ATLAS FDR

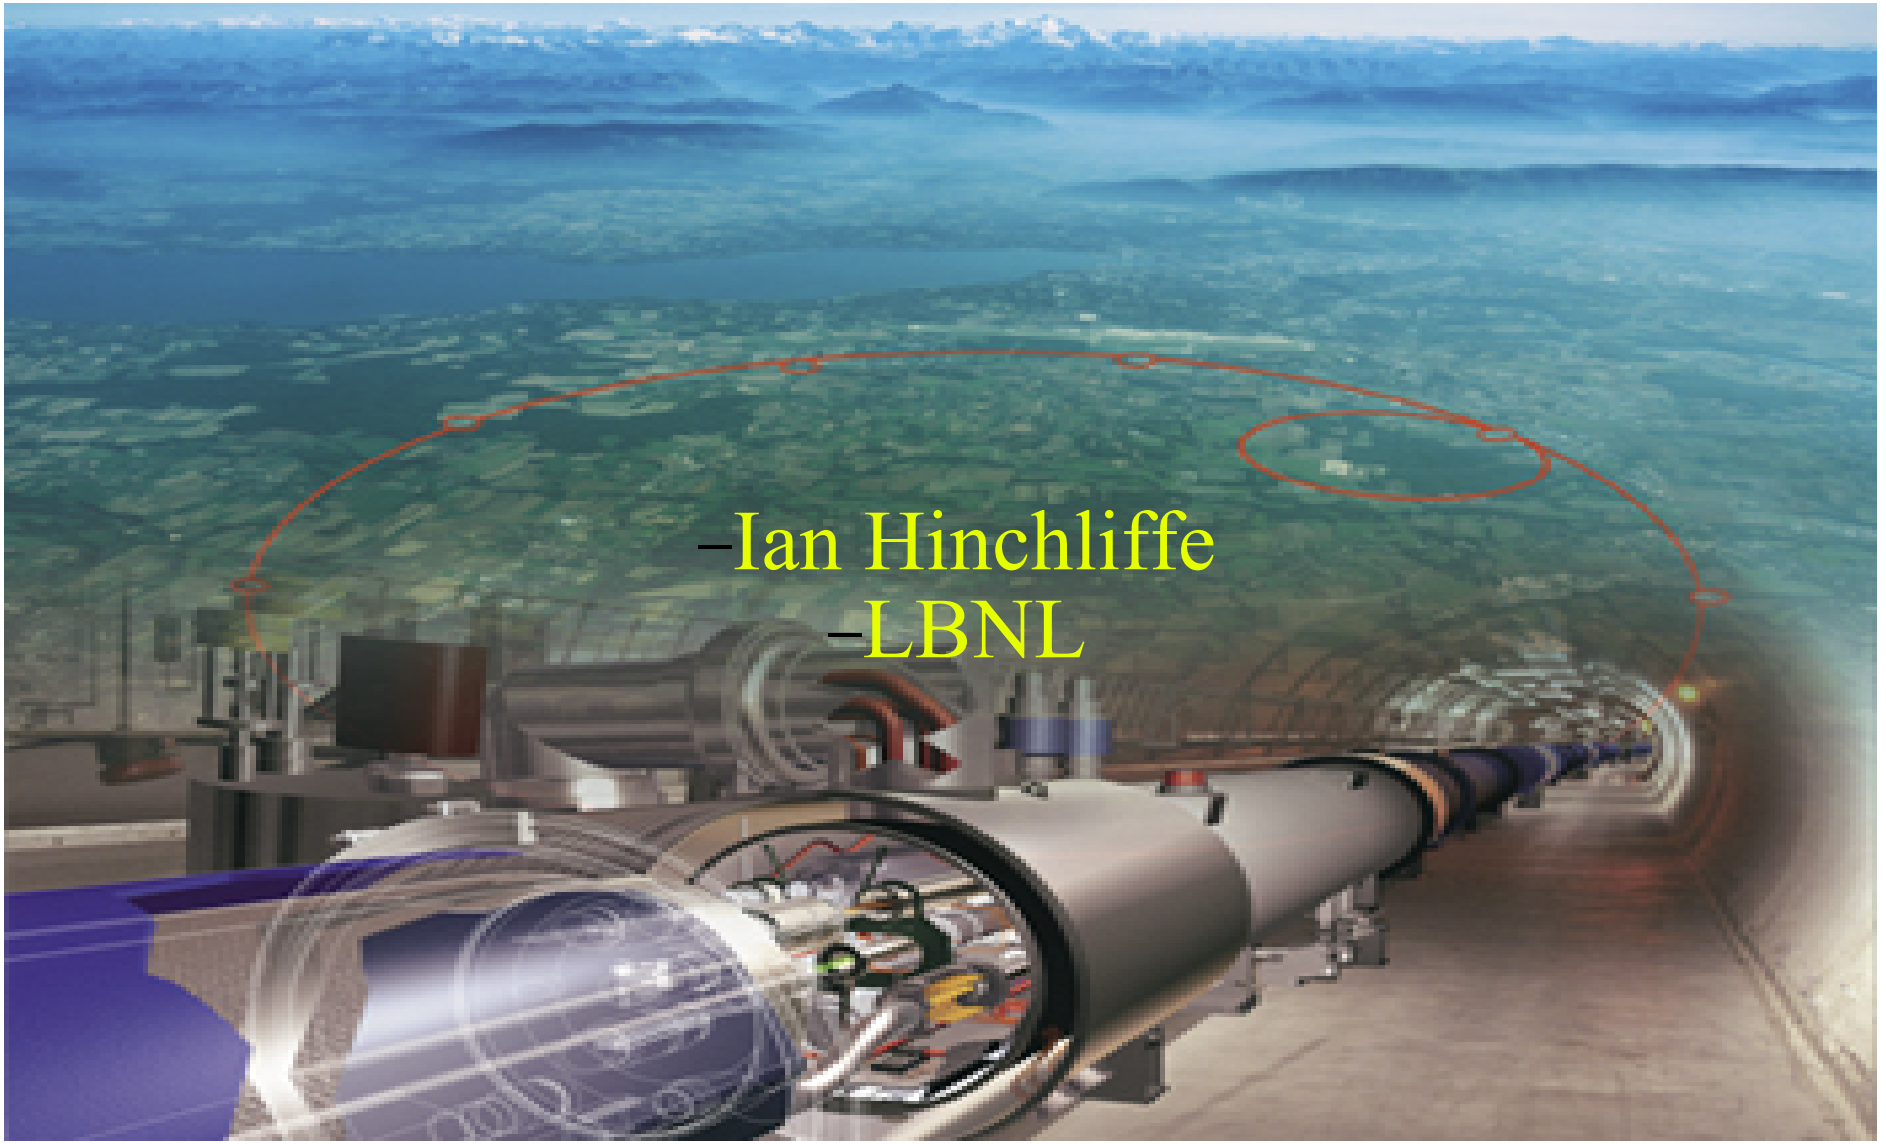

Outline

- Overall scope
  - FDR1: at 10^{31} and 10^{32} luminosity: 1 hr at latter, 8 at former
  - FDR2:at 10^{32} and 10^{33} luminosity: 4hrs at former: 1hr at latter
- Data "played through T0" several times
- Users issues
- Links
- http://indico.cern.ch/categoryDisplay.py?categId=1377
- https://twiki.cern.ch/twiki/bin/view/Atlas/WebHome?topic=FullDressRehearsal

Overall scope

- Test all aspects of offline data flow and processing
  - Realtime tests of hardware, software, databases, storage media, networks, data flow, data integrity and quality checks, calibrations, etc.
  - Identify flaws in our processing models?
  - Identify and correct routine running problems? Where are resources needed?

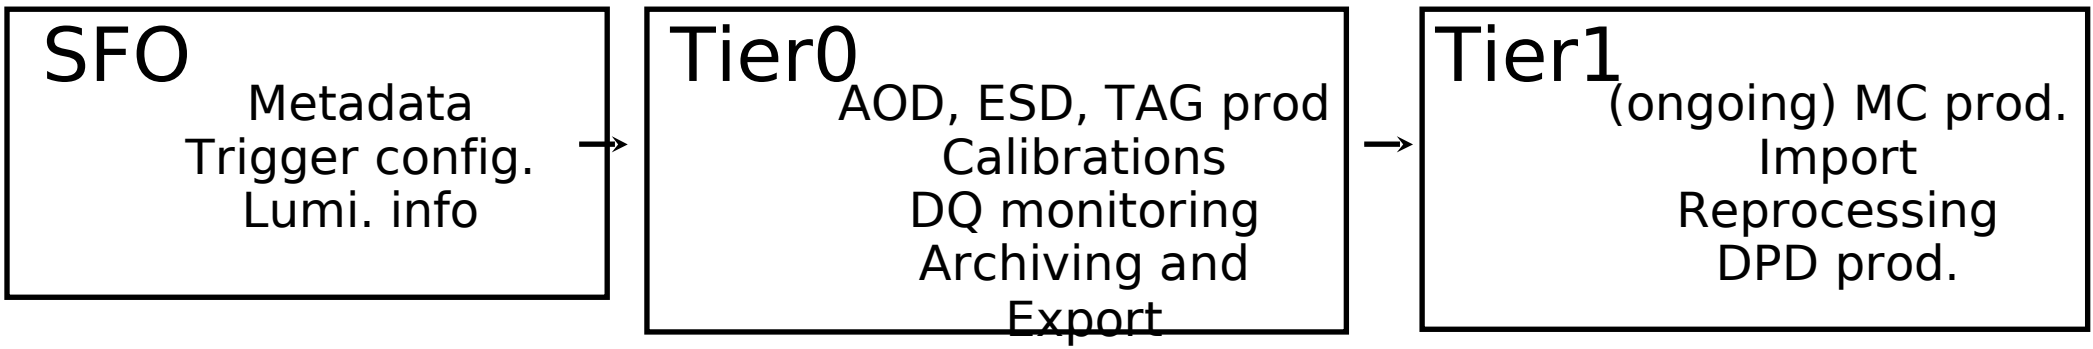

FDR1 lessons and achievements

Ran during early February

Goal: include all significant standard-model processes in known proportions ("mixing") with background events passing trigger chains ("fakes")

- 10 one-hour runs at ~10^{31} $cm^{-1}s^{-1}$ (lumi. varying across run), 1 one-hour run at ~10^{32} $cm^{-1}s^{-1}$ (constant lumi.)
- Produce files in bytestream format, split across streams, SFOs, and luminosity blocks (2 minutes/LB)

- The actual data was
  - 8 runs at ~10^{31} no fakes; 2 runs at ~10^{31} with some e/γ fakes
