Trigger validation

Ricardo Goncalo, RHUL
Physics Validation - 14 September 2007

Memory leaks in 13.0.30 nightlies

- The plot thickens…
- Still work in progress, unfortunately
- Leaks running the trigger on 12.0.3 RDOs are OK
- But running on 12.0.6.5 RDOs, we have ~50MB jumps in some events
  - See plots from Frank Winklmeier next

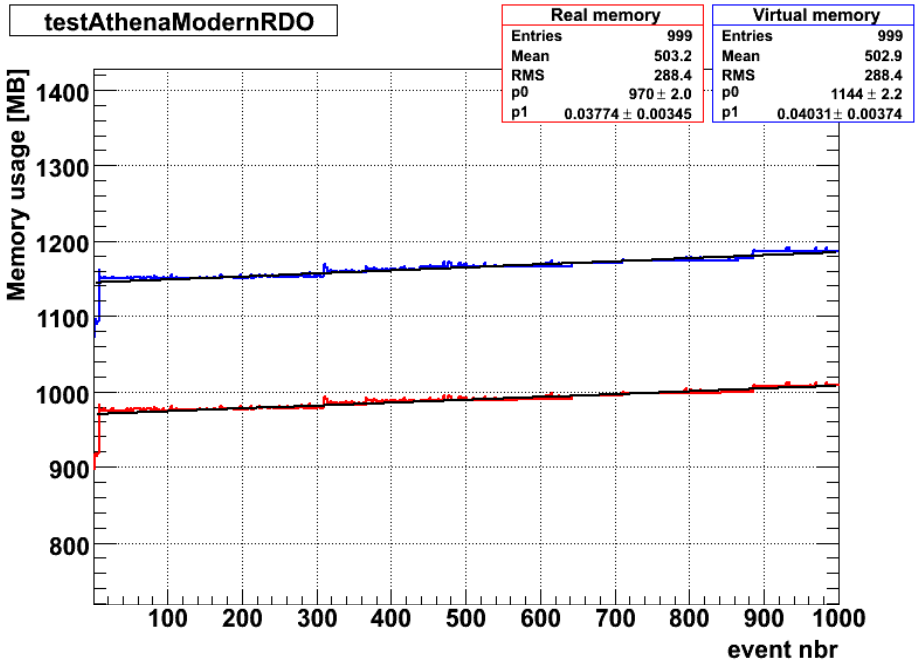

Test	Initial virt.mem.	Leak/event
All slices, no output	~1.1 GB	~20-40 kB
All slices, BS output	~1.0 GB	~100-250 kB
Muon slice	~950 MB	~30 kB
Jet slice	~850 MB	~10 kB

Memory leaks in 13.0.20/30

- Same behavior seen in bugfix, rel_4 as in 13.0.20.2
  - Wouter's tar ball run in rel_4, bugfix:

```
CENTAA_Audit    WARNING  Memory leak!  circa 3.38893 (10 first events
excluded) or 2.25362 (only second half of the job) MBytes/event
```

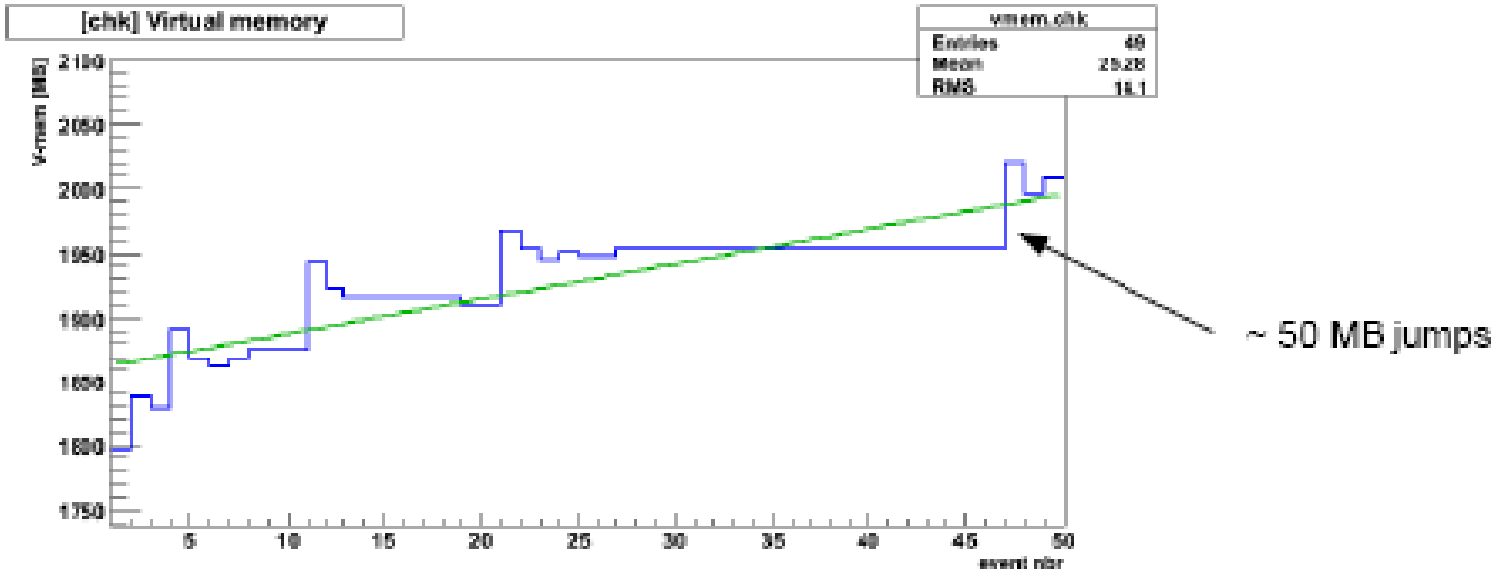

Is it a trigger problem?

- Difference in total vmem before and after EF steering execute:

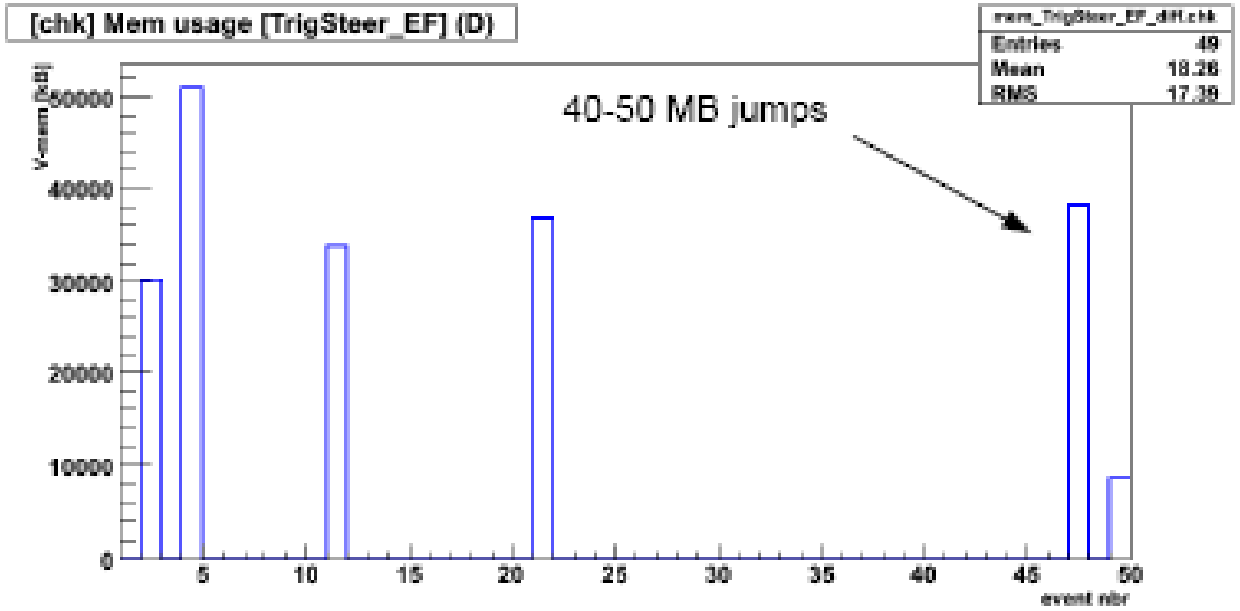

- BUT:
  - EF Steering only runs EF algorithms. We don't find any peaks like that in any of the EF algorithm.

