

# *CDF Offline*



Kevin McFarland  
Rochester/Fermilab  
on behalf of the  
*CDF Offline Operations Group*

May 30, 2002

- 
1. Personnel and Organization
  2. Infrastructure
  3. Data Handling, Farms and Level-3
  4. Reconstruction and Simulation
    - Current releases and plans
  5. Perspectives and Outlook

*Separate talks on CAF and Tracking Topics to follow...*

## Personnel and Organization

- Personnel

- ↪ Pierre Savard has stepped down as co-head of calorimetry reconstruction
  - ★ Beate Heinemann joins Argonne::BobWagner
- ↪ Matt Herndon has joined Avi Yagil as tracking co-head
- ↪ I will step down June 30th

- Organization

- ↪ Spokespeople have proposed a new organization
- ↪ “Offline operations” will cease to exist in its present form
  - ★ Computing hardware systems → “Operations” group
  - ★ Infrastructure, Reconstruction → “Analysis” group
- ↪ Discussion at Executive Board meeting this afternoon

## Infrastructure/Performance Issues

- “I/O performance” in a minute (Liz, Rob K.) . . .
- Compilers
  - ↪ KAI has been purchased by Intel; support is already ramping down, gone by end 2003
  - ↪ gcc Progress (Pasha, Liz&librarians, Art, Chuck, Andy)
    - ★ Code compiles, links, e.g., finds tracks
    - ★ No serious comparison yet with KAI because . . .
    - ★ Libraries are  $\sim 10\times$  bigger; link times  $100\times$  slower
    - ★ We know of specific problems that cause this
      - CD/CDF/D0 task force investigating
      - Testing gcc 3.1 (released May 15)
    - ★ No “Plan B” at the moment, just ideas
- RedHat 7 Migration (Art, Chuck, Andy, Pasha, Farms group)
  - ↪ Running code under RH7 only on farms, CAF
  - ↪ Trailers, offsite partly migrated
  - ↪ *So far*, same code gives same results
    - ★ Encouraging speedy migration of Level-3, desktops
- Optimization (Andy)
  - ↪ +K1 (C++ inlining) now default
  - ↪ Have tested higher levels of compiler optimization

*People working on these issues are lonely heroes of the experiment; desperately need more personnel here*

## Production Monitoring

- New web-based offline monitoring framework ((Pasha, Rob S., Angela)
  - ↪ Check offline code and data
- Compare production output histograms to reference
  - ↪ Hierarchical: pass/stream/validation module/hist
  - ↪ Used to verify stability under RH6→RH7, 4.5.0 vs 4.3.2
  - ↪ *Has identified problems!*

[http://ncdf41.fnal.gov/val/validation\\_C.so](http://ncdf41.fnal.gov/val/validation_C.so)

# Production Monitoring (cont'd)

The screenshot displays a web-based production monitoring interface with several panels:

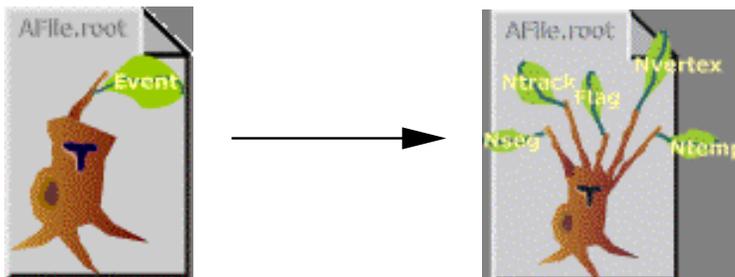
- Top Left Panel:** A sidebar menu for "CDF Offline Validation" with options for "method =", "ValidationPass::MainPage", "Save Validation Results", "Stream" (test2), and "History". A list of runs (138671-140819) is shown with status indicators.
- Top Middle Panel:** A list of validation metrics including energy and counts for west and east detectors, strip occupancy, and eta/phi correlations.
- Top Right Panel:** A plot titled "PES\_Phi" showing a fluctuating signal over time (0 to 5). The y-axis ranges from 0 to 16000. A legend indicates "PE", "Entry", "Mean", and "RMS".
- Middle Left Panel:** A file index for "/cdf/data05/4.5.0pre4a\_611\_02/test1/hist/sum\_hist\_138672.root/CprClusterMon;1/". It lists files such as "PTH1C;1", "Phi(Trk)\_vs\_Phi(MC);1", "Eta(Trk)\_vs\_Eta(MC);1", "Sim\_Pt\_Resolution;1", "QCPR\_vs\_PhyPT\_(track);1", "X\_(Cluster)\_vs\_X\_(trk\_extrapolated);1", "MuonType\_of\_muon\_tracks\_in\_CPR;1", "NHits\_for\_tracks\_in\_CPR\_fiducial\_volume;1", "D0\_of\_tracks\_in\_CPR\_fiducial\_volume;1", "#eta\_of\_hit\_position\_at\_CPR;1", "#phi\_of\_hit\_position\_at\_CPR;1", and "Track\_X\_position\_at\_CPR;1".
- Middle Right Panel:** A plot titled "NHits for tracks in CPR fiducial volume" showing a distribution of hits. The y-axis ranges from 0 to 45000, and the x-axis ranges from 0 to 100. A legend indicates "PE", "Entry", "Mean", and "RMS".

## Event I/O

- Common and correct complaint is “I/O” is too slow
  - ↪ ROOT I/O itself, the ability to read ROOT files, is fast
  - ↪ “Streaming objects”: get object from file to memory
    - ★ Fast for “blocks” of common data (e.g., RAW Banks)
    - ★ Significantly slower for our big heterogeneous reconstruction output objects (e.g., CdfTrack, Cluster/Strip Collections)
  - ↪ “Puffing objects”
    - ★ Complicated calculations required to use objects reside here
    - ★ PAD “puff” or adding geometrical information (design issues?)
    - ★ Can be turned on/off object-by-object; often unnecessary (e.g., slow CalData puff)
    - ★ *Currently all disabled by default*
- Anecdotal performance (I ran stream B stripped electrons, output of 4.3.x production)
  - ↪ Block read  $\sim 5$  GHz-ms
  - Streaming  $\approx 20$  GHz-ms
  - Puffing  $\approx 50$  GHz-ms

## Multi-Branch Events: The idea

- Multi-branch ROOT I/O
  - ↪ ROOT supports a feature of the file structure called “branches”
  - ↪ I/O on each branch *independently*
  - ↪ Our events can be stored in multiple branches



- Analysis examples:
  - A. An event skim that wants to form a new dataset by L1/L2/L3 triggers
    - (a) Read in header branches (small part of event)
    - (b) Make trigger selection
    - (c) Write out dataset
      - ↪ Dominated by “block read” speed, not interpreting event
  - B. Need to redo jet clustering in calorimeter with new algorithm
    - (a) Read in header+calorimeter ( $\sim 10\%$ )
    - (b) Redo clustering; replace jet collections
    - (c) Write out dataset

## Multi-Branch Performance

- Framework (4.6.0) and EDM (4.5.x) supported (Liz, Rob K.)