  - 1 run at ~10^{32} finished late, was processed later

- Used existing CSC simulation CSC-01-02-00 and release 12: not fully consistent (mixed sets of event generators)

FDR1 lessons and achievements

The trigger menu for L = 10^{31} selects mostly fakes: very difficult to simulate.

- The rate of interesting events selected at this lumi. is much lower than the nominal 200 Hz
- Thus, because 8/10 low-lumi. runs do not include enhanced-fake (e.g., EM-like jets) samples, the overall rate is closer to 10 Hz
  - Achieved rate was ~50 Hz for runs with enhanced-fake samples
- Conclusions
  - This was a menu for first data and detector commissioning
  - FDR2 uses simpler menus with far less items and higher trigger rates

FDR1 lessons and achievements

Data played from 5 SFOs during 5-8 Feb (Tues-Fri)

- Same events replayed every day with different run numbers: 3070-3079 is the final sample
- Reconstruction with different AtlasPoint1 patch every day (for monitoring updates and fixes)
  - Flexible: allowed us to run anything at all
  - Dangerous: not reproducible; software not properly validated
  - Final reprocessing in T0 much delayed by need for validated release:
  - Used 13.0.40.4 in early May!!
- Expecting to use 14.1.0.2 for T0 in FDR2 (muon issue?)

