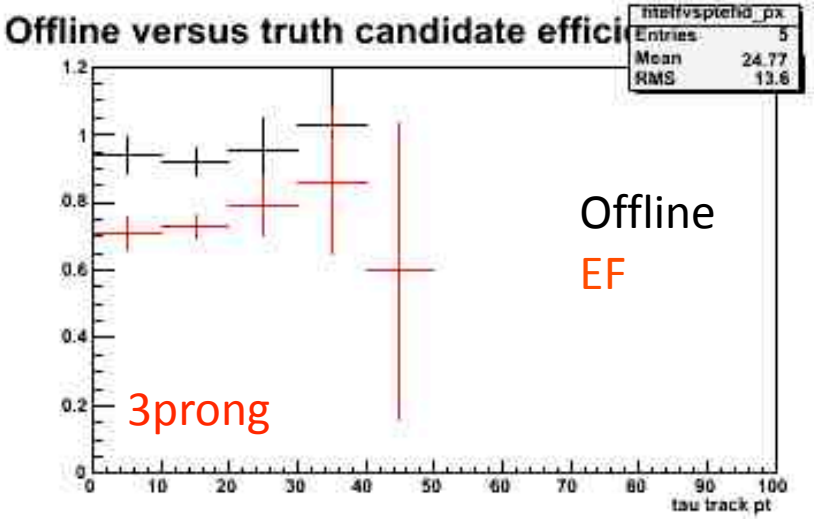

- Detailed commissioning plans:
  - https://twiki.cern.ch/twiki/bin/view/Atlas/TauTriggerCommissioning

B-physics slice

- Progress in extending selection to EF
- Great reduction in EDM size on disk after EDM migration: 14kB/ event to 1.9 kB/event (top events)

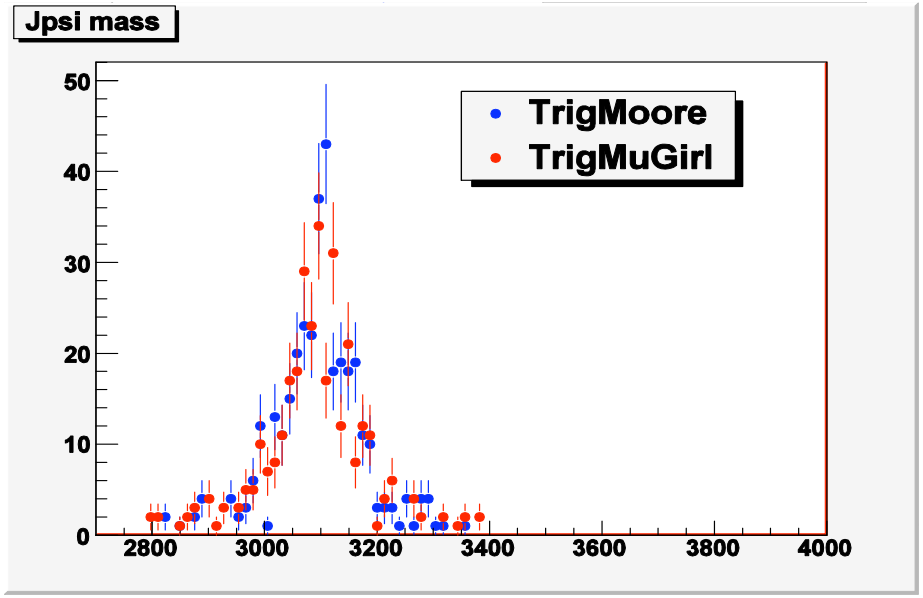

B-jet slice

- Investigated sensitivity to beamspot movement at L2
- With nominal tracking tune:
  - Sensitivity starts at 25µm
  - No discrimination at L2 for displacements > 100µm

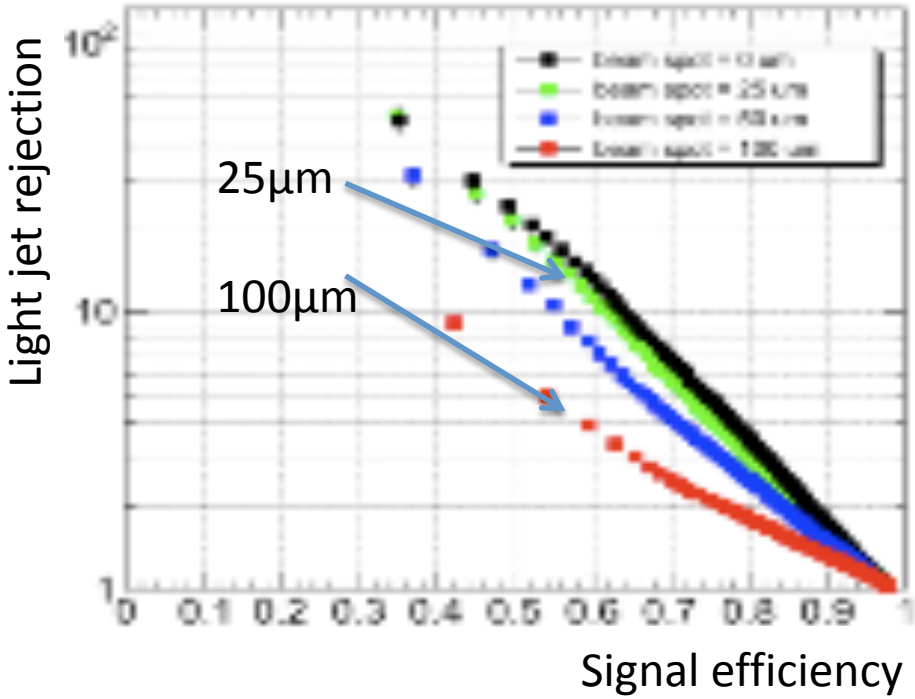

Jet slice

Latest developments:

- Tested slice by re-running on M7 bytestream files
- Work progressing on forward jets in L2 and EF
- Ongoing work on L2 calibration, including forward jets