Trigger tests in bugfix, rel_3

- Standard trigger tests in bugfix, rel_3 don't show this problem.
  - 12.0.3 RDO input, No POOL output

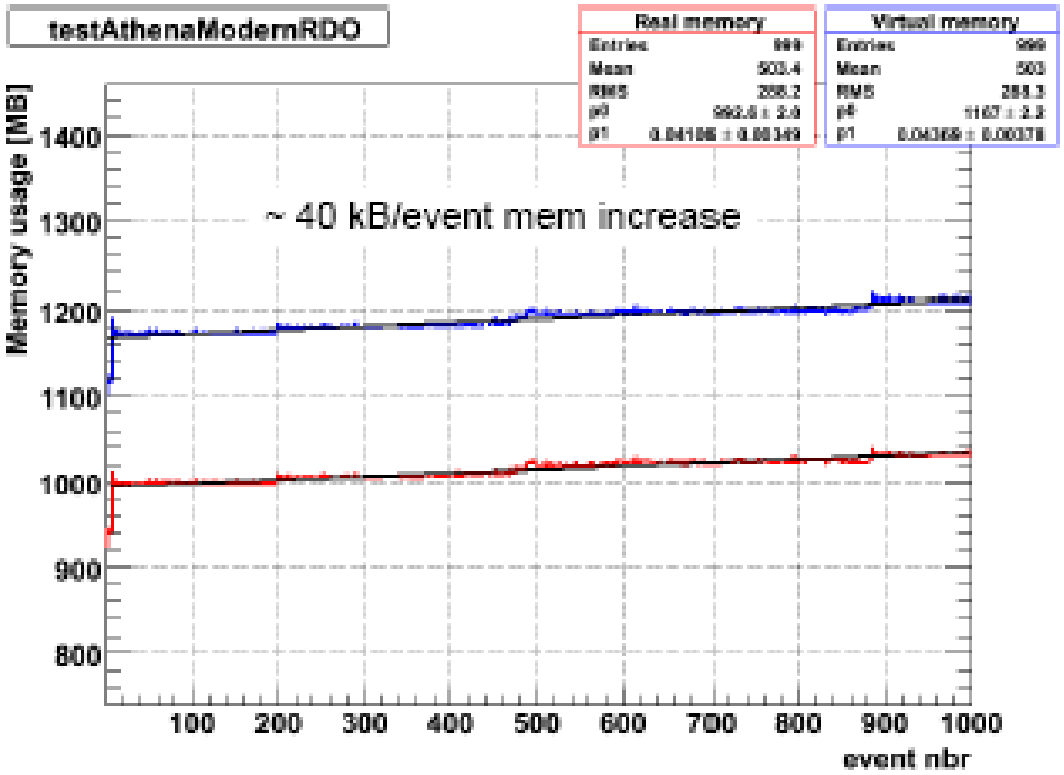

No large peaks
Moderate mem increase

Running trigger on Wouter's file

- Same test as on previous slide but on 12.0.6 RDO

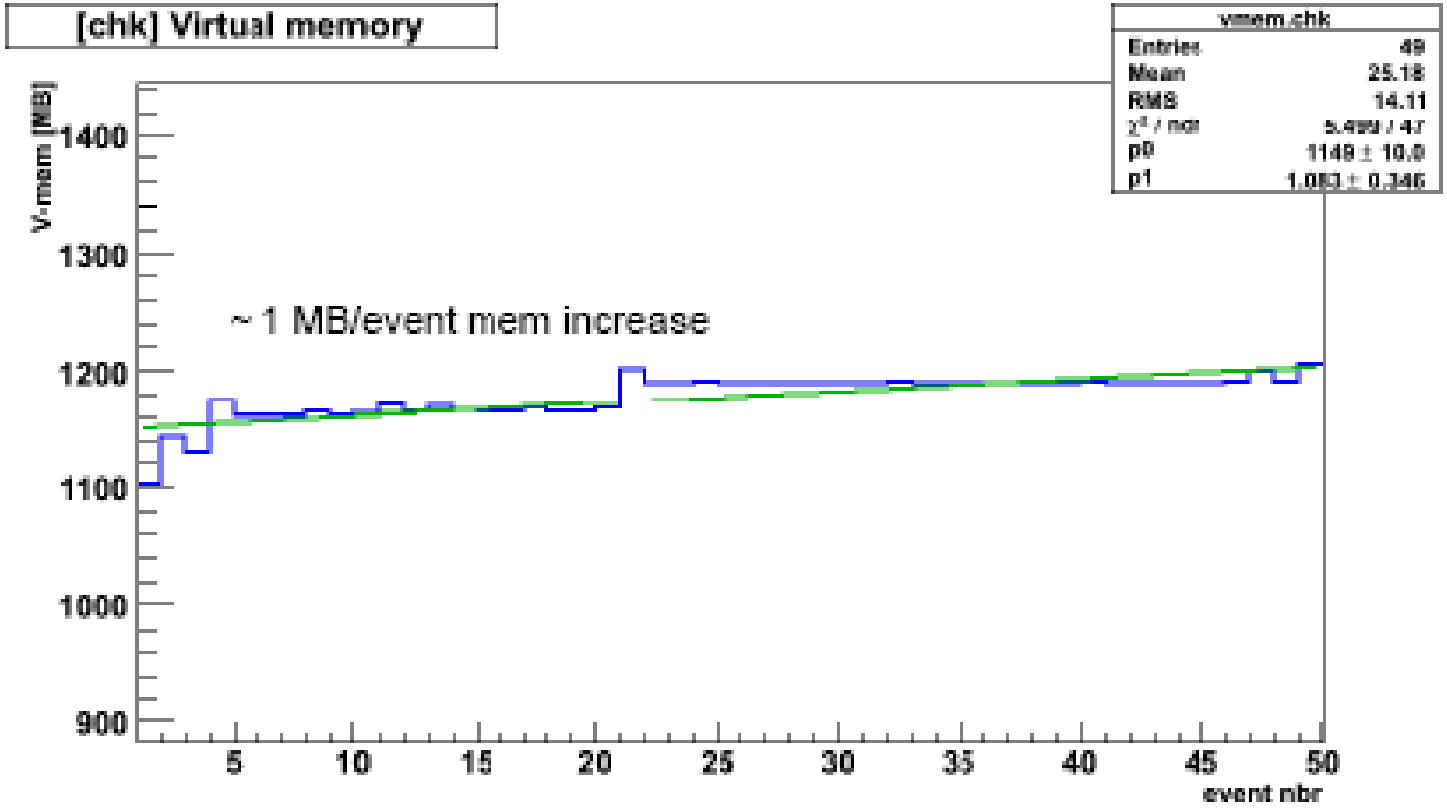

- Seems specific to 12.0.6 RDOs.