Schedule and software versions FDR2:

- "Data prep" on Grid and BNL
  - Simulation with release 13.0.40, Full Geant, CSC-02-00-00 and CSC-02-01-00
  - Production on Grid: digitization run in separate jobs
    - 100 or 1K events per RDO file. Approx 40M events produced
    - Data "mixing" done at BNL: 13.0.40.4 and 13.0.40.5
    - Bytestream files moved to CERN.
    - Note FDR1 mixing was done at CERN: not a success
- Events loaded to SFO's
- Shipped to T0 and reconstructed with 14.1.0.2
- ESD/AOD/TAG exported

FDR2 structure

- Run 4: same as run3 with addition of e/gamma fakes.
  - Have 10M filtered jets: candidates for e/gamma triggers
  - 0.6% of these pass trigger
    - Must prefilter in separate job: 2K jobs of 10hrs each. Ran last weekend
    - Disaster since the event following the one that passed was written out (don't understand why)
    - Recovery plan: extract run/event numbers for passed ones, subtract one and refilerter: 2K jobs at 15 mins each (mostly data staging time). Expect to do this before Monday
  - Expect to start run 4 on Monday
  - Data should arrive in time
- Run 8: 10^{33}. CSC-02-01-00: 30 sec luminosity blocks, 75 ns pileup, beam gas, beam halo, displaced primary vertex: the 'go for broke' run
  - Test job running to get rates and timings
  - Some data still missing: dq2 issue, being worked on.
  - Data will arrive at CERN during FDR2 run week.....
- Note that "detector imperfections" are added in some lumi blocks

Data content of FDR2:

- 5 "runs" with different events: mostly see below. Each is 1hr
  - 6 data streams. Muon, Jets, Egamma, Bphys, Minbias,Express
  - 5 SFO's
  - 1 raw file per sfo per stream per luminosity block
- Runs 1 and 2: 10^{32}. CSC-02-00-00: 1 min luminosity blocks: 13.0.40.4
  - Run 1 is complete, data at CERN: run 2 is 90% complete, will declare completion on Sunday, data transfer to CERN in process.
  - Some data missing
    - A few b samples were not delivered to BNL in time
    - Some triggers underpopulated (later)
- Run 3: 10^{32}. CSC-02-01-00: 30 second luminosity blocks: 13.0.40.5 (forward CSC's present)
  - Started yesterday, 15% complete, expect "completion" by Monday, data transfer starts Wednesday.
- Full list of input simulated samples attached

The “rate problem” I : muons

- Isolated muons are “easy” to get right.
- Non-isolated muons come from
  - Charm and bottom inside jets
  - Could prefilter to generate these more efficiently
    - This is what pythiaB does
    - Must not let these events pass other triggers (e.g. Jets)
    - These are present in FDR2
  - K/pi decay in flight (approx ½ of non isolated muons)
    - Tried to fix this with modified G4 to increase K/pi decay rate: “Biased” min-bias
    - Probability of getting muon to pass any trigger about 1% (0.01% in real min bias)
    - Attempted to get these in FDR1, failed due to data handling: too ambitoius for FDR2

The “rate problem” II : e/gamma and tau

- Most objects passing these triggers are not real e,gamma or tau
  - They are jets
- Prefilter events before simulation
  - Will get bias
  - Accept narrow jets
  - Strategy has been used for many years by e/gamma performance to get enough rate to study algorithms
- Tau fakes are easier than e/gamma as the trigger performance is worse
- We had a set of these prefiltered samples for FDR1:but could not use them
  - Could not use them due to data hanlding and too complex trigger menu
  - FDR2 has a sample targetted to E25/gamma20 etc.
  - 10M biased simulated events yields 60K events that pass
  - These are added to Run4 and increase trigger rate by 100 Hz

Comment on Express stream

Express stream was used successfully in FDR1

- Primary purposes: data quality and calibrations
- Sufficient to detect problems and details, allowing timely validation of data (if DQ is operational)
- Monitoring histograms (detector and performance) available hours after run appeared; shifters spotted problems shortly thereafter
- Histograms moved to AFS to avoid overloading castor with too many requests
- ESDs and AODs not exported—consider temporary storage so users can access them if needed for DQ.
- No luminosity in express stream

Data imperfections found in FDR1

Problems were introduced for short time intervals to test data-quality monitoring:

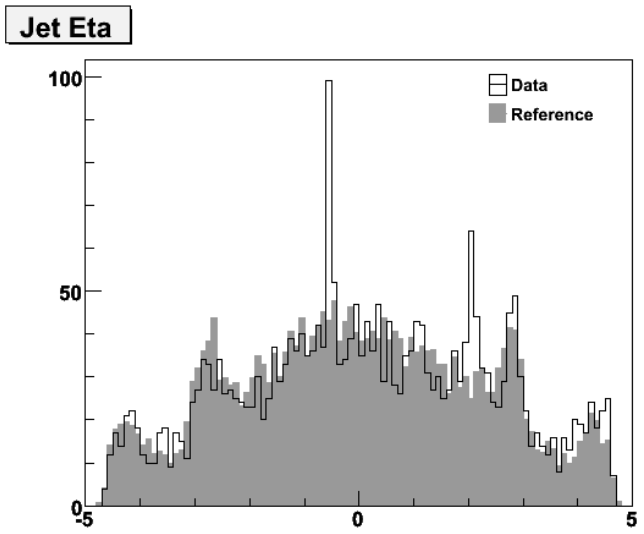

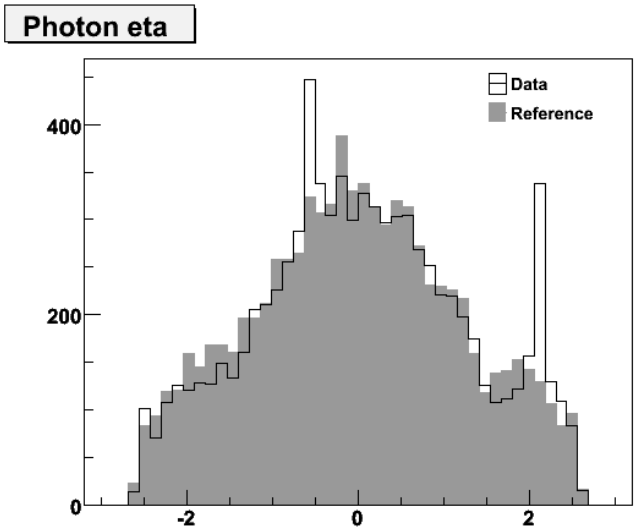

Run 3062, minutes 20-3
hot LAr cells

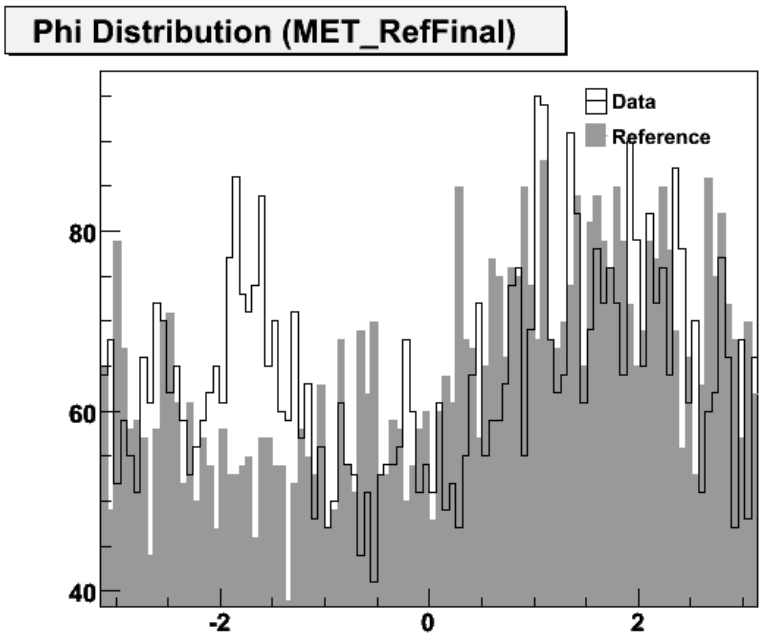

Run 3062,
dead LAr crate (or crates?)

Data quality (from FDR1)

Data quality checked by central shifters and system experts

Bad data was injected into the samples and had to be found

- Included: Pixel, LAr, Tile, MDT, RPC, L1RPC, ID alignment, e/gamma, jets, missing E_T, muon tracks
- Two shifts per day using desk in Tier0 control room
- Exercised ability to spot problems with 10-minute granularity
- Lots of room for improvement; core functionality came online during the week.
- Much learned in FDR1. Expect to do better next week

Bugs found in FDR1

Shoulders at ~150 GeV and ~600 GeV

TRT timing not configured for one wheel (shift of 15ns)