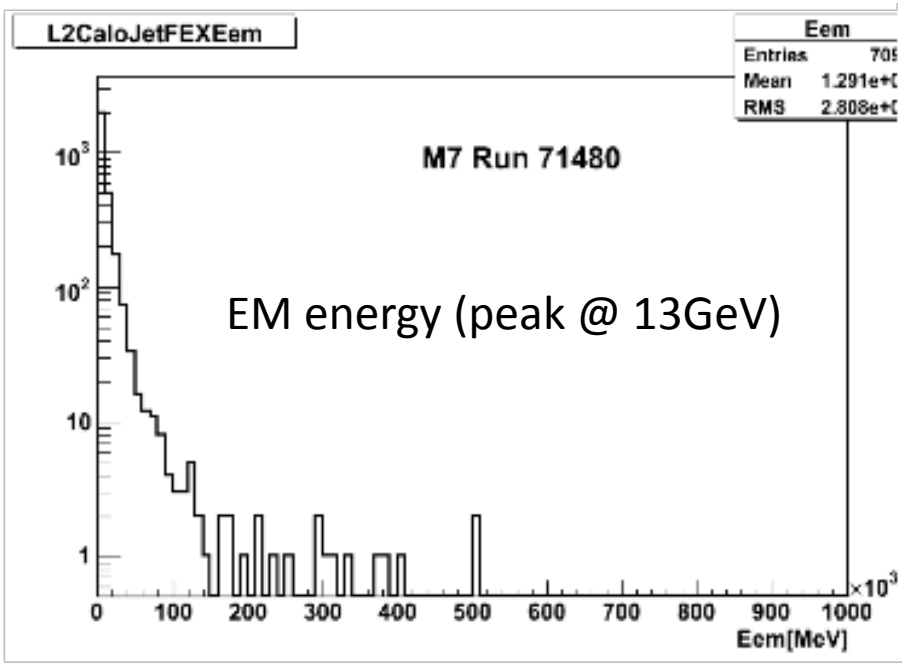

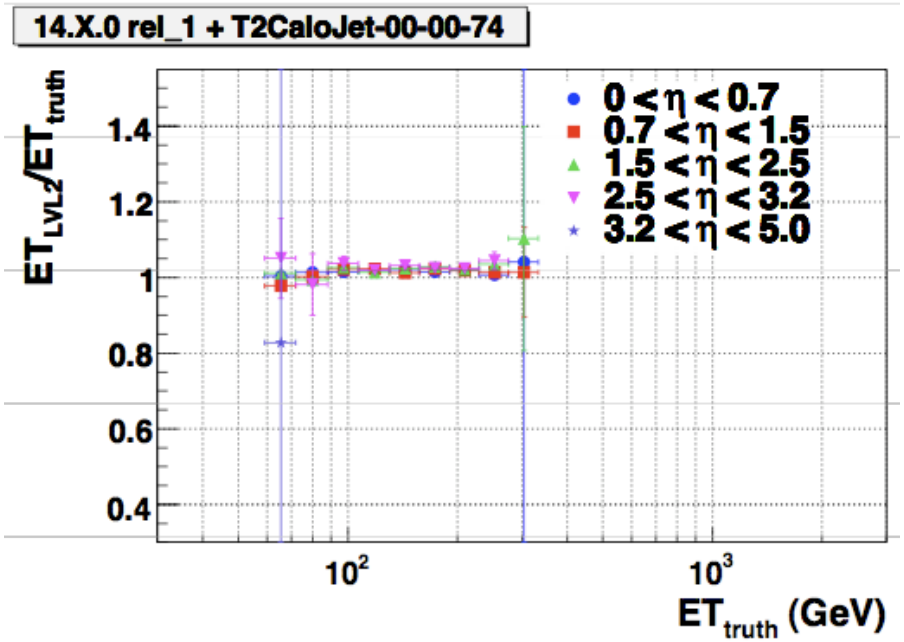

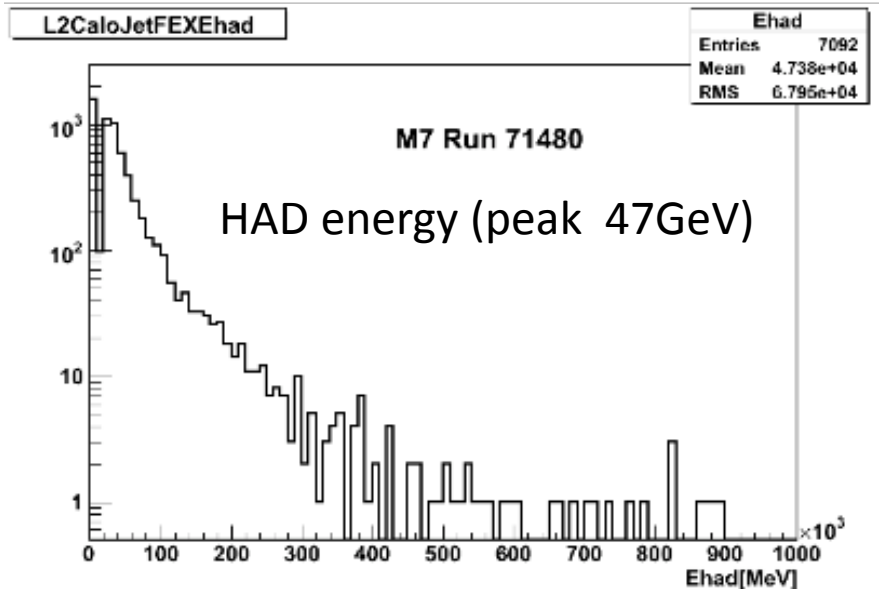

Missing ET

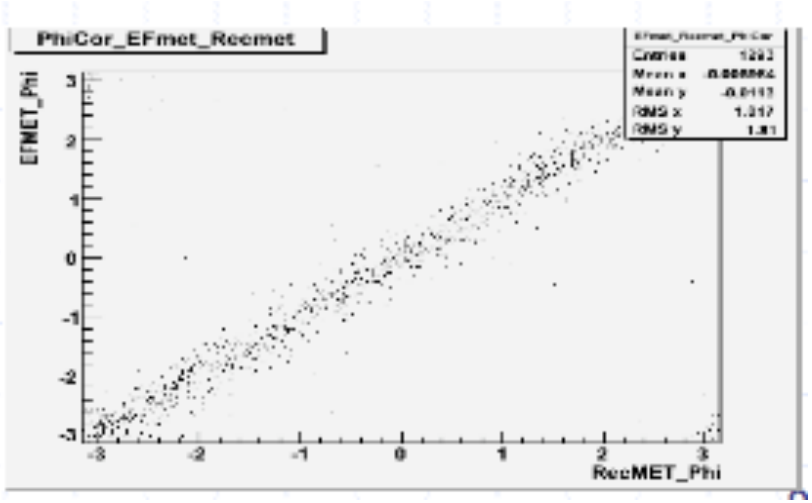

- HLT algorithm runs at EF in “unseeded” mode
- Algorithm:
  - Runs by default from calorimeter cells
  - Can also use partial energy sums done in the Lar Front-End Boards to speed up algorithm
  - Muon correction applied in hypothesis algorithm
- Remaining issues:
  - Calibration:
    - 2 constants per calorimeter sampling + 2 constants for muon component
    - Studying methods to calibrate from real data
  - Monitoring:
- For 2008 run, plan to exercise slice starting from simple and robust algorithm
- For 2009 run, plan to make algorithm more robust against detector effects:
  - Channel-based checks against noise
  - Add intelligence to algorithm to reject main sources of fake E_T^{miss}
  - Optimize data access to improve time performance (essential for high luminosity)

Minimum Bias Slice

- In early running, important to trigger on p-p collisions
- Later, all bunch crossings will have at least one p-p interaction
- Triggers based on:
  - Inner detector (SpacePoints counts, track counts)
  - MBTS – Minimum Bias Trigger Scintillators – Scintillators on the inside of endcap calorimeter
  - BCM – Beam Conditions Monintor – small detectors attached to ID supports
  - LUCID – luminosity detector near beampipe 17m from IP
- Many recent developments of BCM and Lucid
- Ongoing work on MBTS timing to discriminate against beam halo