↪ Ready for 4.6.x production!

- Stand-alone tests (stream A prod. output) (Liz)

Operation	GHz-ms/ev	Comment
Write overhead	1	unmeasurable?
Read overhead	3	
Block read	11	
Block write	55	disk I/O component?
Stream header	4	large trigger banks
Stream electrons	1	unmeasurable?
Stream RAW, Level-3	20	
Stream all	50	

- Bonus: 8% size reduction (compression)
- B group has tested in stripping jobs (Paus et al)
- Current assignments of event to “branch sets”
  - ↪ Level3, Large RAW data to separate branch sets
  - ↪ Trigger results, header to primary (header) branch set
  - ↪ Each reconstructed class to separate branch set (some small ones merged)
  - ↪ User can control I/O at level of branch set

## PADs

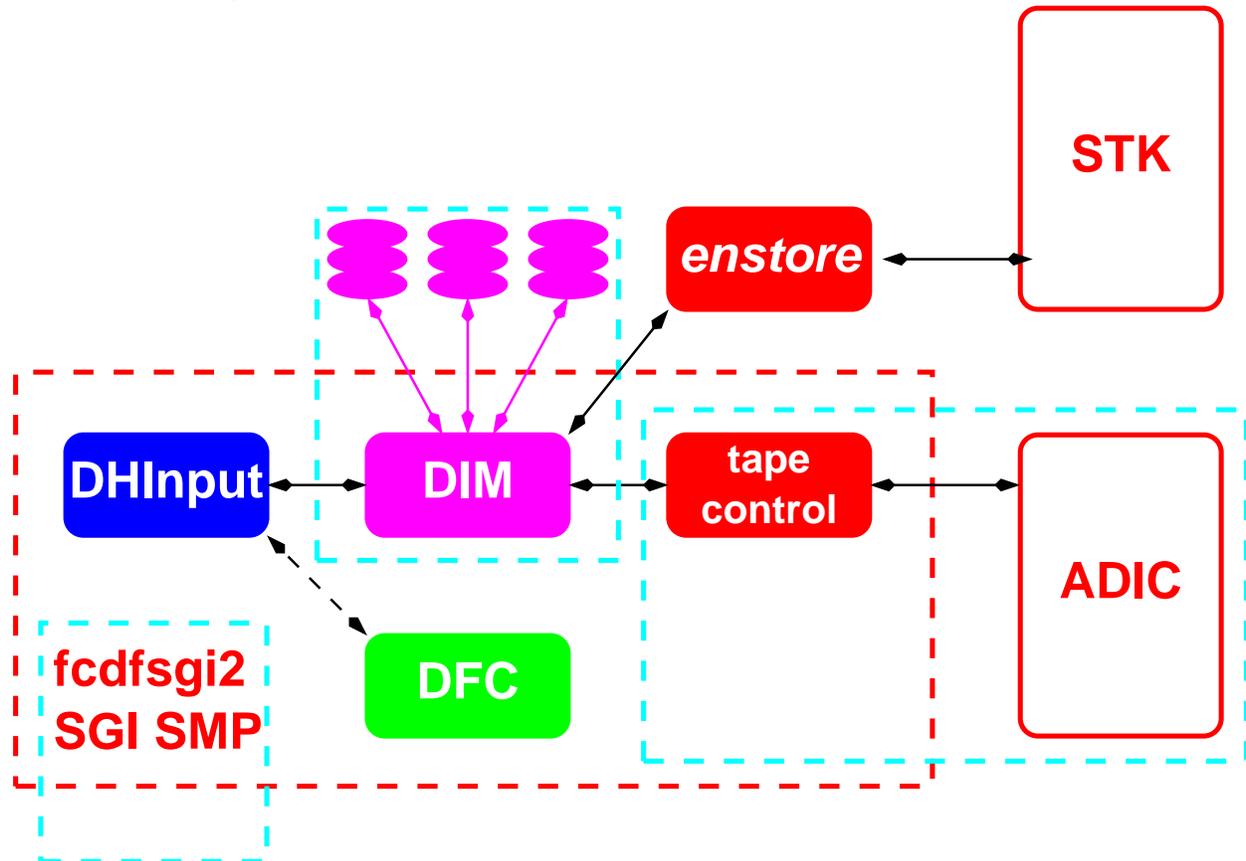
- Recent PADs progress is disappointing
  - ↪ COTD→COTQ is roughly a factor of two, with no information loss
    - ★ But no one has validated results
  - ↪ SIXQ not available
  - ↪ TRKQ: puffing much too slow, need to reassess strategy? (Ivan heroic in getting us this far...)
- However, pruned defTracks view for 4.6.0 (Jason)
- Trigger summary object for 4.6.0 (Liz, Kevin)
- This is big problem is how to optimize use of compressed, summary objects in real analysis  
*This is a hard problem to attack in a vacuum*
  - ↪ Recommendation: informed by summer analysis experience and first results of analysis from 4.6.x production, we should review this to set direction
  - ↪ Available expert effort is limited; need to prioritize directions based on experience

## Databases

- Current focus: how to avoid single point bottleneck in serving data to offline processes
  - ↪ Installing more (Linux) Oracle replicas
  - ↪ Long term remote solution: MySQL? (Waters)
- Summer upgrade to Oracle v9.2
- Lots of work to do in offline access software
  - ↪ Calibration API in good shape
    - ★ Adding more flexible access to Calibration sets
  - ↪ RunSummaryData object ready soon, including GoodRun filter
  - ↪ Slow control access – only worker is now a summer student
  - ↪ Work to make trigger information globally available to offline modules ongoing