Memory leaks in 13.0.20

- 13.0.20.x to be used for detector paper performance section due to delays in 13.0.30

- But crashes due to mem.leaks at >200 events
  - Best-case scenario is that 1 leak corresponds to most of the total leaks
  - One large leak may have been identified so far, but still looking

Running in AthenaMT:

•Events preloading into memory causes jump

•Large difference between 13.0.20 and 13.0.30

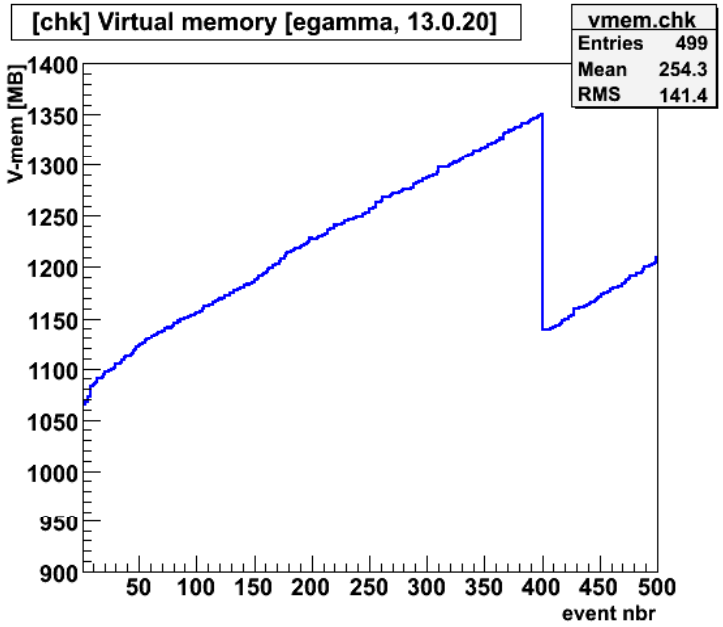

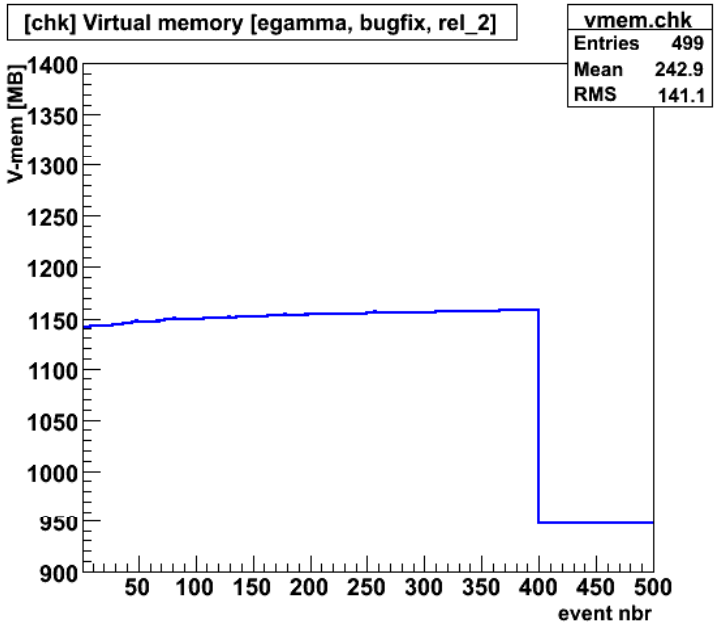

Mem.leak summary

It’s a mistery!... still working on it…

Checks on slides by Joe Foster from previous meeting (11/9/07): Top Physics Validation: 13.0.20.2 vs 12.0.6.5: Trigger Info

- Use CBNTAA ntuples from these MCatNLO ttbar datasets:
  - valid2_misal1_mc12.005200.T1_McAtNlo_Jimmy.recon.NTUP.v13002002
    - 1600 events
    - EVGEN + SIMU = 12.0.6.5; RECO = 13.0.20.2
    - Event weights are still 0. *Should* get fixed in next release. (Thanks Borut!)
  - trig1_misal1_mc12.005200.T1_McAtNlo_Jimmy.recon.NTUP.v12000605
    - 8500 events
    - EVGEN + SIMU + RECO = 12.0.6.5
- Plot Physics parameters from L1, L2, 'Trig' info, ignoring weights.
  - No info about trigger variables at https://twiki.cern.ch/twiki/bin/view/Atlas/CBNTVariableList
  - Not always sure if differences are a problem or an improvement.
    - L1ET_JetEtSum, L1ET_EtMissHits may be problems.