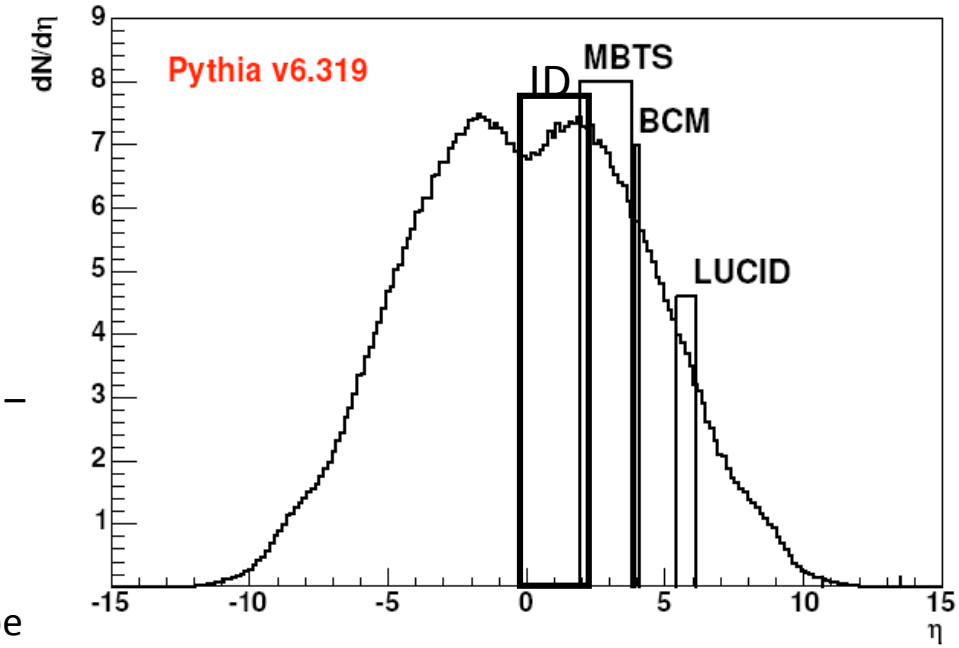

- BCM performance study ongoing
- Lucid studies waiting for simulation validation

Online performance

Technical runs and Mx cosmics runs

Technical (and 24h) runs

Monte Carlo Events

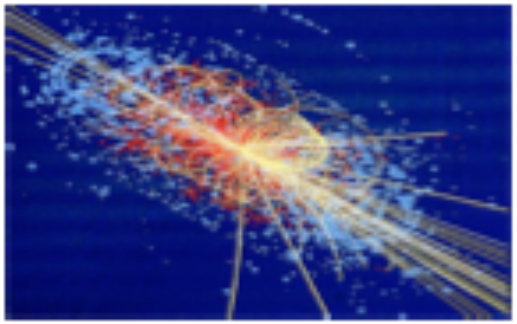

Playback mode

- Data is pre-loaded into the Read Out System
- DAQ/HLT plays back the data through the whole system except Level 1
- Allows testing the system in realistic environment

Last tech.run in April
Currently “24h” periods

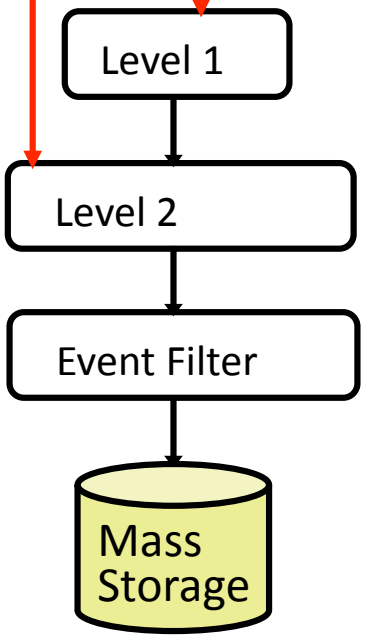

Mx or Sub-detector runs

Cosmic Real Events

Real mode

- Real data comes from the detectors
- Allows testing the whole system including Level 1
- Tests the software with real “imperfect” data
- Limited to cosmic-ray events

M7 cosmics run in May
M8 about to start

April technical run

- 10^{31} trigger menu on L1-accepted minimum bias sample:
  - **33 ms @ L2** (40 ms nominal)
  - **142 ms @ EF** (1 s nominal)
  - L2 has little margin – may be necessary to re-allocate EF CPUs to L2
- **High-rate** measurements:
  - 4 L2 racks + event builder could handle 30kHz L1 rate (minimum bias events)
- Move **prescale before each level**: a L2 time improvement of ~20% for the 10^{31} trigger menu and ~100% for the 10^{32} trigger