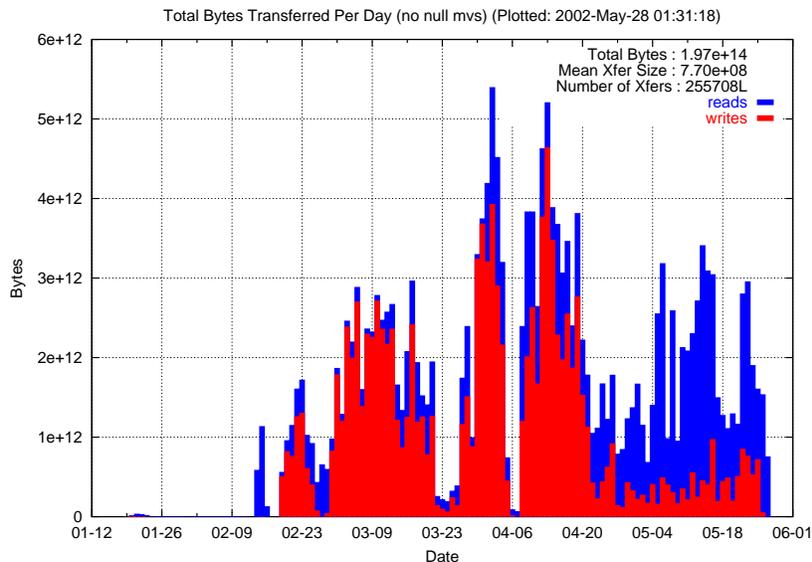
## DH: Enstore at CDF

Operating old and new in parallel



- **Enstore:** network-attached tape (ISD)
- STK silo, 10 mover nodes (18 Feb)

## DH: Enstore Migration



✓ all data in ADIC copied to CDFEN  
(D. Litvintsev)

- avg write 14 MB/s (65 MB/s peak)
- total 100 TB by 22 April (3 months)

✓ production farms I/O (24 April)

✓ raw data logging (1 May)

- kinks worked out in CDF DH and Enstore
- great collaboration with ISD

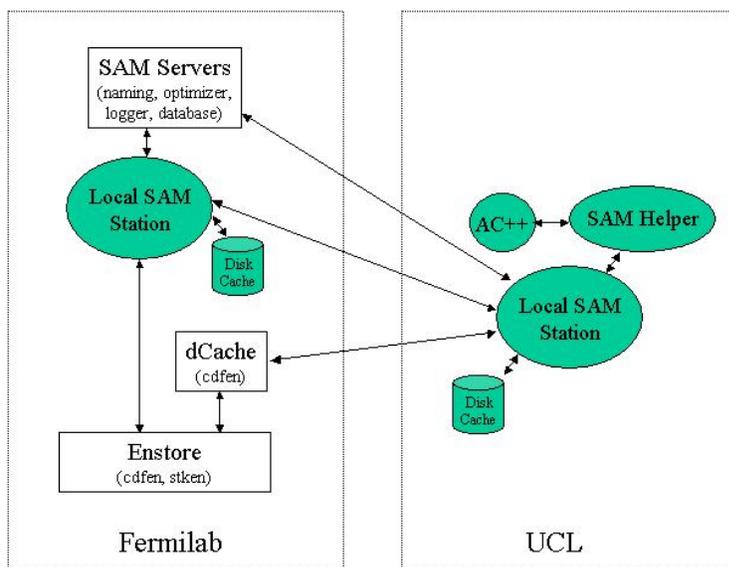
⇒ much smoother operation

## DH: Developments

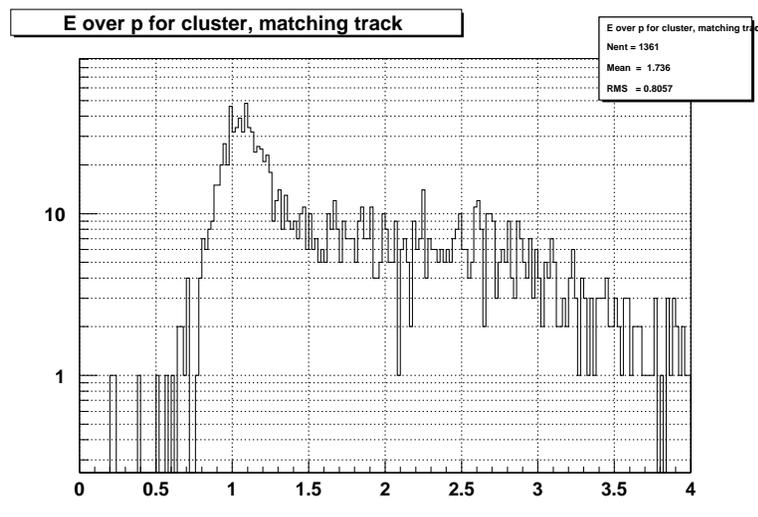
- **network-attached disk caches**  
dCache, DESY/ISD product (Enstore folks)  
(CDF side: R. Kennedy, J. Tseng)
  - ✓ AC++ interface
  - ✓ prototype read pools
    - write pools in development
- **remote computing, metadata management**  
SAM, originally DØ/CD project
  - active collaboration  
(R. St. Denis, D. Waters, T. Huffman, A. Sill, R. Herber, R. Glosson, D. Litvintsev,...)
  - ✓ AC++, dCache interface
    - DFC→SAM metadata underway
    - ↔ prototype next week
    - long-term direction for overall DH
- **CDF notes, user guides in preparation**
- Contact DH group to play on the edge  
[cdfdh@fnal.gov](mailto:cdfdh@fnal.gov)

## DH: SAMple Analysis

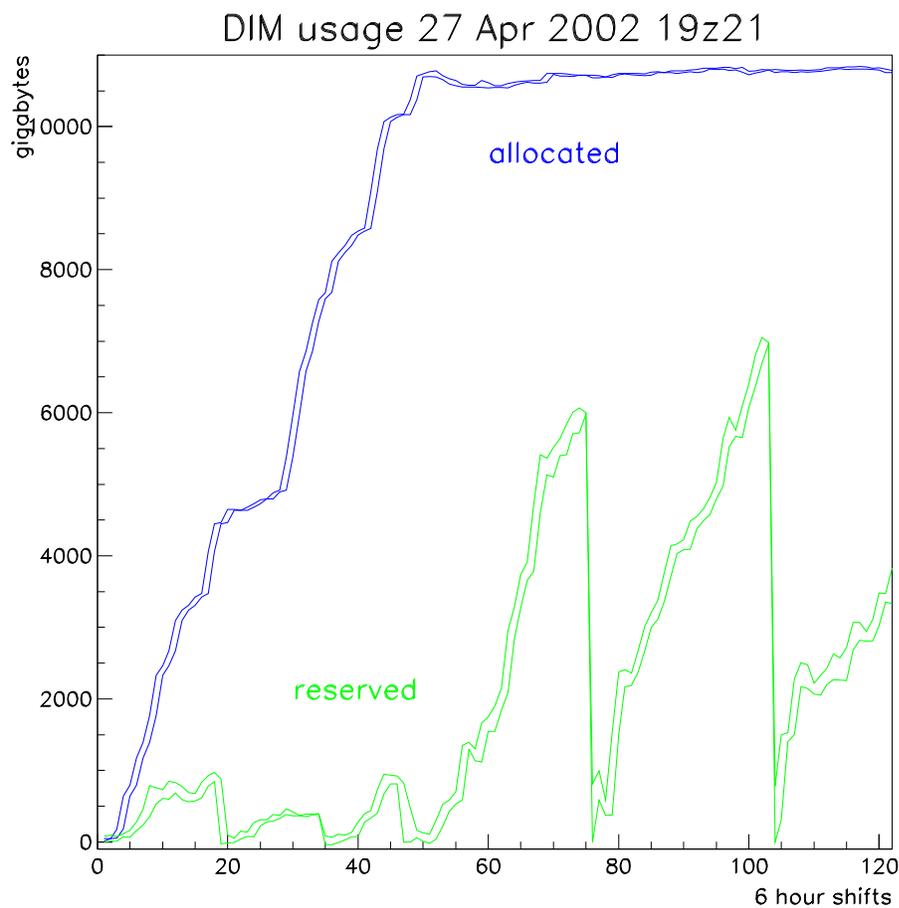
- several SAM stations already deployed
- used to analyze data drawn from Fermilab