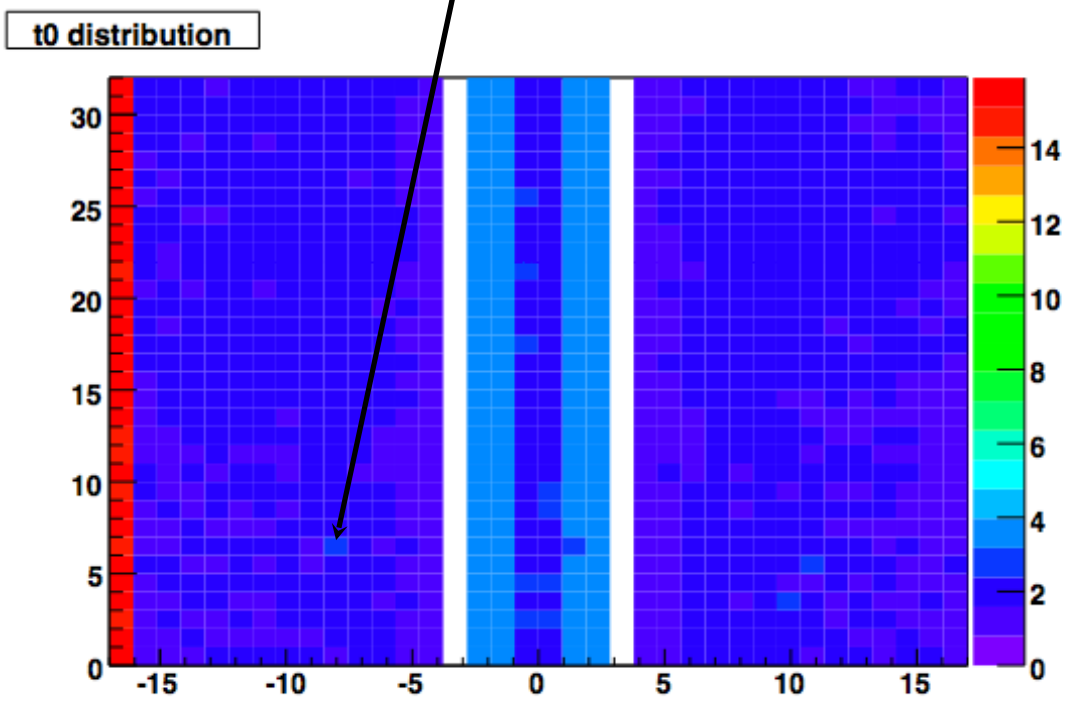

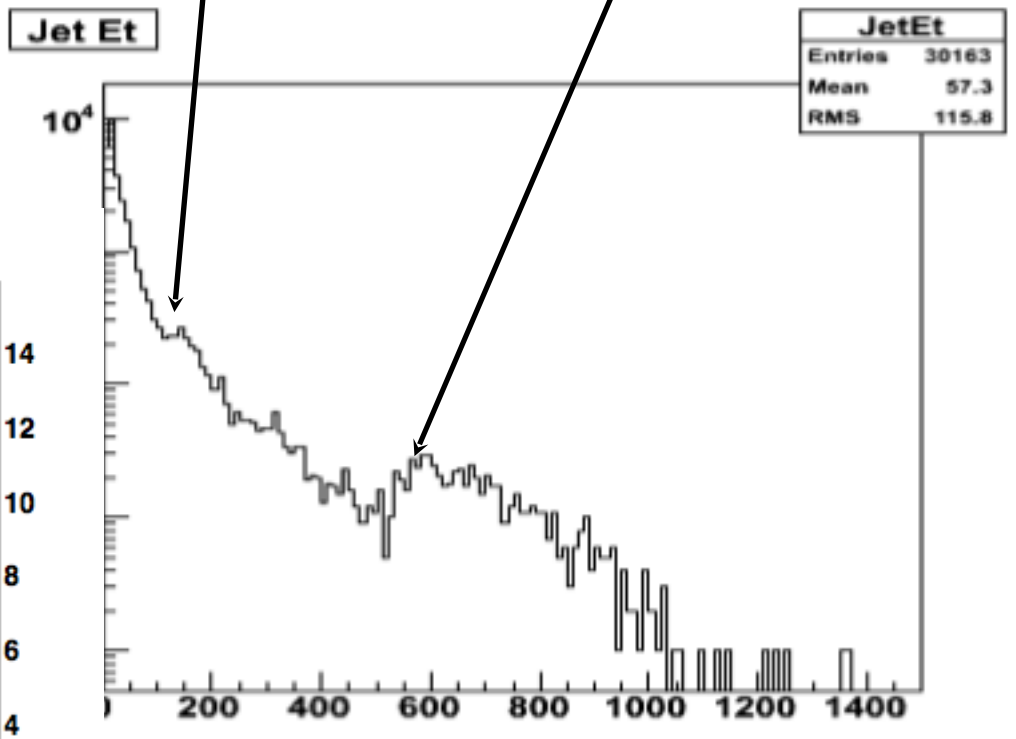

Calibration

Tested the calibration loop: determine calibration constants within 24 hrs of data being recorded and apply new constants during bulk reconstruction

- Pixel calibration planned to run on express stream
  - No code for FDR1.postponed until FDR2
- TRT calibration running on ID-alignment stream
  - First of many iterations completed on Tier0 promptly; remaining iterations finishing on lxbatch

ID Alignment

First test of a calibration stream

- An efficient and dedicated effort had this ready for testing at the beginning of January
- Used by both ID alignment and TRT calibration in FDR1
- Uses event fragments from L2 ("L2_trck10i_calib")
  - Isolated tracks with pT > 10GeV (5GeV for FDR1)
  - Select tracks at 60 Hz (L = 10^{31}) with no additional TDAQ load
  - 600k tracks selected for FDR1 (from a dijet sample; uses 26GB)
- Attempted to include 50k cosmic events; postponed due to software incompatibilities