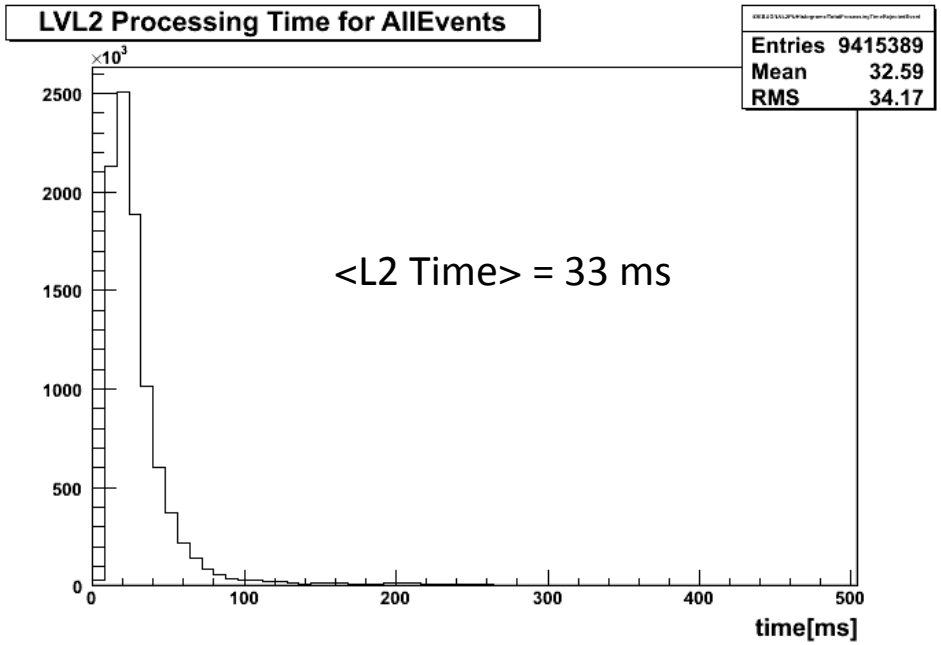

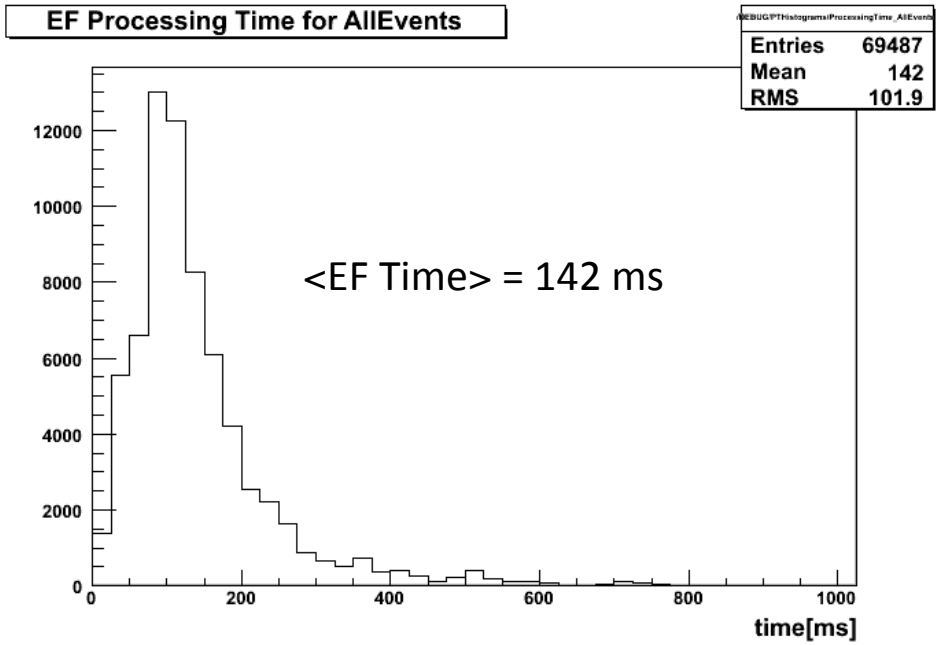

Highlights of M7 cosmics run

- The new TriggerTool for the trigger DB is being used:
- Streaming tested

Stream Name	Events	GB
physics ptdummy	19083902	18453.947
physics HOT_Cosmic_L1Calo	525	5.651
physics HLT_Cosmics_RNDM0	3187952	708.867
physics HLT_Cosmics_NIM4	104823	530.711
physics HLT_Cosmics_NIM0	217398	1610.096
physics HLT_Cosmics_MU3	10323004	29838.986
physics HLT_Cosmics_MU2	3047392	10160.149
physics HLT_Cosmics_MU1	13978720	6595.953
physics HLT_Cosmics_J5	1004	10.808
physics HLT_Cosmic_MBTS	421503	3212.528
physics HLT_Cosmic_L1Calo	1481424	7895.872
physics HLT_Cosmic_EM_L1Calo	354900	2606.481
physics HLT_Cosmic_AllTePixel	2173	10.778
physics HLT_Cosmic_AllTeIDSelected	8547	44.401
debug hlterror	683035	3114.299
debug efd	1339	1.462
debug L2ForcedAccept	13933	46.067
calibration HLT_Cosmics_RNDM0	111233	1036.929
calibration HLT_Cosmics_NIM29	461	0.143

Conclusions

- The migration of the trigger EDM to flat containers was quite smooth and will save a lot of work later
- Configuration machinery is now standard for trigger and can start to be used offline
- The trigger code has been continuously improved and tested and is basically on track for the LHC startup

To this!....

Backup slides

- Three trigger levels:
- Level 1:
  - Hardware based
  - Calorimeter and muons only
  - Latency 2.5 μs
  - Output rate ~75 kHz

- Level 2: ~500 farm nodes(*)
  - Only detector **"Regions of Interest" (RoI) processed - Seeded by level 1**
  - Fast reconstruction
  - Average execution time ~40 ms(*)
  - Output rate up to ~2 kHz

- Event Builder: ~100 farm nodes(*)

- Event Filter (EF):~1600 farm nodes(*)
  - Seeded by level 2
  - Potential full event access
  - Offline algorithms
  - Average execution time ~4 s(*)
  - Output rate up to ~200 Hz

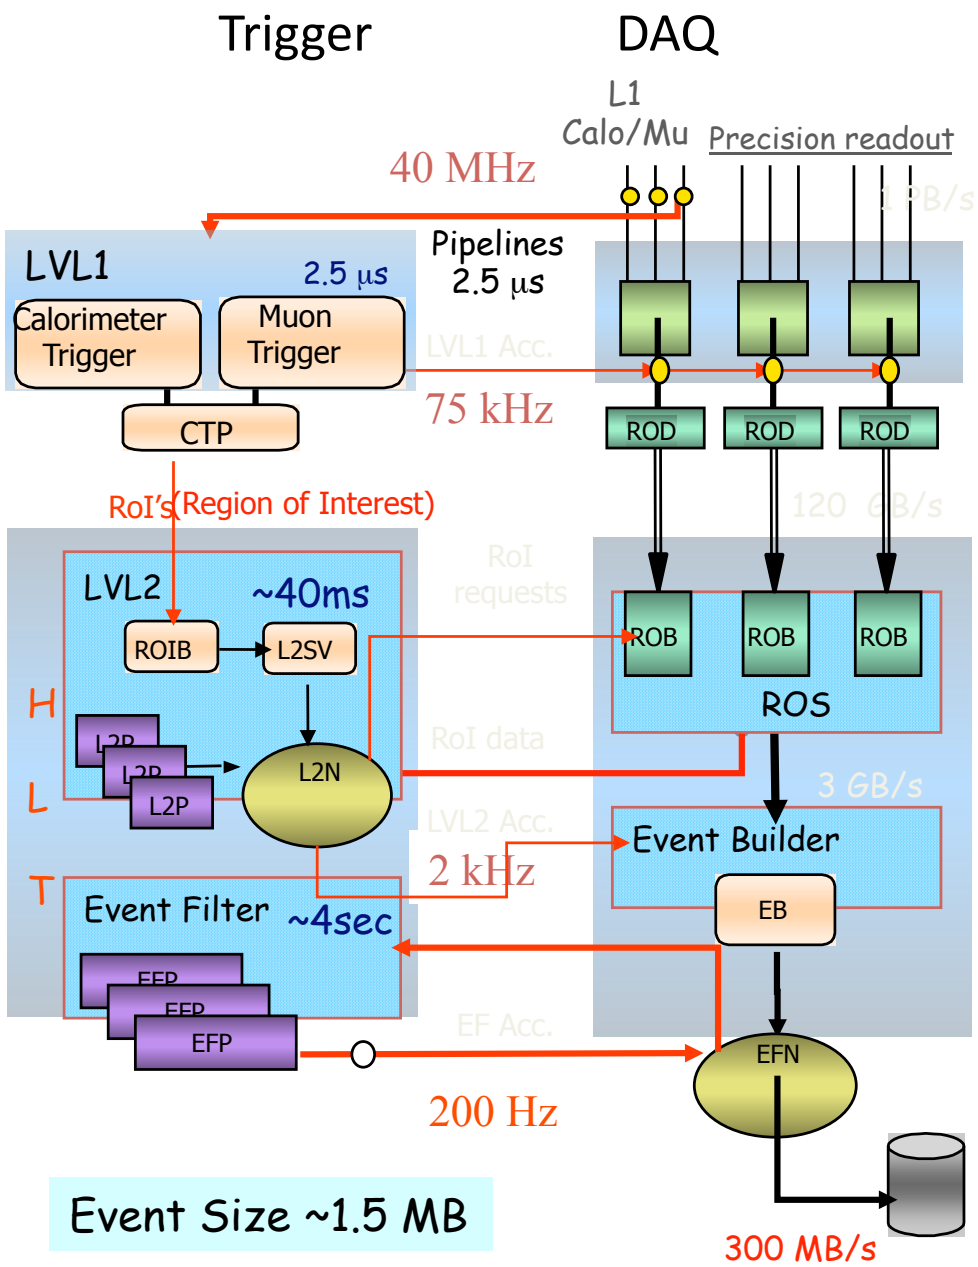

(*) 8CPU/node (four dual-CPU cores at ~2GHz)