⇒  $E/p$ : AC++ in London, data from CDFEN



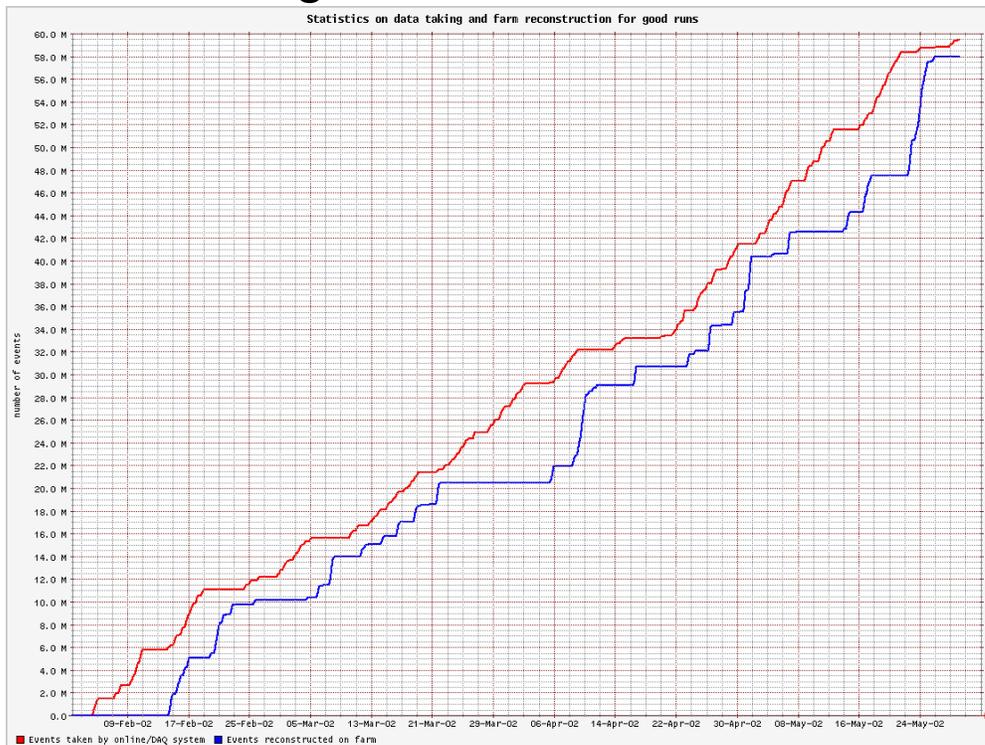
## DH: DIM Usage



- Cache miss rate 28%
  - limited datasets
- ⇒ probably not pushing system yet
- still need more readily decentralized system

## Production Farm

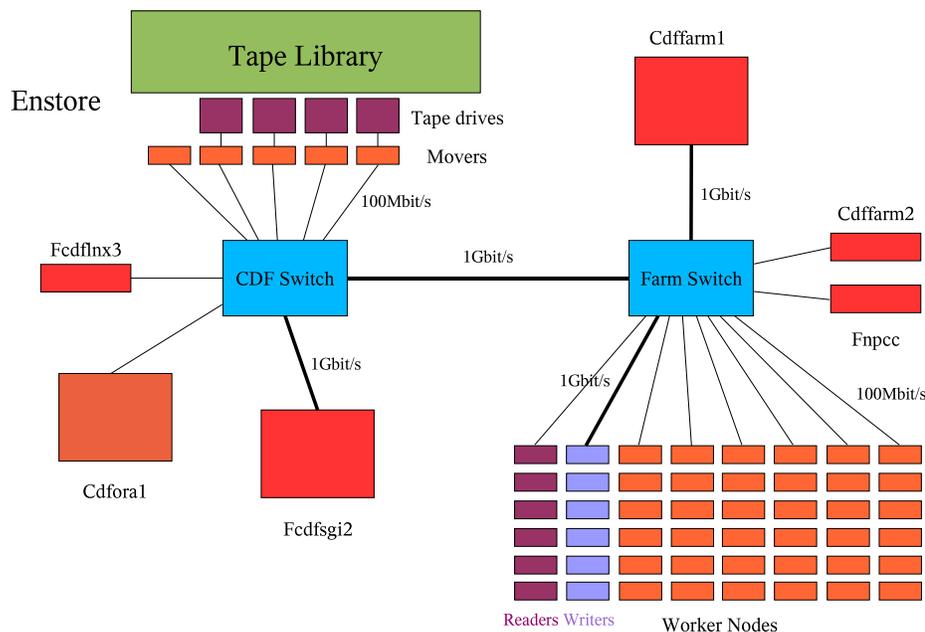
- 58 of 59M events (as of 5/27/02) since 8x20 processed using 4.3.2 ProductionExe



- 10M events reprocessed with 4.5.0
  - ↪ Plan 20M events with 4.5.2 soon
- Large MC production batches have been run; operations still big effort
- Hardware
  - ↪ 169 dual PCs in production
  - ↪ 4 concatenation nodes with Gb links for Enstore
  - ↪ Capacity: 3.5M events/day (at 70% efficiency)

## Production Farm I/O

- Farms have made transition to Enstore DH



- All farms input/output through Enstore  
Input still indirect, through `sgi1`
- Current output bandwidth (old limit) is very far from limiting factor now
- Peak aggregate bandwidth will be 60–70 MB/s
  - ↪ Practical operation at 35 MB/s, 150 Hz
  - ↪ A CPU bound farm! Kudos to Farms and DH groups!