ID Alignment

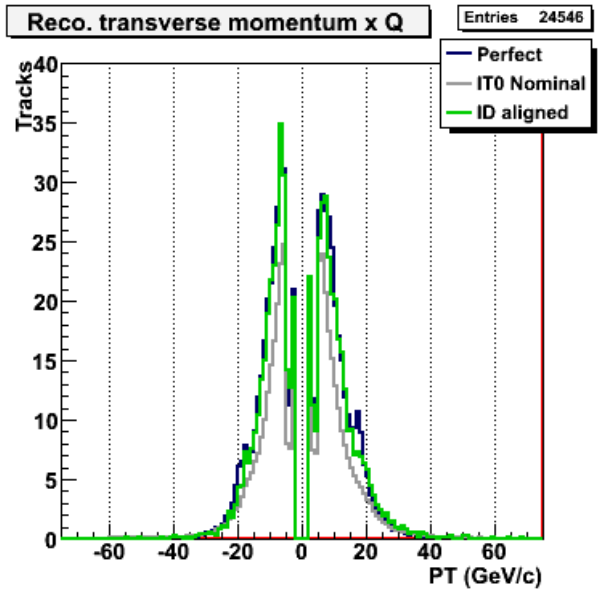

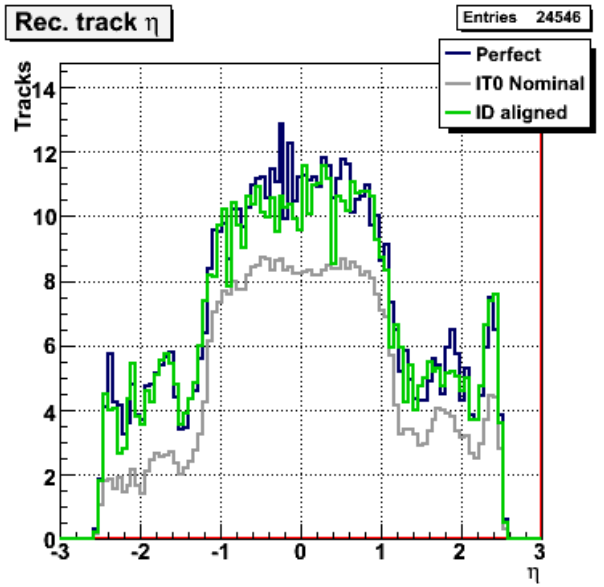

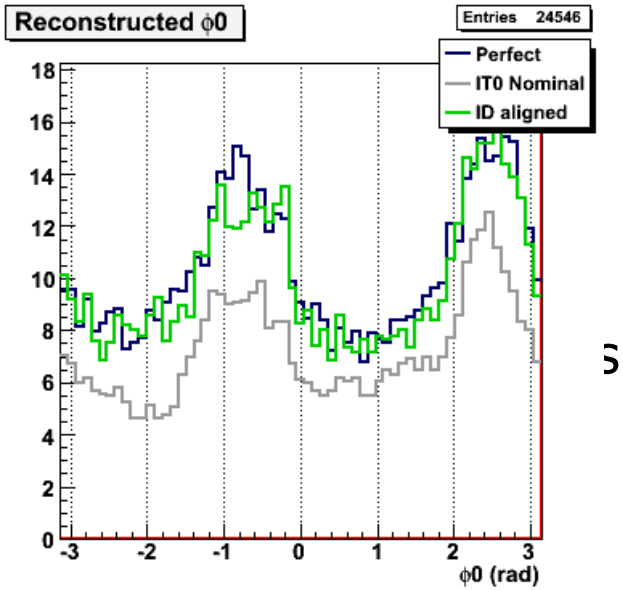

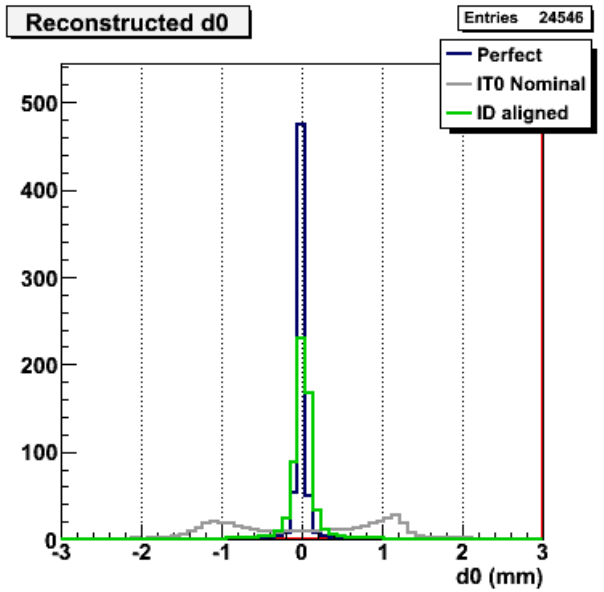

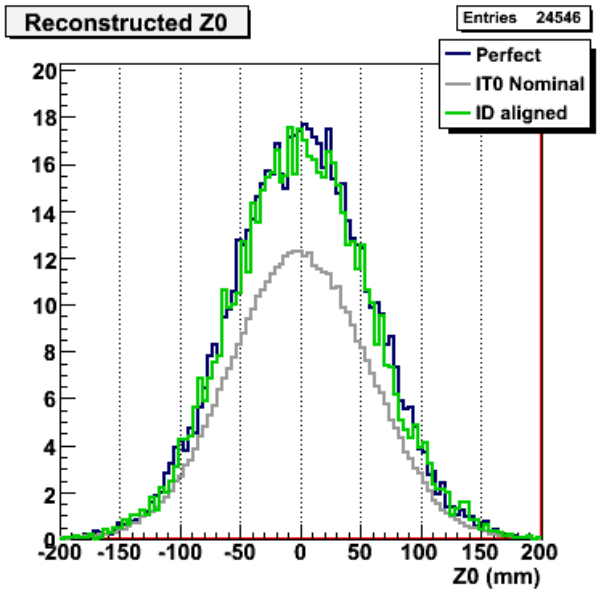

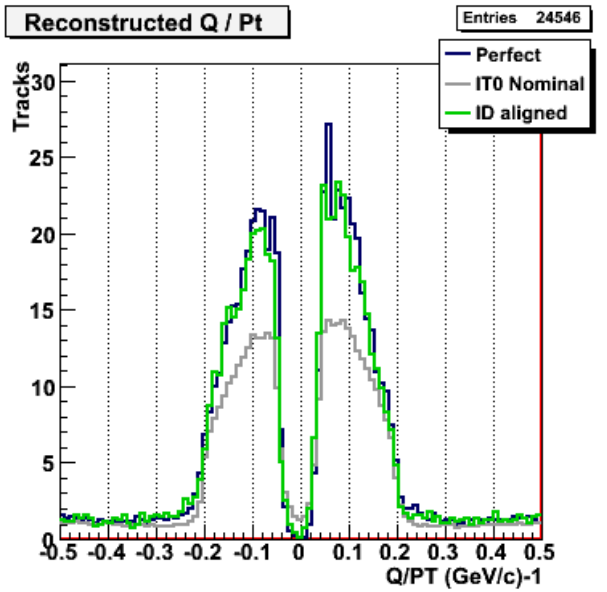

Participation outside CERN?

Information flow to those outside CERN should be improved

- Ideally, a portal where collaborators can check realtime status of
  - Runs recorded
  - Runs with express stream processed
  - Runs waiting for data-quality signoff
  - List of bulk streams processed, and then exported
- How can offsite colleagues participate in meaningful data checks?
- Not just an issue for T0 processing

Trigger rates for FDR2

- Run 1 and 2

- 82 Hz jets (recall comment regarding 3J50 and 4J35)
- 11 Hz: e/gamma
- 51Hz Muons
- 30 Hz B-phys
- 18 Hz Min bias
- 16 Hz Express

- Run 3 and 4: same as run 1 and 2
  - except for fakes added: E/gamma is ~130HZ
  - B-phys and muons slightly bigger: missing samples arrived

Trigger rates for FDR2

- Run 8: Preliminary numbers from test using run 1 data and run8 trigger
  - Egamma: 21 Hz
  - muon 120
  - jet : 65
  - Bphys : 2 (limited by rate out of Lv2)
  - Express : 50
  - Minbias 50

- Will prescale mu20 by factor of 5 to reduce muons to 40Hz
- May drop some samples that only enter Bphys: will speed up jobs at little cost