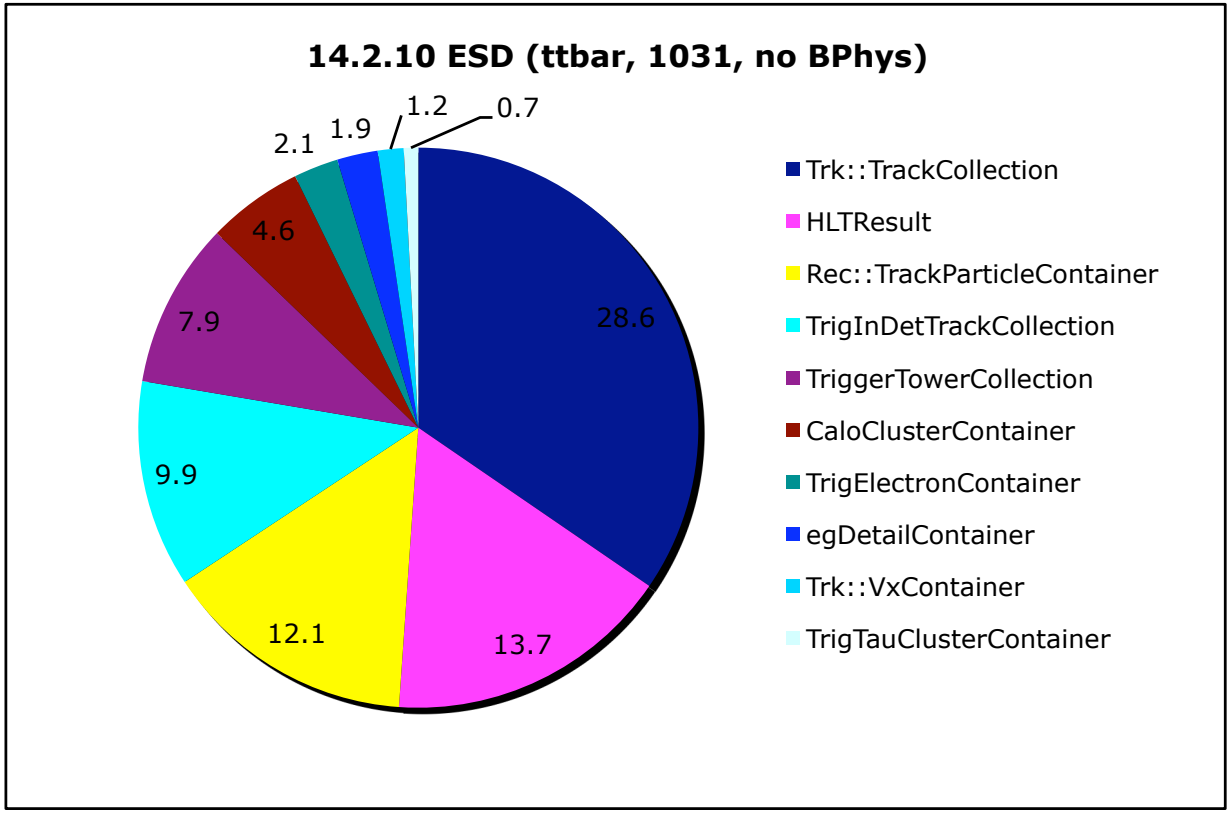
14.2.10 ESD (ttbar, 1031, no BPhys)
28.6
13.7
12.1
9.9
7.9
4.6
2.1
1.9
1.2
0.7
Trk::TrackCollection
HLTResult
Rec::TrackParticleContainer
TrigInDetTrackCollection
TriggerTowerCollection
CaloClusterContainer
TrigElectronContainer
egDetailContainer
Trk::VxContainer
TrigTauClusterContainer