## Level-3 Status

- Running 4.3.2; plan to upgrade to 4.6.x when available
  - ↪ Then, Si tracking (discussed later)
- Path interference problem found (diphoton candidates failed electron trigger)
  - ↪ Technical problem was that both filter and reconstruction code picked up wrong information from the event
    - ★ Code policy and existing technical hooks were ignored
  - ↪ Inadequate validation allowed it to go undiscovered
  - ↪ Currently running unvalidated but more conservative “all reconstruction” approach

## Level-3 Status

- Review of code (Savard, Veramendi) found
  - ↪ Inconsistent error logging
  - ↪ Overloading of variables for cuts (Lshr)
  - ↪ Inability of code to flag when needed reconstruction failed
  - ↪ Unspecified selection of data from event
  - ↪ Too difficult to test Level-3
  - ↪ Need for more atomic filters (generic filters too complex)
- Level-3/Offline group and review are working out technical solutions
  - ↪ Many offline changes for 4.6.0
  - ↪ L3 test platform (Farrukh, Dave, Jonas, Jeff)
  - ↪ “all reconstruction” vs “path” validation done (Farrukh)
  - ↪ Working on “single path” vs “all path” tests

## Level-3 Silicon Calibrations

F.Azfar, A.Dominguez, M.Herndon, P.Maksimovic, J.Nielsen, S.Nahn,  
A.Sill, D.Waters, I.Volobouev ... and others

### Operational Issues :

- Level3 runs reconstruction from a text replica of the Oracle database.
- The main silicon calibration tables are represented in the text database :

```
level3@b0dap31 6> ls -al calib/3242/DATA/SI*
8939755 May 14 16:11 calib/3242/DATA/SICHIPPED.144630.1.80690.COMplete.UNDEFINED.
      6 May 14 16:11 calib/3242/DATA/SIFIBOFFSETS.111193.1.22743.COMplete.UNDEFIN
150873 May 14 16:11 calib/3242/DATA/SISTRIPDH.144308.2.80528.COMplete.UNDEFINED.
      ...
```

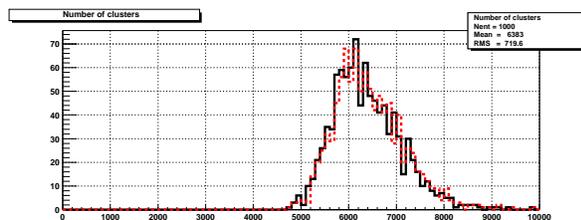
- Silicon reconstruction is a special case because :
  - ↪ Calibration uses more complex database “Views”
  - ↪ Need pedestal subtraction mode from USED\_SETS table.

Currently the text database export and distribution schemes support neither of these.

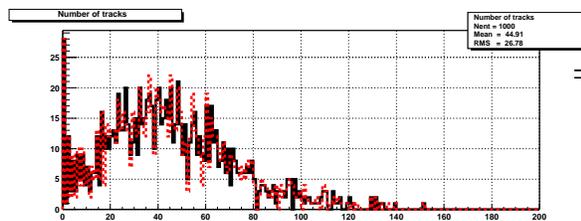
- A “Level3” mode has been developed which assumes pedestal subtraction on the FIBs, avoiding these difficulties.
- Exporting calibration tables in their present form into text database : tar-balls increase by over 50% to approx 1.7 MB
- The time the export takes increases from 10s ⇒ 45s
  - ↪ Develop versions of the tables without the backend state multiplicity factor of 4?
  - ↪ Leaving out the pedestal information altogether in the Level3 versions of the calibration tables, provided this is shown to be safe?

## Level-3 Silicon: Reconstruction

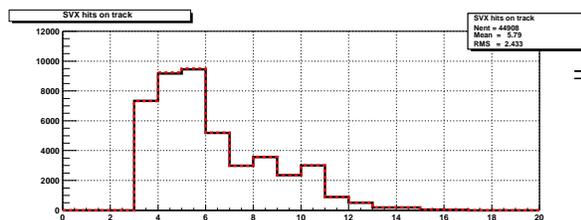
- There are different scenarios for what calibrations are to be applied at Level3.
- POSSIBLE BENEFITS : smaller calibration tar-balls, faster run starts and faster silicon clustering event times, if certain calibrations can be skipped
- POSSIBLE PIT-FALLS : complications if silicon Level3 reconstruction is not IDENTICAL to offline, for example in the calculation of trigger efficiencies.



BLACK : full offline  
RED : do not re-do  
pedestal subtraction



⇒ Small differences in track multiplicities.



⇒ Hit content of tracks identical

### Bottom Line:

- Silicon reconstruction now ready to be run at Level3 in special “Level3” mode.
- This is probably adequate long term, but this is to be confirmed in detailed tracking studies.
- Several ideas around for how to improve performance. Lots more work needs to be done.

# Muons

- Highlights of Progress

- ↪ IMU reconstruction in place; tuning (Bellinger et al)

- PLOT**

- ↪ Stubless muons done,  
“ready to confuse the world” – Ken

- ↪ Cosmic ray tagging (Anyes et al)

- PLOT**

- ↪ Documentation: CDF 5870. Yea, muons!

- ↪ GEANT Extrapolator ready for 4.6.0 (Andreas, Slava)

- ↪ CMU, CMP alignment (Victoria)