L1 Trigger: EM, Jets

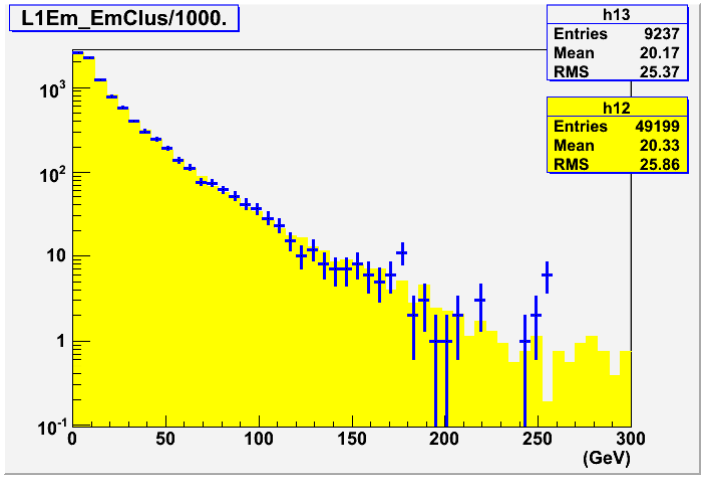

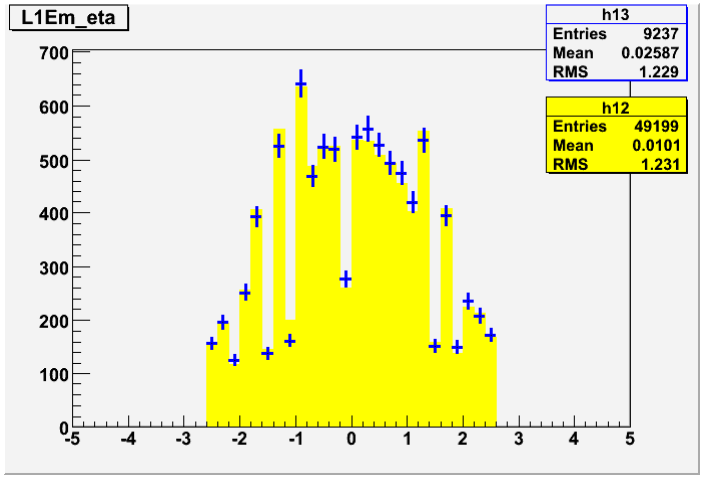

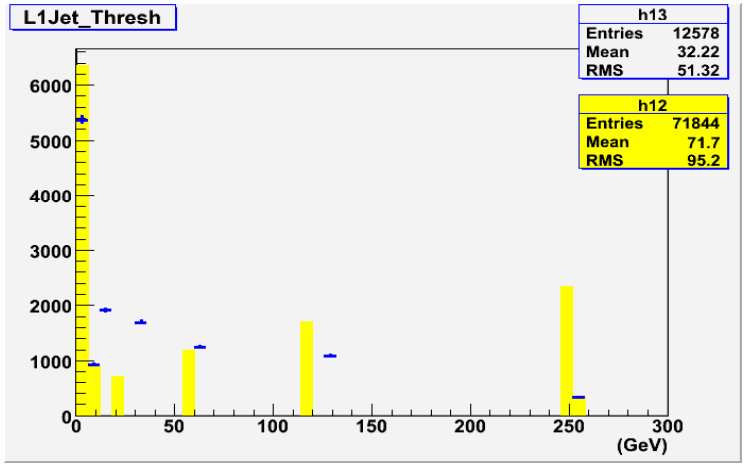

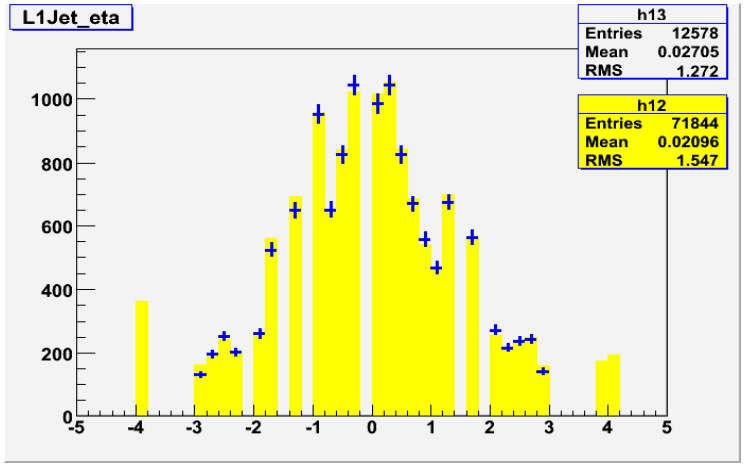

page 2: fine

bottom left: some L1 jet thresholds changed from 12.0.6.5 to 13.0.20

L1 ET, Missing ET

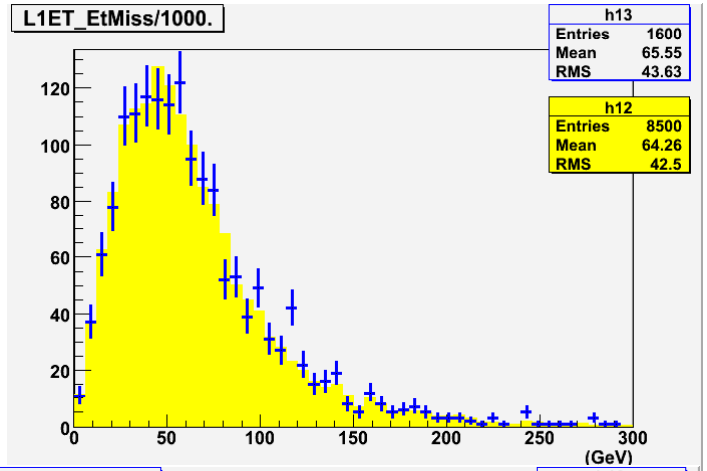

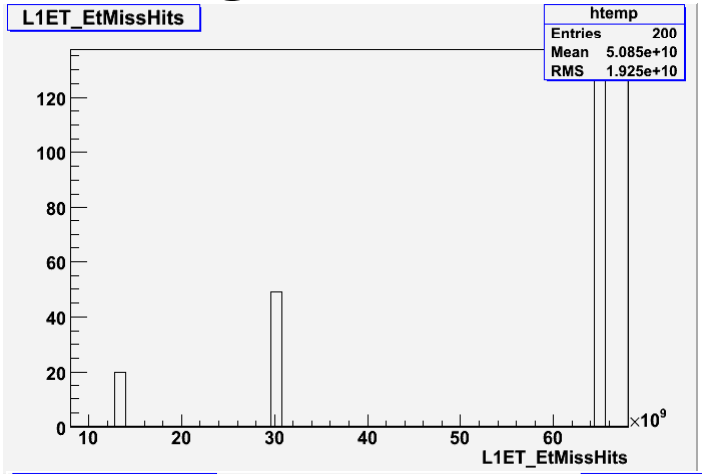

r13

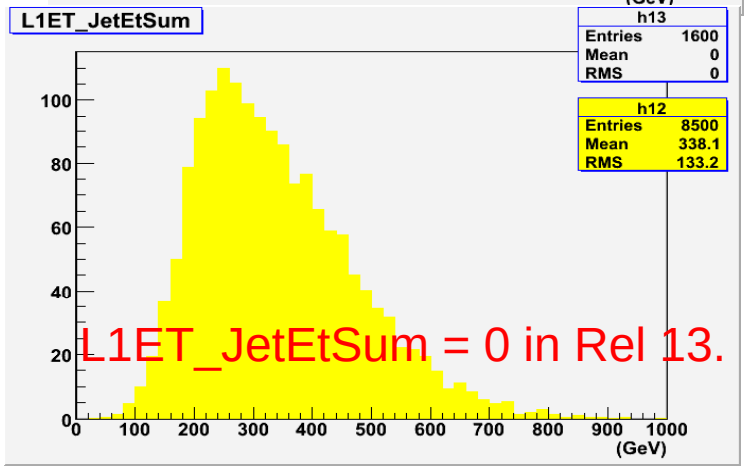

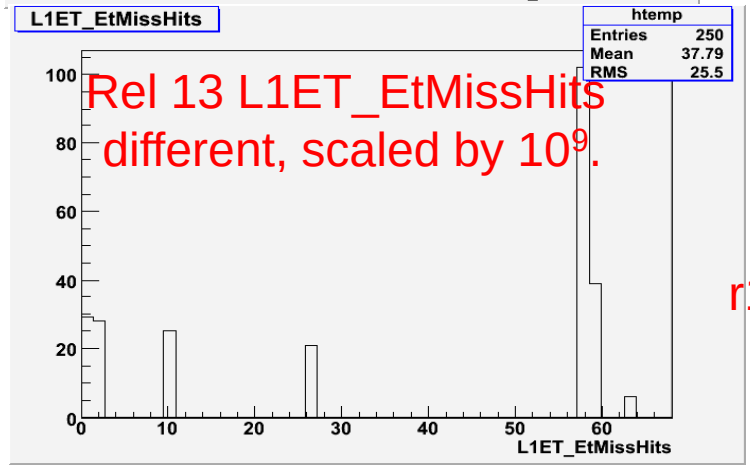

r12

page 3: all understood

top right: conversion error in CBNT/ROOT when going from int to 64-bit; tag to fix this, from Alan Watson, is in validation

left: L1 JetEtSum not defined in 13.0.20 menu; several thresholds exist in 13.0.30 menu