Configuration Database

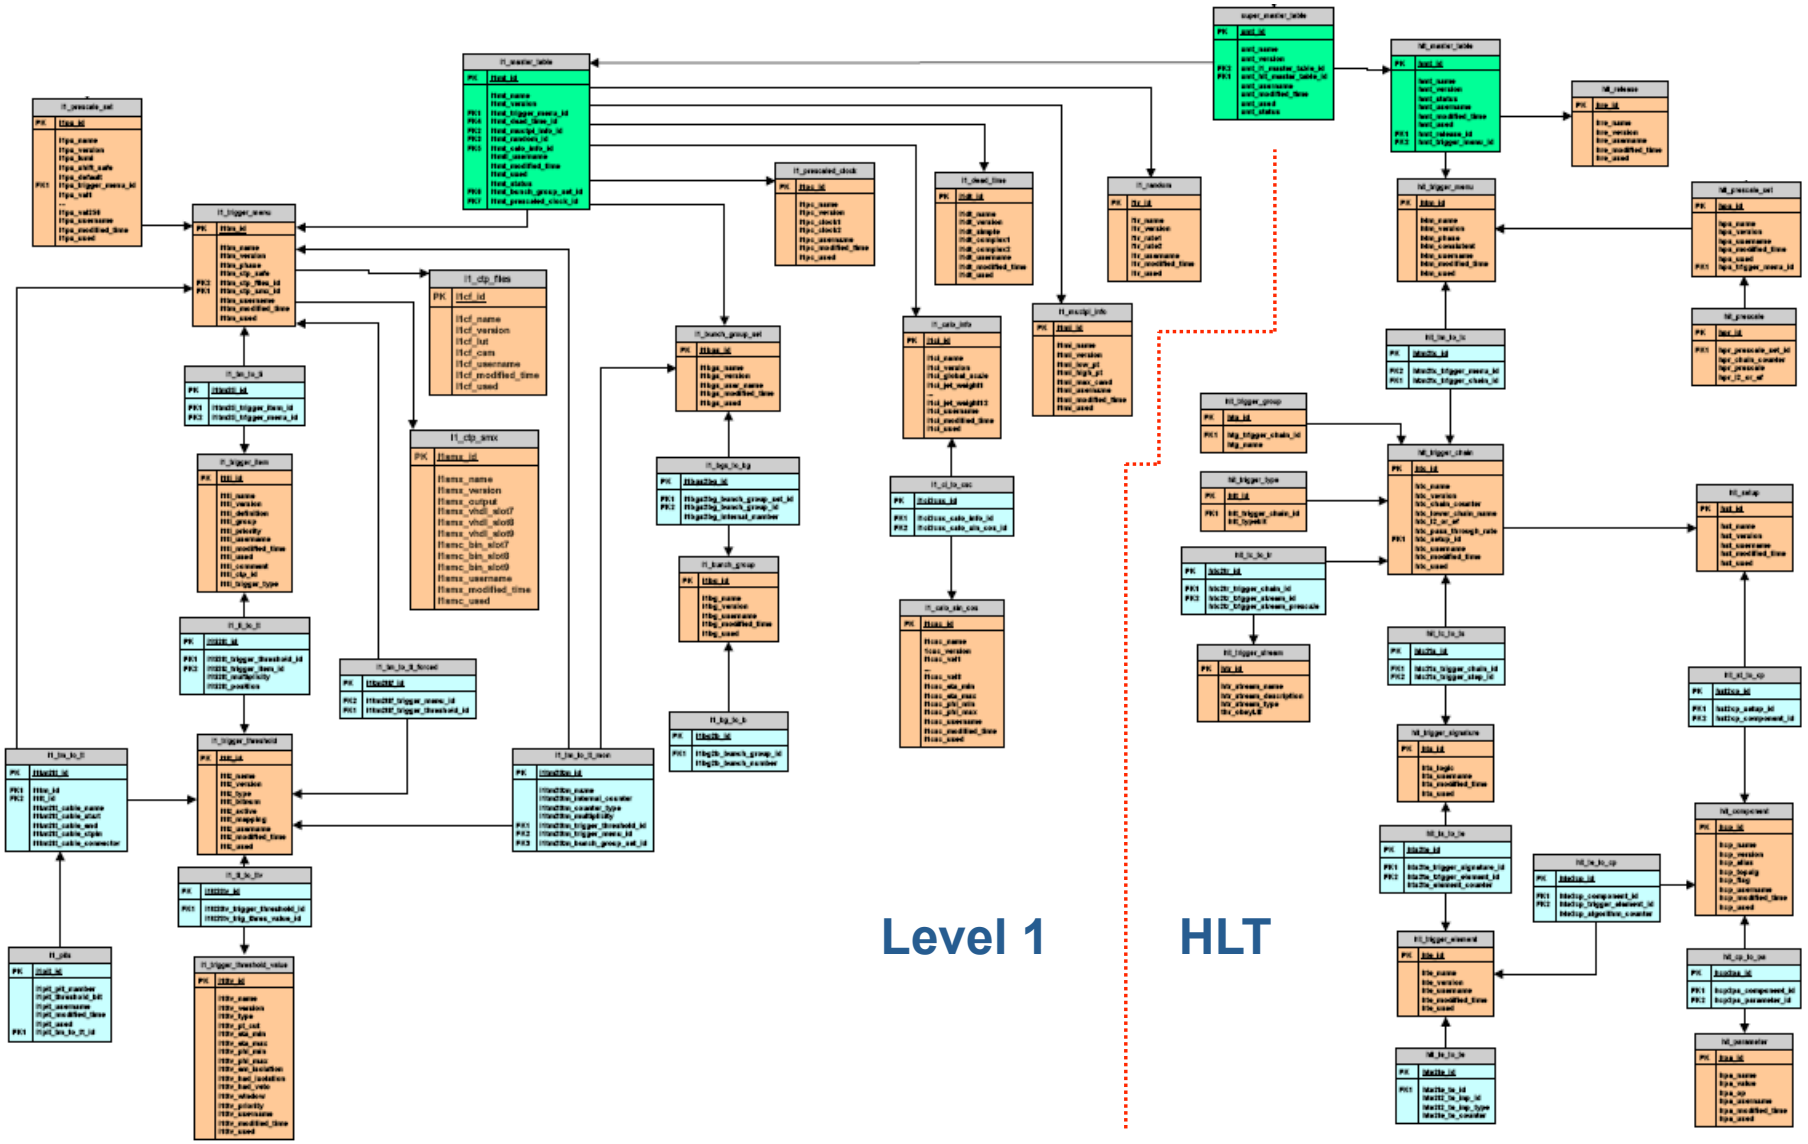

Modifying a menu with TriggerTool

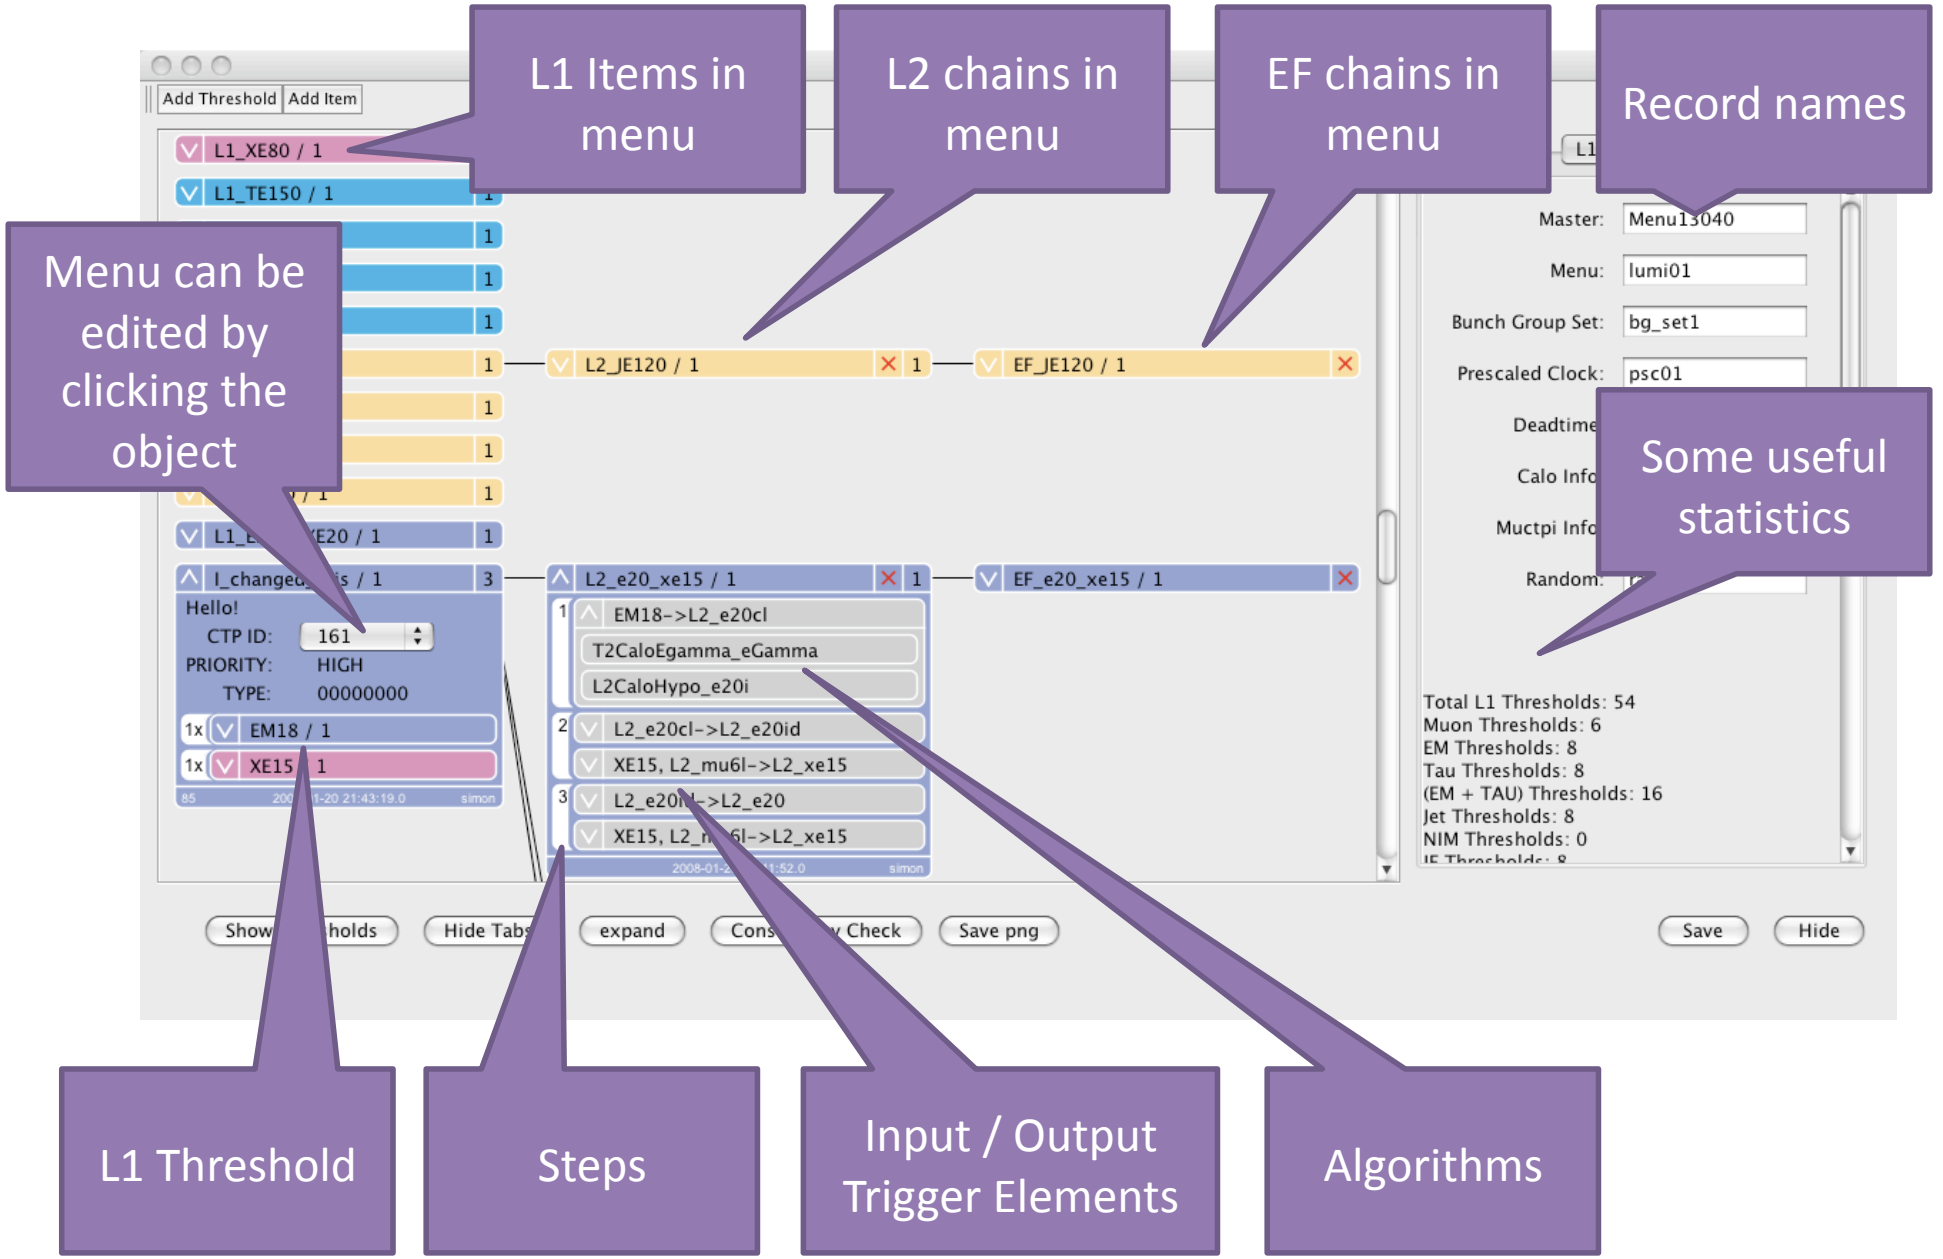