- Areas of activity

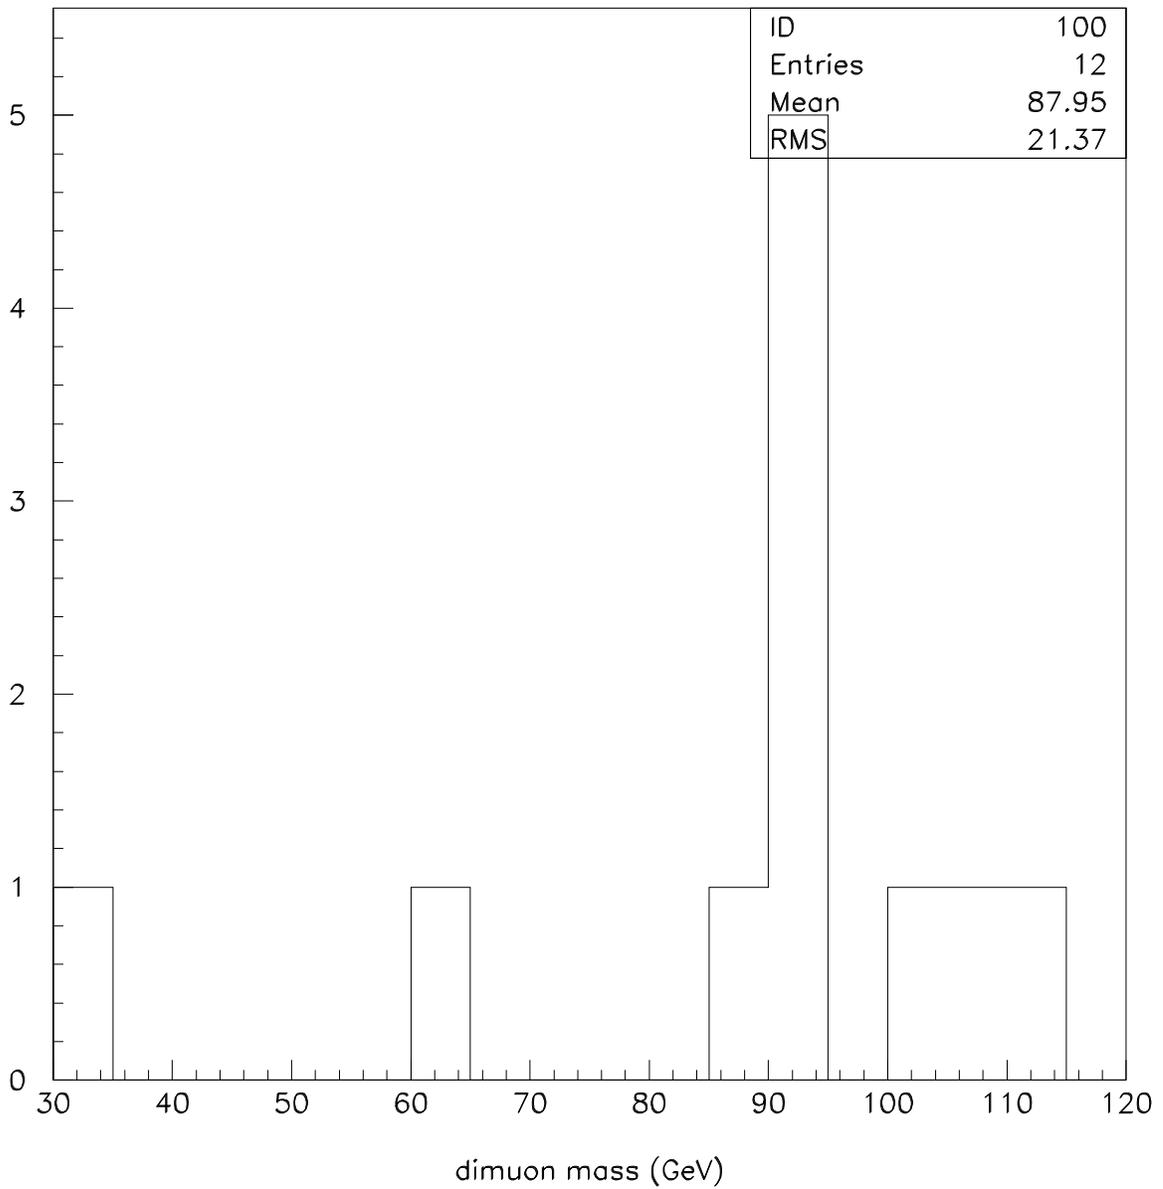
- ↪ Monitoring code in production

- ↪ Adding scintillator hits to muons (Victoria)

- ↪ Reconstruction efficiency (Ken, Dagenhart)

- ★ Fiducial volume tool for  
CMU, CMP, CMX in place (Dagenhart)

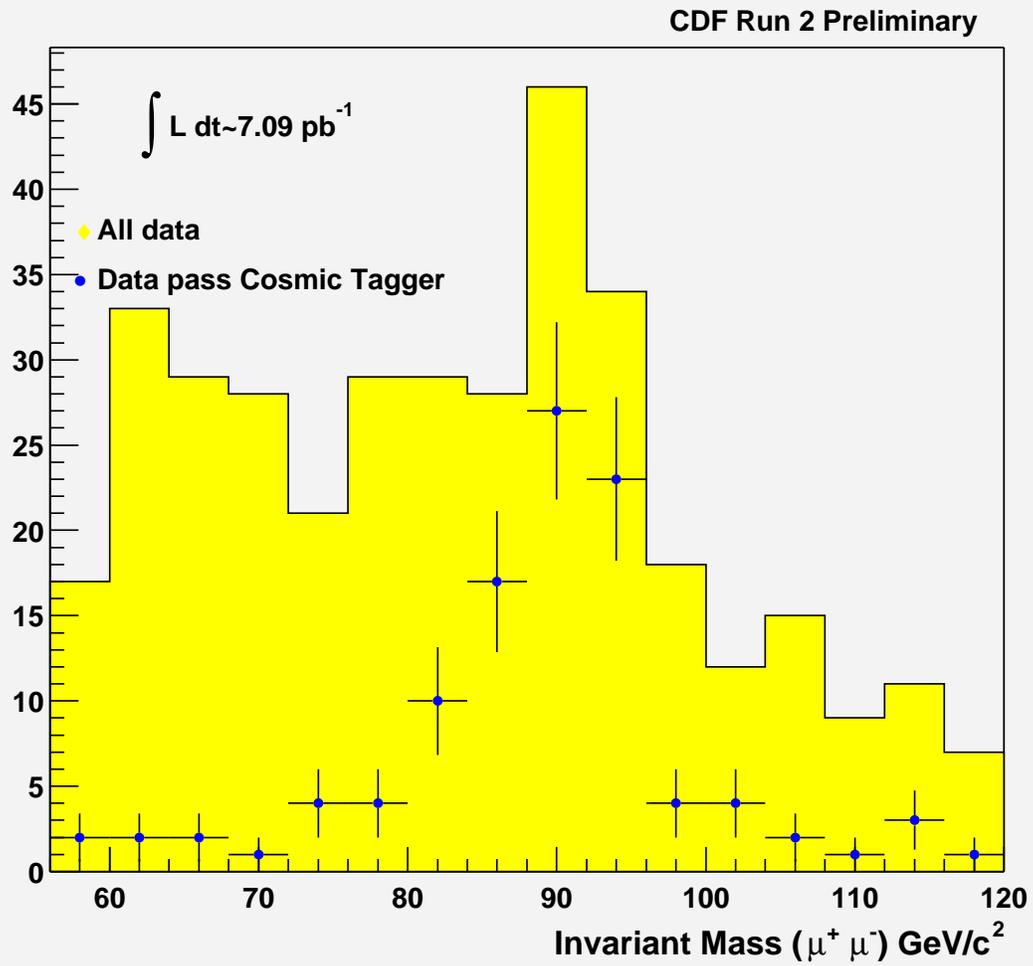
# Muons (cont'd)