L2 EM, Hadron Calo EL2 Clusters

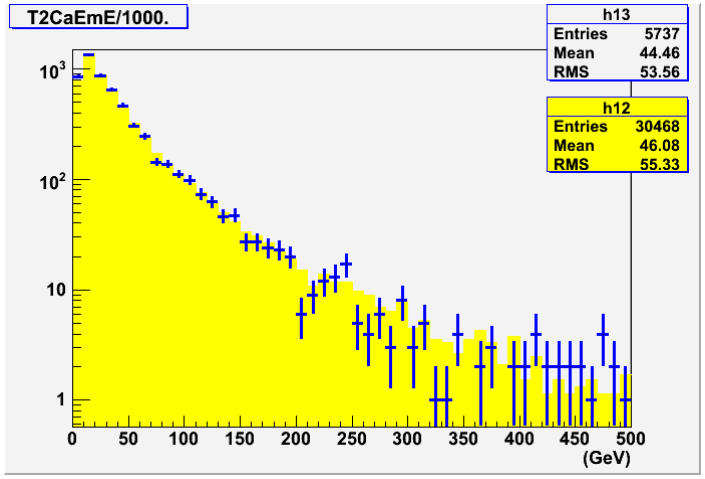

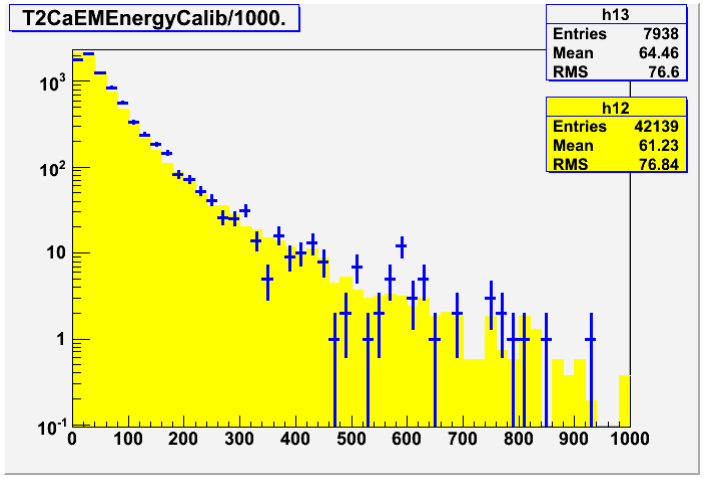

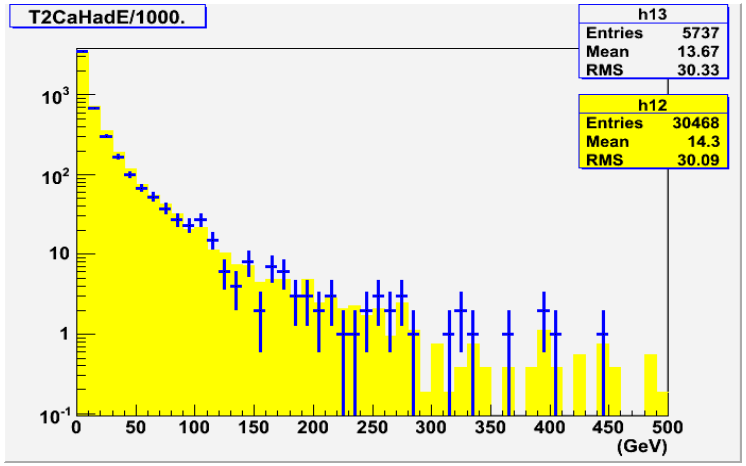

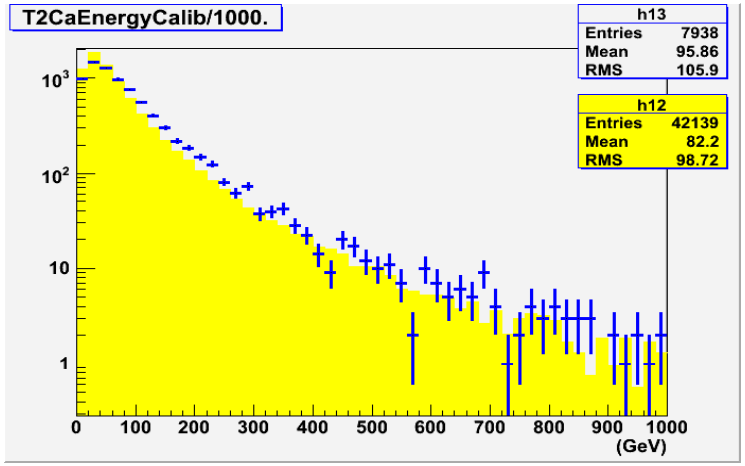

page 4: all fine?

let's say that the disagreement in T2CaEnergyCalib is not significant and may be due to several changes in reconstruction (anyone who knows better please shout)

L2 Jets

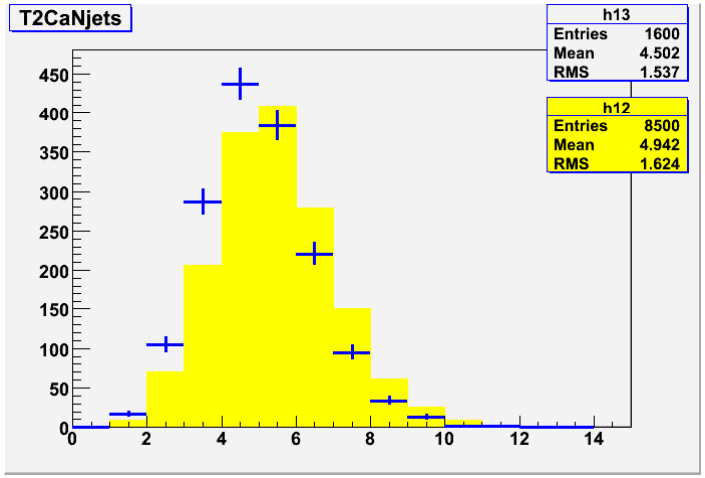

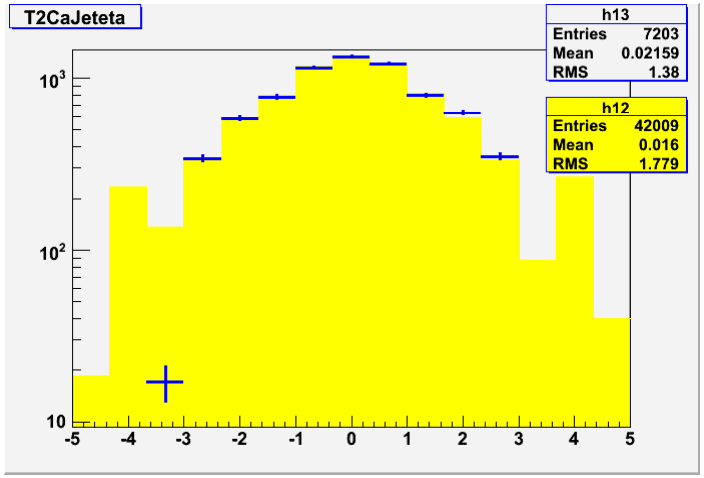

page 5: understood; Patricia Conde verified that the jet energy scale was correct to within 2% (next slide)

left: different number of jets at L2 due to changes in the L1 jet thresholds seen in page 2; the L1 thresholds were raised between 12.0.6.5 and 13.0.20, in general, which agrees with smaller nr. of L2 jets in 13.0.20

right: no forward jets were defined in 13.0.20 (they exist again in 13.0.30); therefore no events in 13.0.20 apart from 1 badly calculated value