Some user results from FDR1

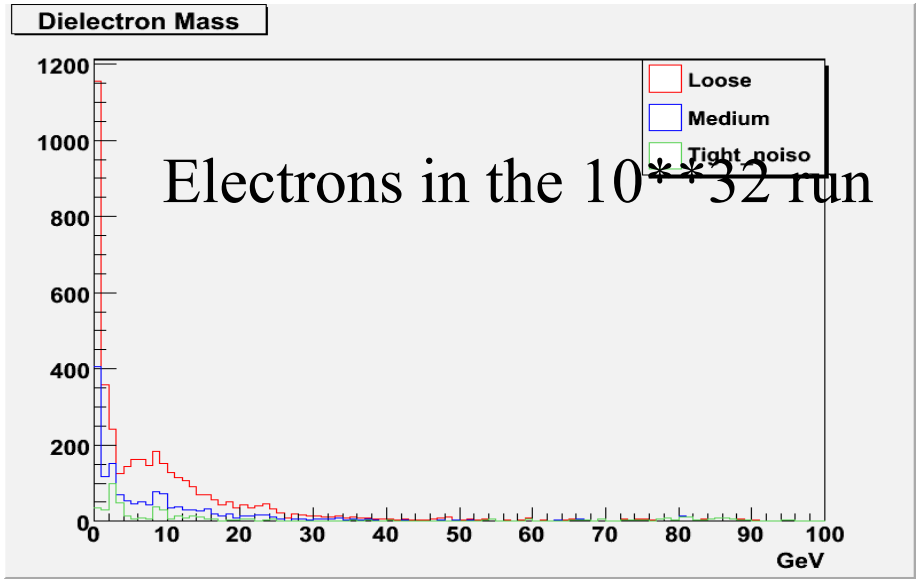

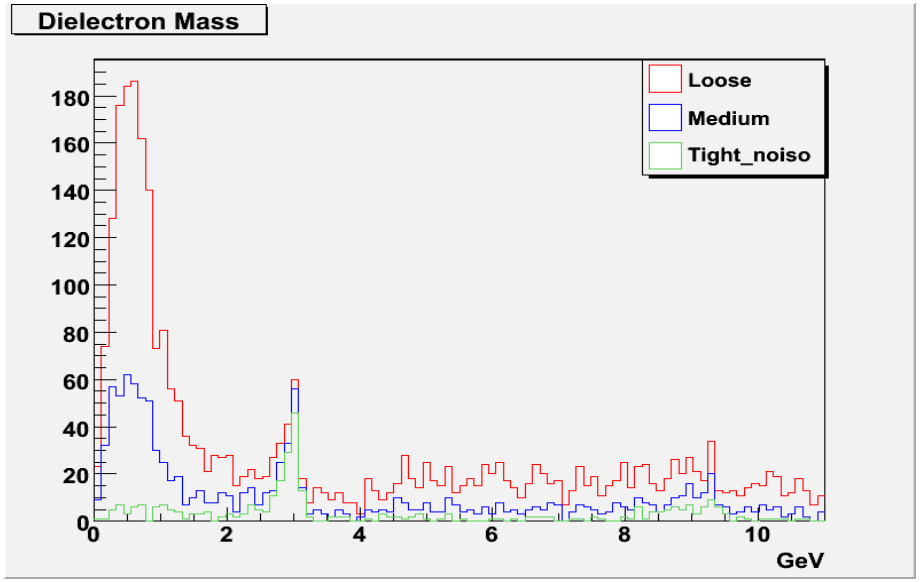

	Loose	Medium	Tight
Total	4239	1694	619
J/ψ	288 = 7.4%	196 = 5.0%	137 = 3.5%
Upsilon	383	164	72