BMU-CMU/P  $Z$  Candidates

# Muons (cont'd)

## Z Invariant Mass



# Calorimetry

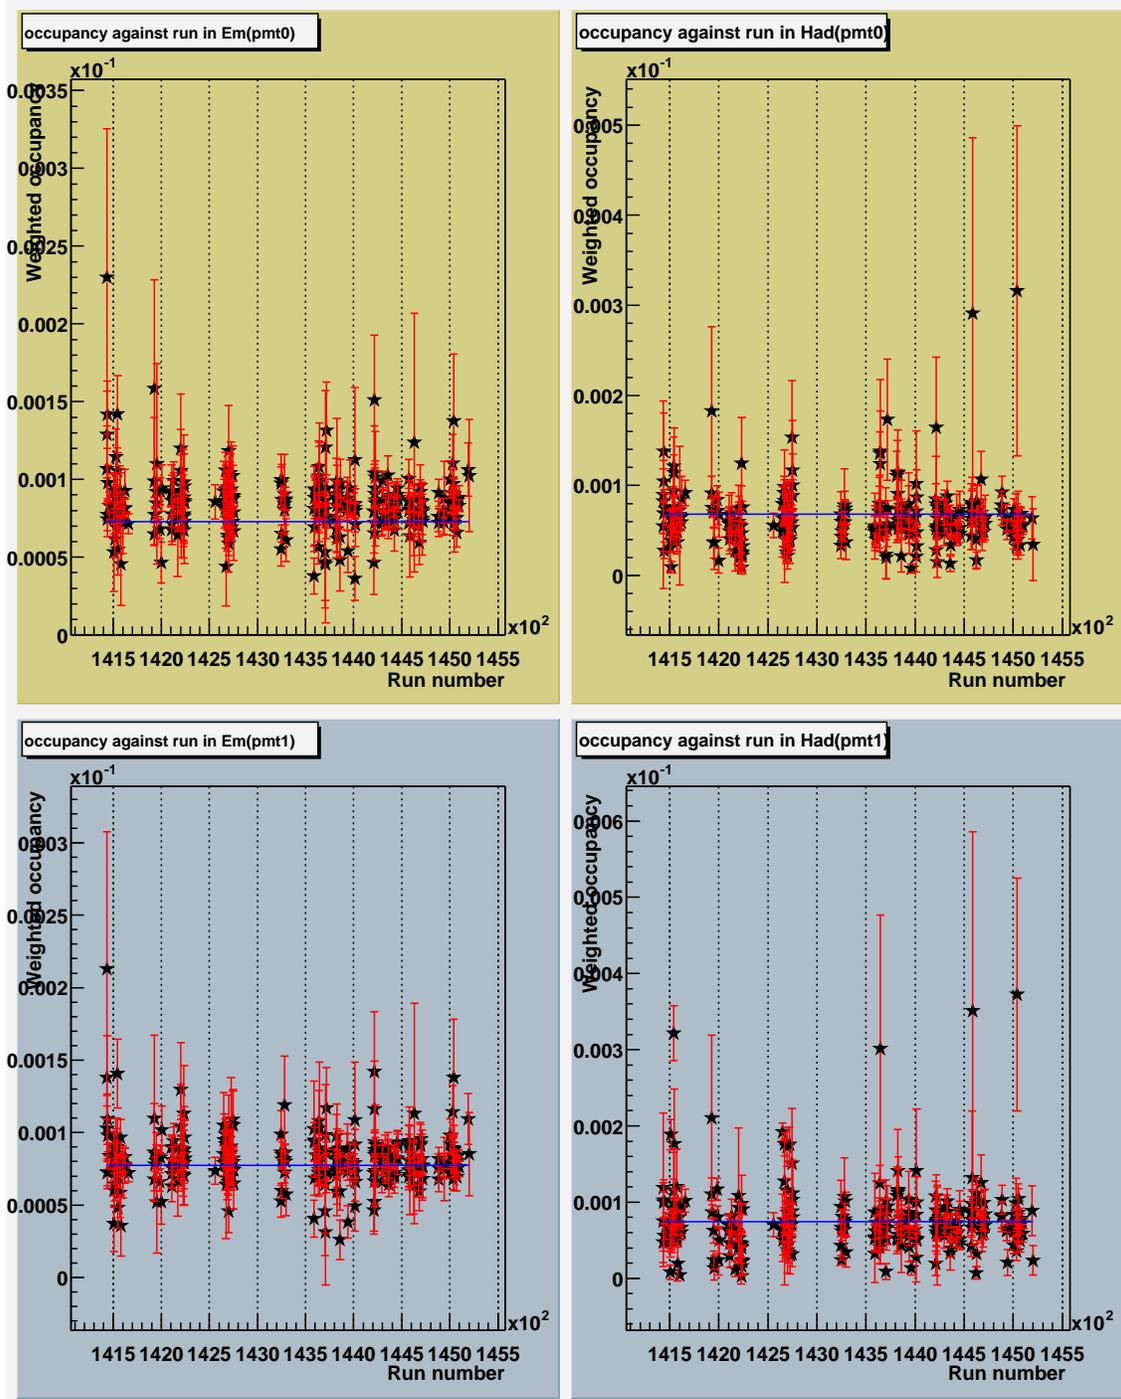
- Highlights of Progress

- ↪ Tower-by-tower corrections (8 GeV electron  $E/p$ ) (Eva)
  - ★ Clear improvement in  $\cancel{E}_T$  (Dmitri)
  - ★ Measuring CHA corrections with  $\mu\text{s}$  (Dmitri)
  - ★ PEM, WHA, PHA? Si standalone. . .
- ↪ Sophisticated production monitor ((Martin, Costas)  
**PLOT**
- ↪ New hadron TDC calibration
- ↪ MiniPlug in CalData ((Koji, Beate)
- ↪ Recent improvements in CEM clustering
- ↪ Track and unbiased Lshr available to Level-3 (4.6.0)

- Areas of activity

- ↪ Tower-by-tower to database (Olga, Beate, Dustin)
- ↪ Streaming speed of CESQ, PESQ (Costas, Michael)
- ↪ Generator-level JetClu (Jean-Francois)
- ↪ Beam offset corrected  $\cancel{E}_T$  (McQueen)