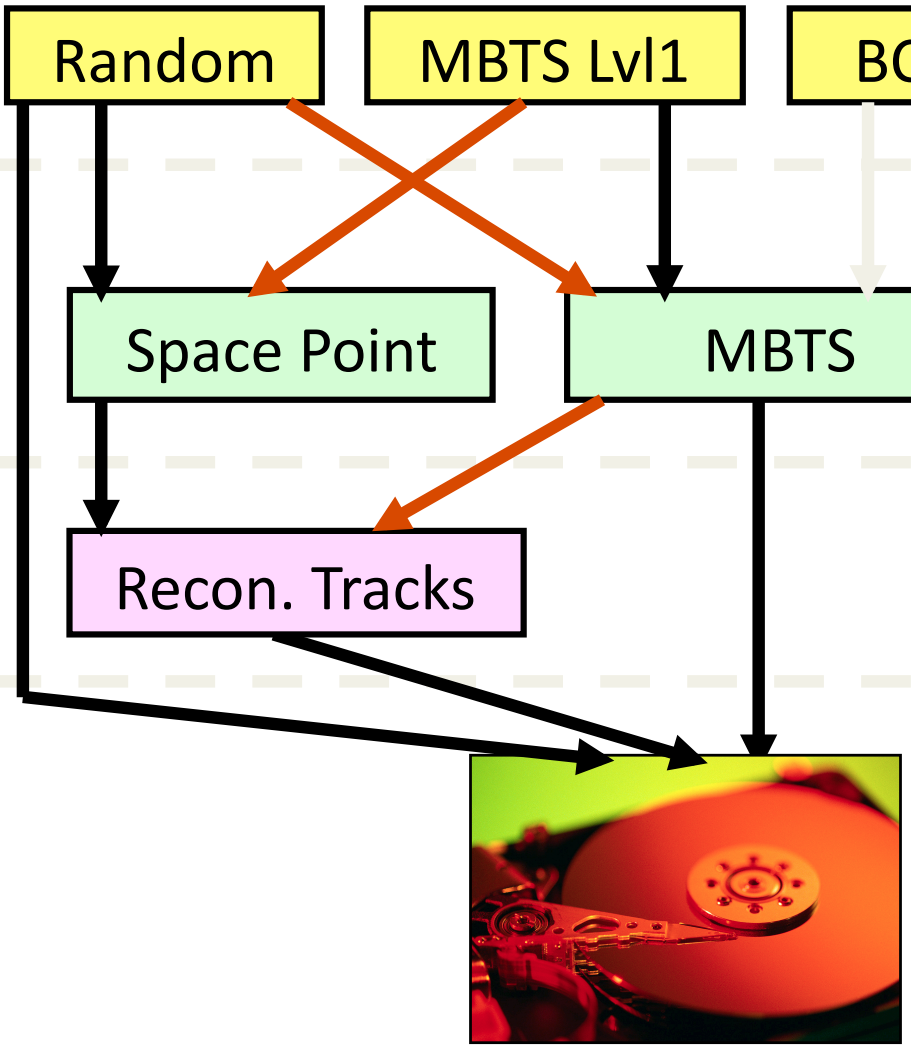
Random
MBTS Lvl1
BCM Lvl1
LUCID Lvl1
Space Point
MBTS
Recon. Tracks