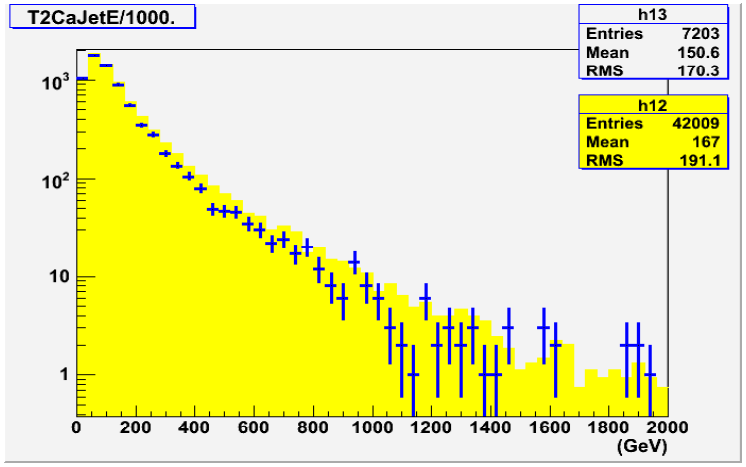

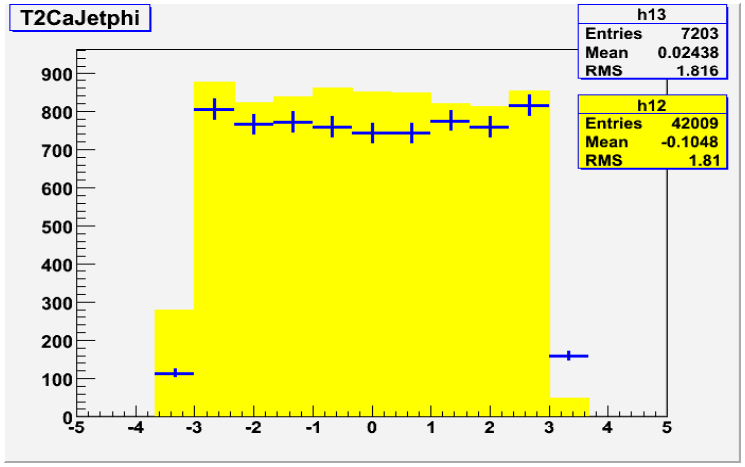

left: slightly smaller number of jets per event may explain small disagreement in histograms (I assume the histograms are normalised to the total nr. of events in the 12.0.6.5 sample)

right: also consistent with smaller number of jets per event; shape for small phi is different – may be statistically significant

More on L2 jets – from Patricia Conde

- Energy scale ruled out as a source of disagreement in previous pages
  - Shown here in different eta bins
  - Consistent with 12.0.6.5 E scale