Nelson

Some user results from FDR1

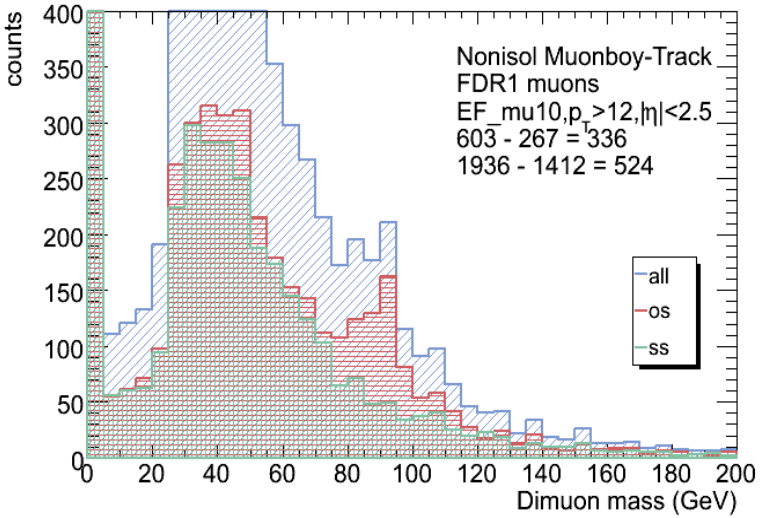

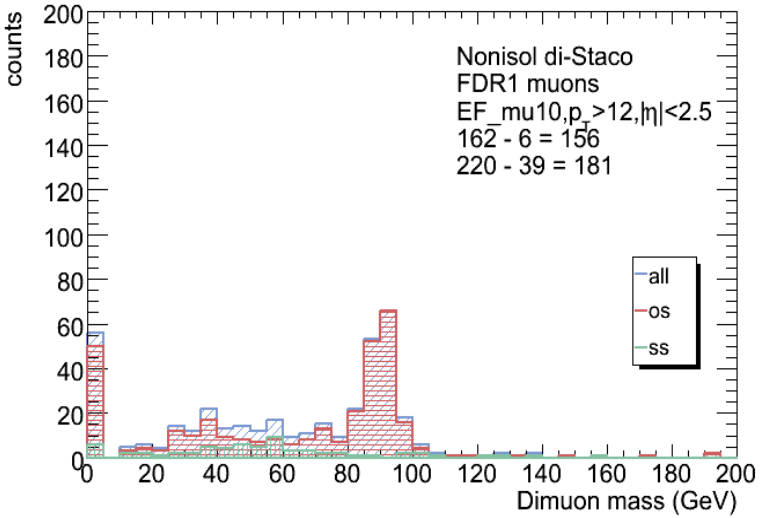

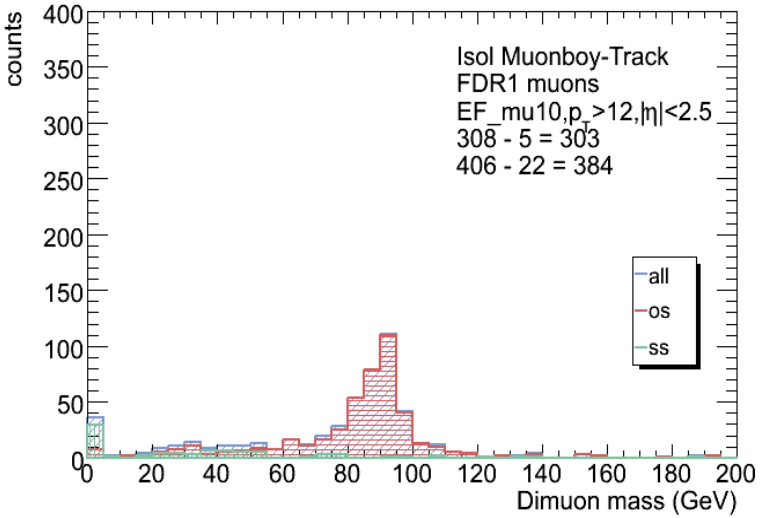

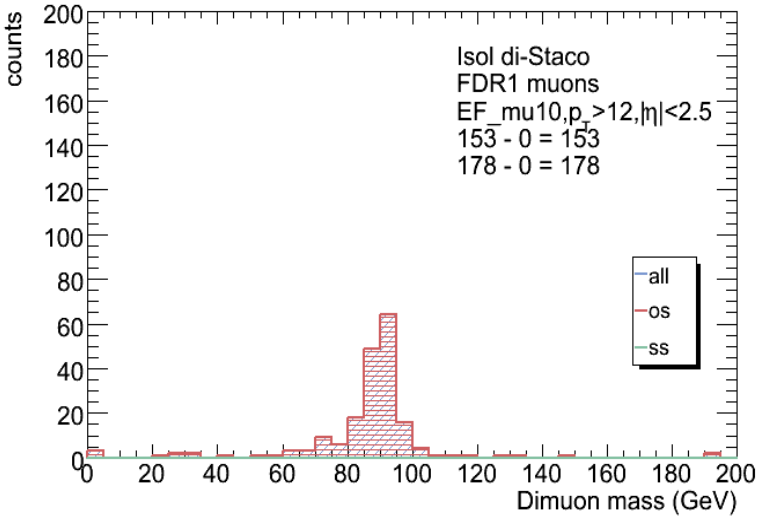

Adams

Dimuon mass: tracker and spectrometer combined

Cross section

- For $Z \rightarrow \mu\mu$ where both muons have
  - $|\eta| < 2.5$
  - $pT > 12$ GeV/c

Adams

Luminosity

- $L = 10e^{31}$ /cm^2/s * 7200 s = 72 nb-1

Trigger efficiency

- Assume 0.80 ± 0.10 from Monte Carlo studies

Cross section

- $\sigma = N/L/\varepsilon trig$
  $= (209 \text{ events}) / (72 \text{ nb-1}) / (0.80)$
  $= 3.6$ nb
- Uncertainty is maybe 20% dominated by luminosity

DPD Production

For FDR-1, DPD production was not a success

- First reconstruction was with unvalidated (transitory) nightlies, not really usable for analysis (though very useful for quick studies)
- Problems in production systems and transforms being fixed (but release 13 closed, so now need to use rel 14)
- Also little data, so AODs usable for studies

For FDR-2

- Investigating option to make first-version primary DPDs at Tier-0 (infrastructure & mechanisms) - David Cote + Tier-0 team, Manuel++
  - For early data, this should be focused mainly at "performance DPDs"
  - Establishing mechanisms is highest priority
  - Time is very short
- Anticipate that rel 14 DPD making in prod sys should be debugged by then - should also be able to make primary DPDs at Tier-1 sites, as always planned

Conclusions

- We are in '"fire drill" mode, hence the quality of this talk
- Data should be much more interesting to users
  - Higher integrated luminosity than FDR1
  - More trigger rate (optimized menu)
- Next week is Tier 0 processing
- Hope to learn more valuable lessons

Extra stuff

The “fake problem” Use case I

- Case number 1: A trigger chain with no prescales and rejections of r1, r2 and r3 at each level. We expect to have M events in the FDR passing this chain
- We start with a sample of N events and end up with N1=N/r1/r2/r3 after running the trigger code. If M>N1 we do not have enough events. This is the "fake" problem. For all "physics samples" we have data so that N1>M. The Trigger decision is that the event passed the chain in the menu. If you rerun the trigger decision on that event it will give True for all the events that end up in the FDR data set. For unprescaled triggers the solution is more events or simulating events that pass the trigger with higher efficiency such as the min bias with enhanced k/pi decay to populate the muon trigger

The “fake problem” Use case II

- Case number 2. A trigger chain with rejections of r1, r2 and r3 at each level and a prescale of factor of f at level 1.
- If we ran the trigger code as is we would get N2=N/r1/r2/r3/f. (If f is 5 and r1=1, Lv1 passes events 1,6,11 etc) If M>N2 we are OK. But if M>N1 and M<N2 we can "solve" the "fake problem" as follows. Instead of prescaling by factor f, prescale by factor of t so that the number of events we pass out of Lv1 is N/r1/t, and the chain passes N3=N/r1/r2/r3/t=M. The rates are right: in effect we have rejected events that were never simulated. The event is marked as having passed the correct trigger chain. What we are doing gives the right result on the average. But if you ran the trigger tool again on any particular event which we passed, there is only a probability of t/f (<1) that this event that would pass this time. (To follow the example if t=3 we would pass events 1,4,7,10)

The “fake problem” Use case III

- Use case number 3: Two separate trigger chains, with Lv1, Lv2 and Lv3 rejections of r1,r2 and r3 and s1, s2 and s3. Prescales f1 and f2 on Lv1 and lv2 of chain one.
- To be very simple assume r1=r2=r3=1 and s1=0.5 and s2=s3=1. f1=2 and f2=3 There are 12 possible trigger patterns labeled by (Sf2f3) as follows (y=fired N=failed -= not relevant): all the remaining triggers fire on every event.
  - YYY,YN-,YYN,YN-,YYN,YN-,NYY,NN-,NYN,NN-,NYN,NN-
- If we have 12 events using the "and" of R and S: we pass 7 events (1-6 and 7 of the patterns) of which 5 pass S alone, one passes R alone and one passes both. You can start an analysis with the events that pass S and measure f1*f2=6 as you would do for real data. If S is feeding a stream and R feeding another inclusive stream you get 6 in the S and 2 in the R stream and 5/12 events have been wasted.
- Suppose we only have 6 events available half of which pass S and half fail. We need to use these as efficiently as possible, so we use them to replace events 1,3,5 and 7,9,11. Of these two passes S alone, one pass R alone and one passes both. You get two in the R stream and and three in the S stream. We have only wasted two events but the ratio is off. In practice we waste 33% of the data, and would waste 42% with the trigger table. If you now start with the events that pass S and "measure" f1*f2, you get 3. You have actually got f2 as I pushed the problem into f1. This is a general result. We tried to use the right thing for the Lv2 prescales.