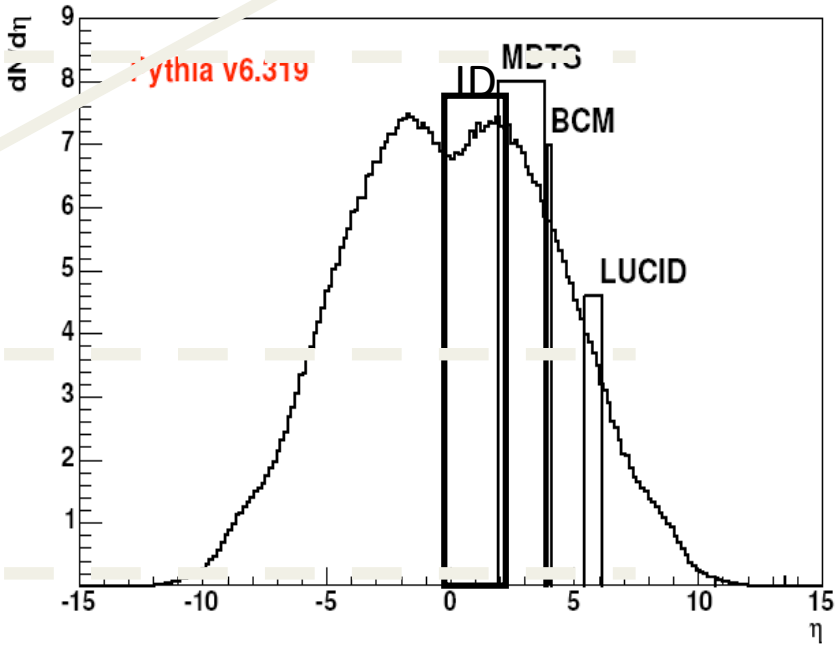
Pythia v6.319
ID
MBTS
BCM
LUCID
η

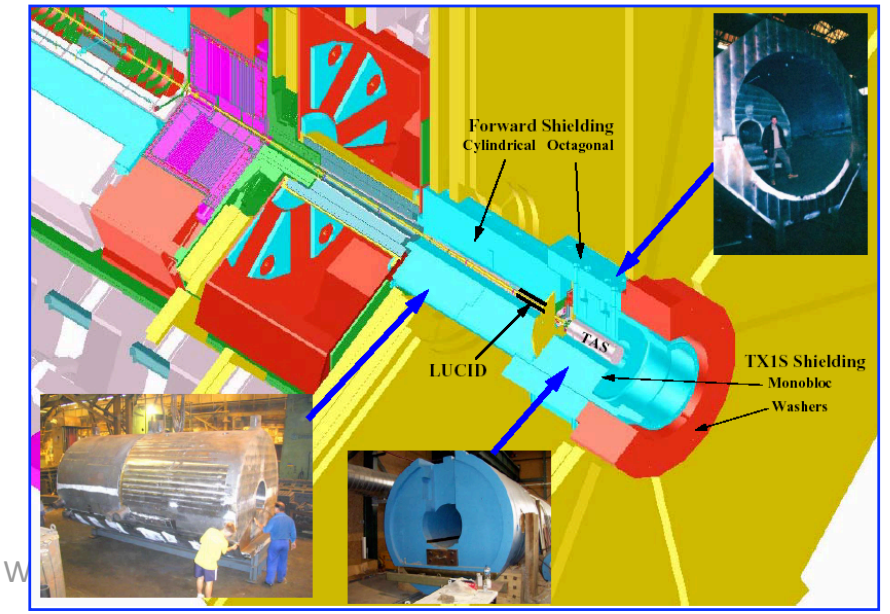
Forward Shielding
Cylindrical Octagonal
LUCID
TAS
TX1S Shielding
Monobloc
Washers

April Technical Run (cont.)

- Trigger pre-scaling was implemented to be done after the chains have been processed:
  - It has advantages but isn't it a waste of time?
- L2 time performance results of normal vs early pre-scaling tests:

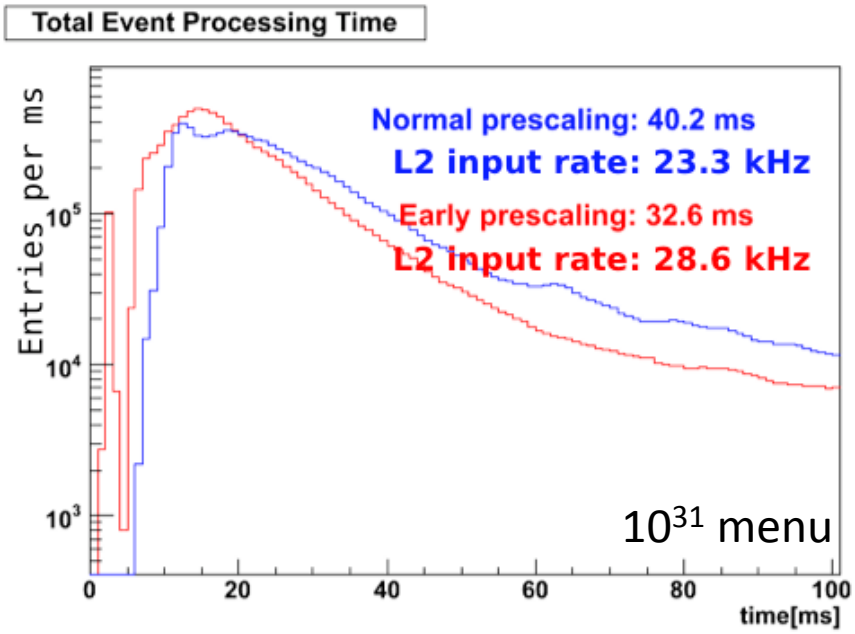

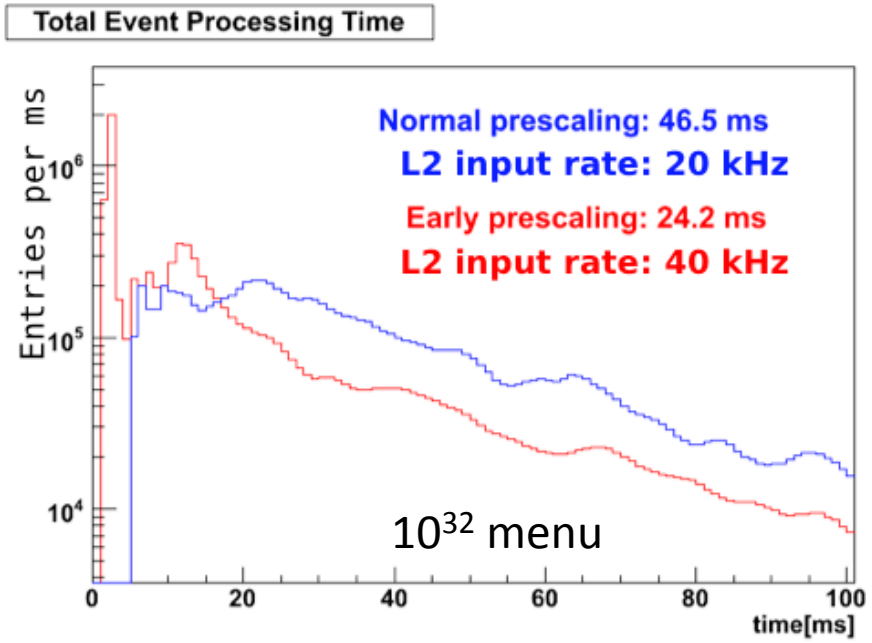

The test showed a L2 time improvement of ~20% for the 10^{31} trigger menu and ~100% for the 10^{32} trigger menu (faked by events not processed)

Einstein image from:

http://www.hetemeel.com/einsteinform.php