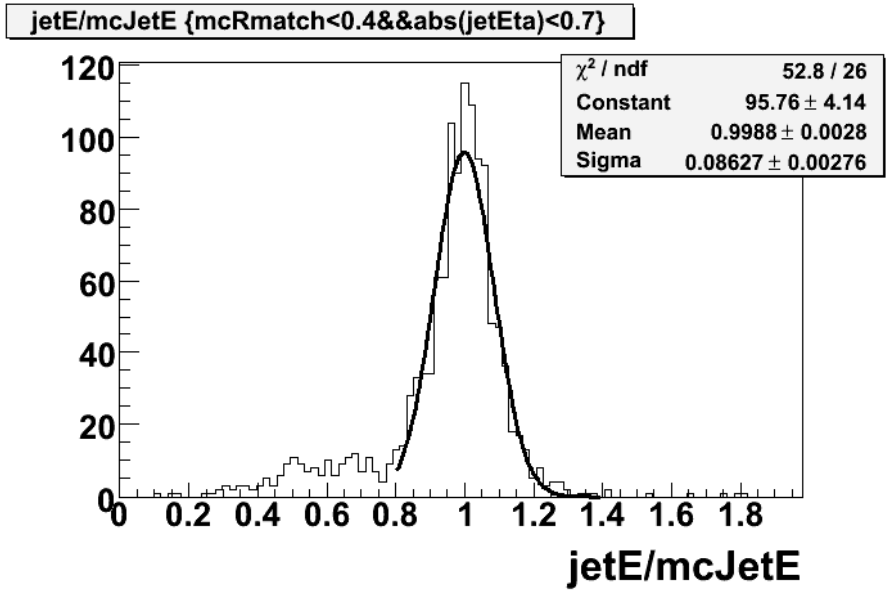

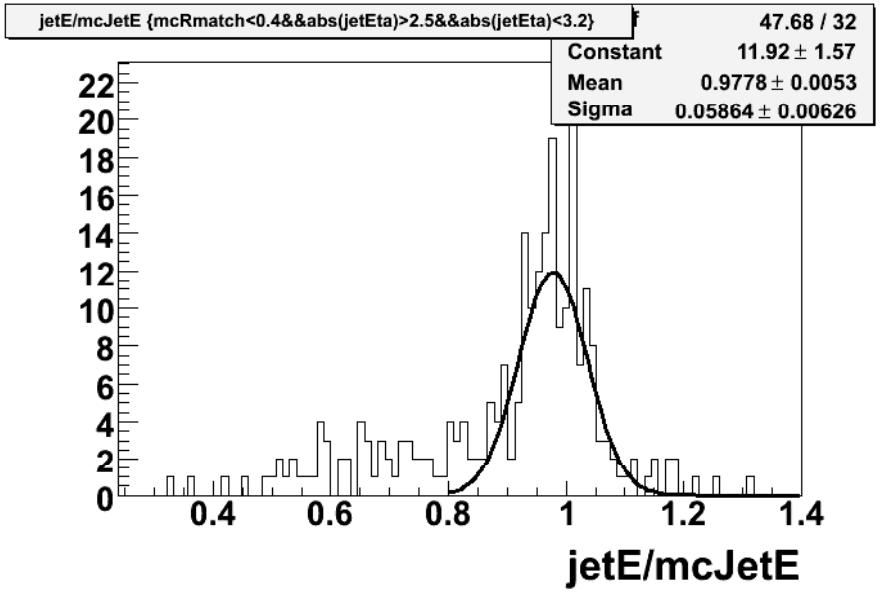

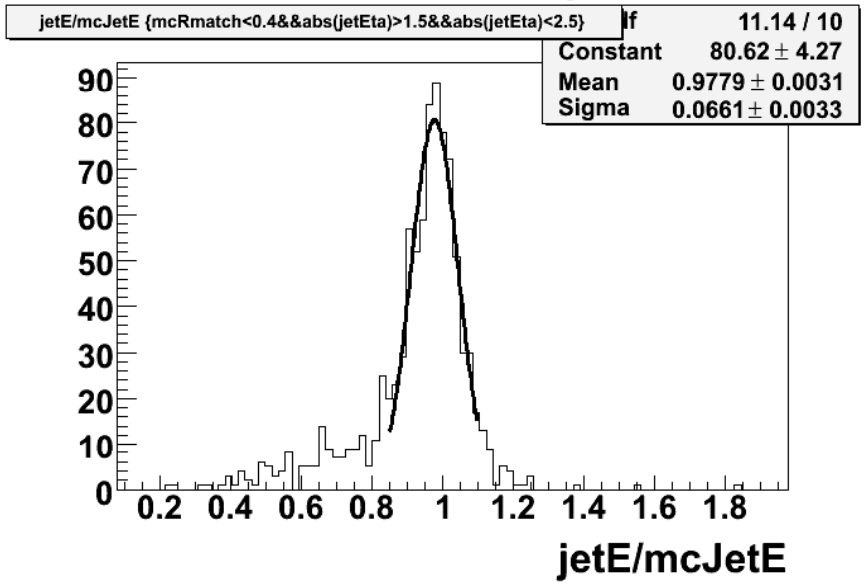

'Trig' Clusters

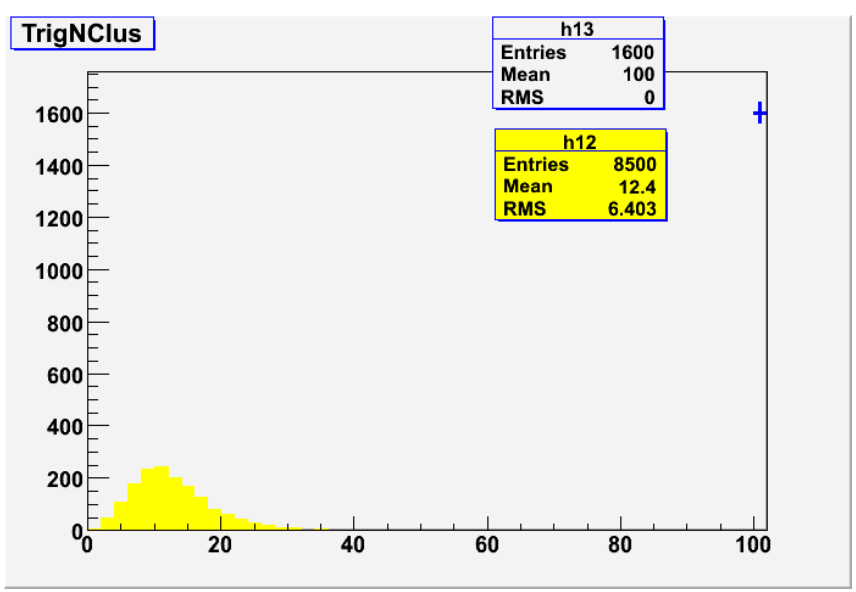

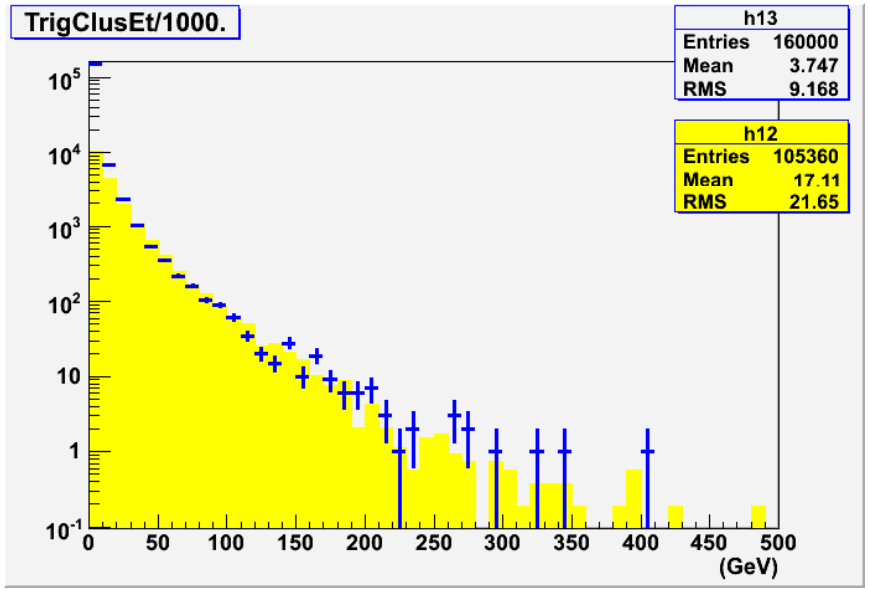

page 6: could not be reproduced so far in 12.0.; the suspicion is that clusters are being built for lower ET in 13.0.20

left: all (ttbar) have more than 100 clusters, which is the max number which can be stored in the CBNT block

right: good agreement for ET>20; disagreement for lower ET (?)

More on TrigClusters – from Joe Foster

- The low-ET are the ones populating the high-eta regions
- Here, low-ET means around (or less than) 2GeV, i.e. first bin of right-hand-side plot in previous slice

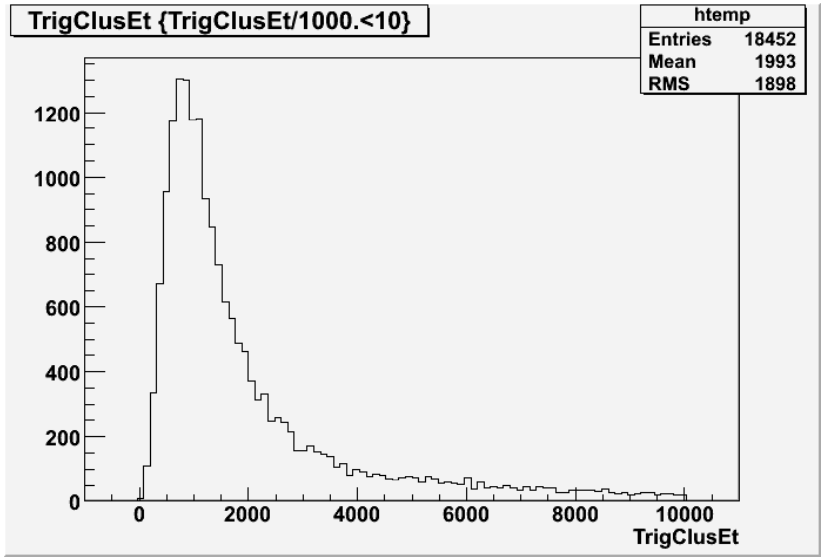

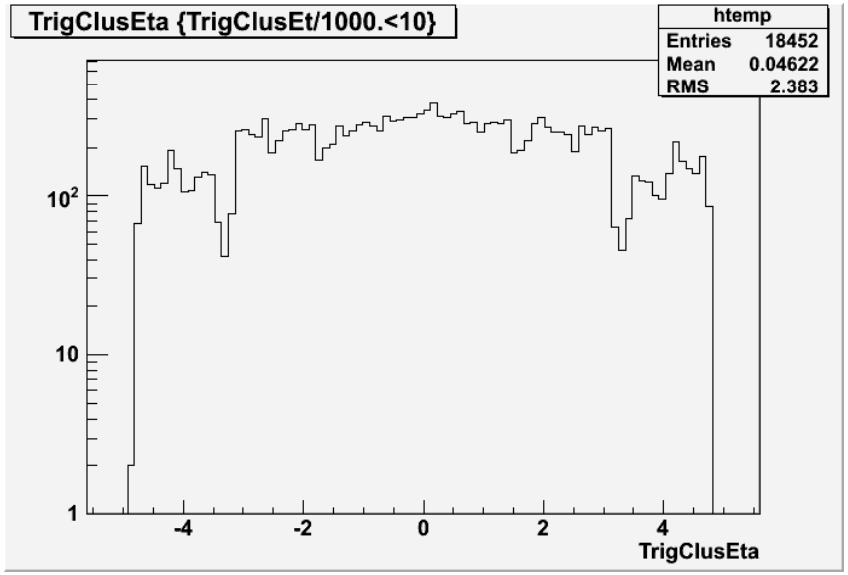

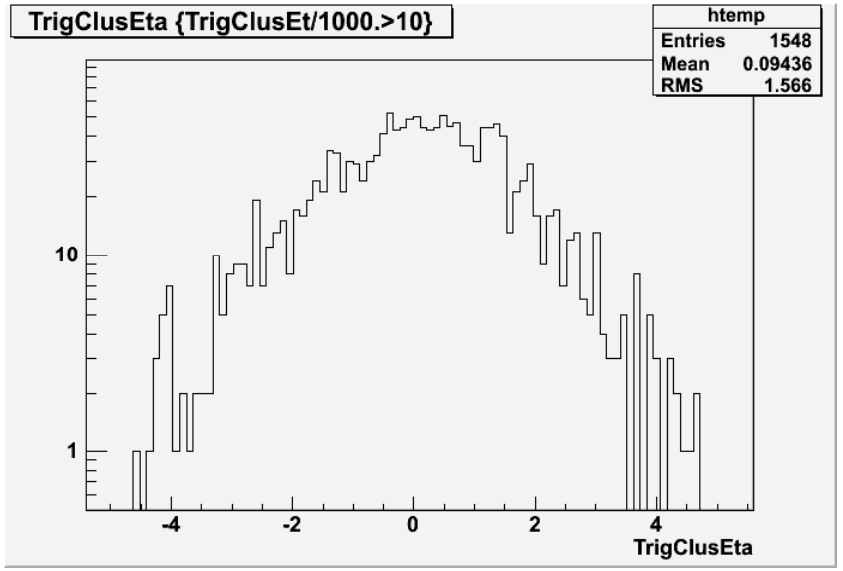

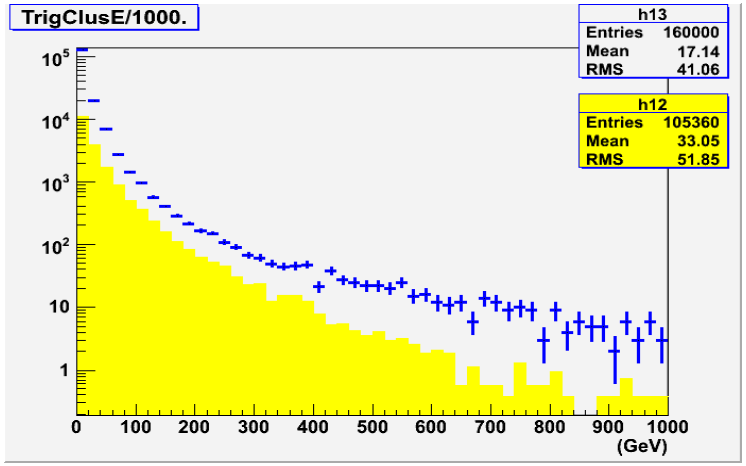

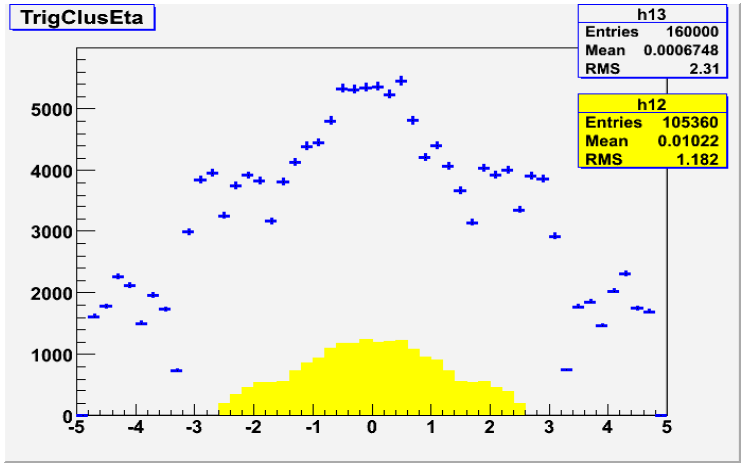

left: hard to interpret, since lower ET may correspond to very-forward high-E

right: most of the excess clusters will have low ET and fall in the first couple of bins of top right plot
Cibran Santamarina working on it using the data files used by Joe.

'Trig' Jets

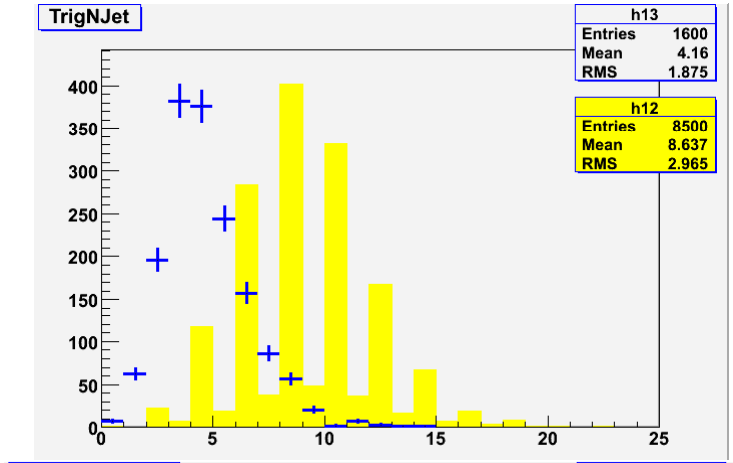

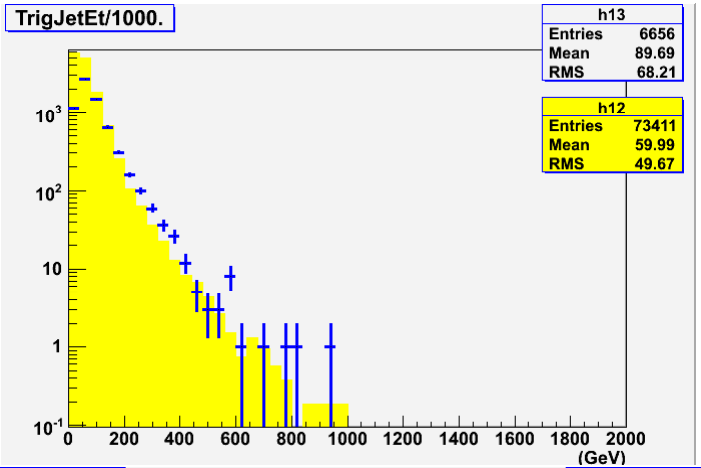

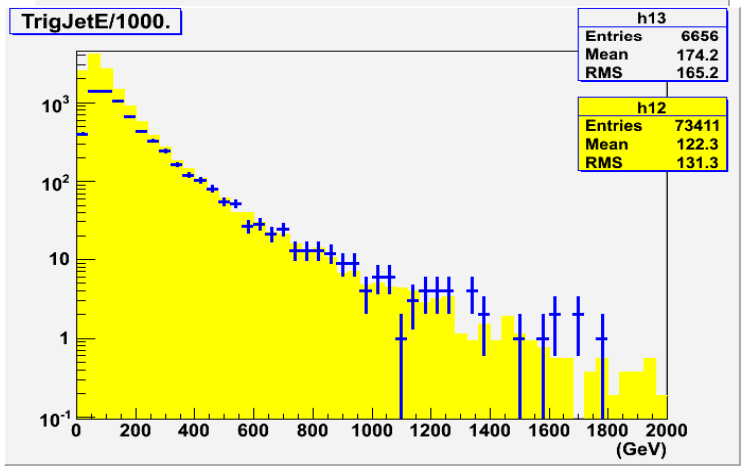

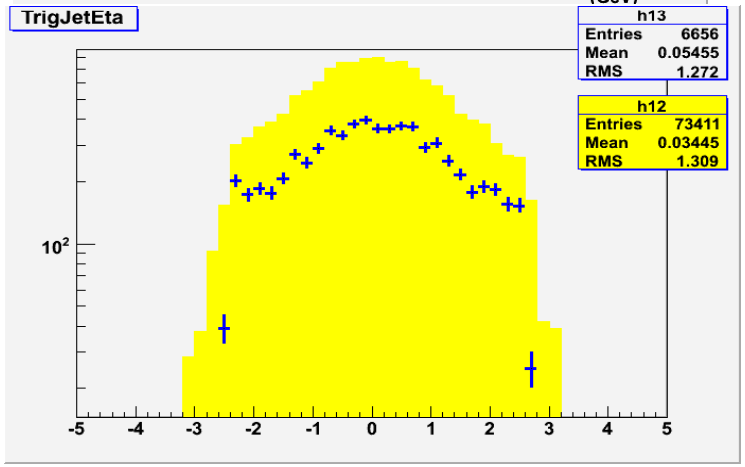

page 7: understood

•top left: different menu in 13.0.20 and 12.0.6; having corresponding cone and Kt jets was ~doubling the number if jets found in 12.0.6.5 and creating spikes

•top right and bottom left: these changes in the menu seem to affect mostly the low-ET region, where most jets are found

•bottom right: the same factor of ~2 appears here

Performance summary

- 12.0.6.5 was at the end of a development cycle

- 13.0.20 performance is poorer
  - Menus less complete than 12.0.6.5
  - …and less complete than 13.0.30
  - Not as user-friendly and not as well tested
  - Caught in the middle of at least one big transition in the trigger software