# Calorimetry



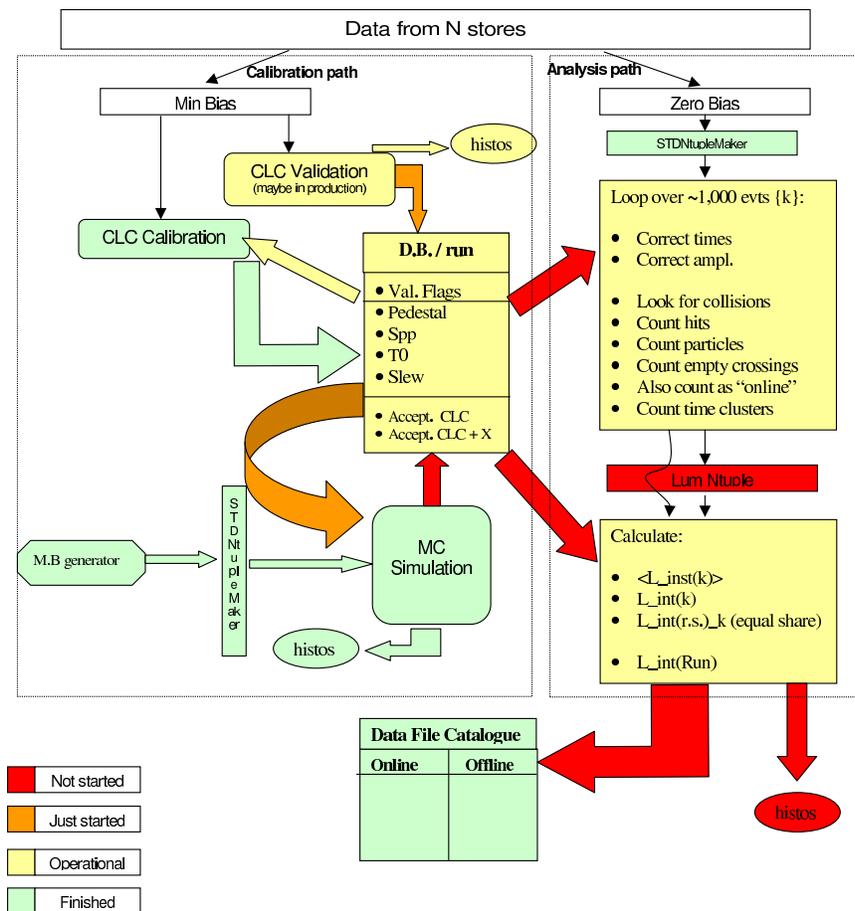
## CLC Offline: Lumi Calc

### Major Issue is Getting Lumi Calc into Offline Reconstruction

- Calibration path
  - ↪ Validation of data in Zero bias/Min. bias data set (gtst3/gcrs0)
  - ↪ Calibration:
    - ★ Amplitude (pedestals, gain — Single Particle Peak calibration)
    - ★ Timing (T0, slewing, stretching coeff.)
  - ↪ Calculation of CLC acceptance (MC)
- Analysis path
  - ↪ Reconstruction of raw data:
    - ★ Amplitude in number of SPP
    - ★ Timing correction, fine timing measurements ( $\sigma_T \sim 100$  ps)
  - ↪ Luminosity calculation:
    - ★ Counting of hits
    - ★ Counting of particles
    - ★ ...
  - ↪ Filling of DFC

## CLC (cont'd)

### Plan for automating the Offline Luminosity Reconstruction



- All critical modules are operational
- We continue debugging DB interface
- In 2–3 weeks we will be ready for the complete test
- Job-level, DHOutput tool for lumi sum (Ratnikov)
- Module level tools for inst. lumi, lumi by run in progress

# TOF

- Stable detector operation - all channels read out
- First calibration constants in database
  - ↪ *average* timing resolution  $\sim 110$  ps
  - ↪ Design goal was “100 ps”
- Use TOF in your analysis with version 4.5.0 or later...
  - ↪ Getting started:

[http://fcdfhome.fnal.gov/usr/mulhearn/tof\\_howto/index.shtml](http://fcdfhome.fnal.gov/usr/mulhearn/tof_howto/index.shtml)

- Version 4.6.0 will be the preferred release:
  - ↪ extra user-friendly interface features
  - ↪ *Access TOF information with less than 10 lines of code!!!*
  - ↪ support for non-track based timing queries:  
“Given  $z$ , what is  $T$ ?”
- *Caveats:*
  - ↪ Reconstruction/calibration not fully certified yet
  - ↪ Achieves  $2\sigma$   $K$ - $\pi$  separation for  $p < 0.98$  GeV/ $c$
  - ↪ Efficiency is low:  $\sim 40\%$ , part from  $\eta$  coverage, part from TOF-track matching criteria
- Goals for 4.7.0:
  - ↪ Fully certified calibration implementation
  - ↪ improved track matching algorithm for increased efficiency

## Simulation

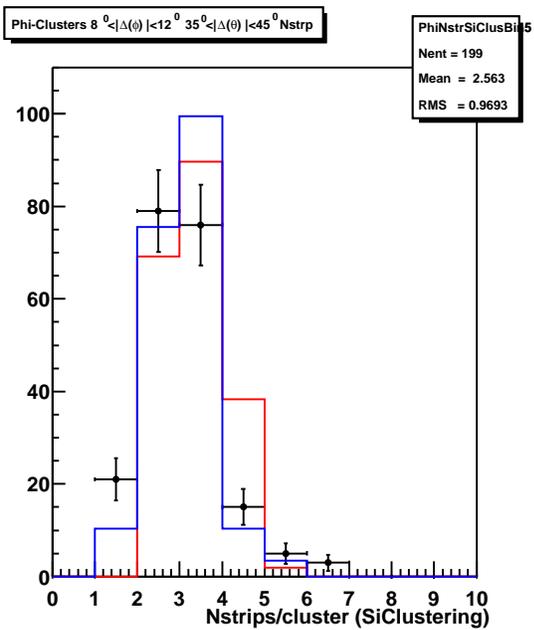
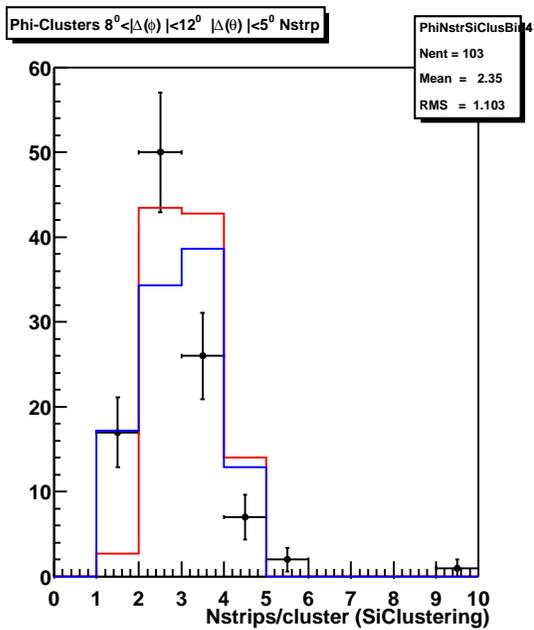
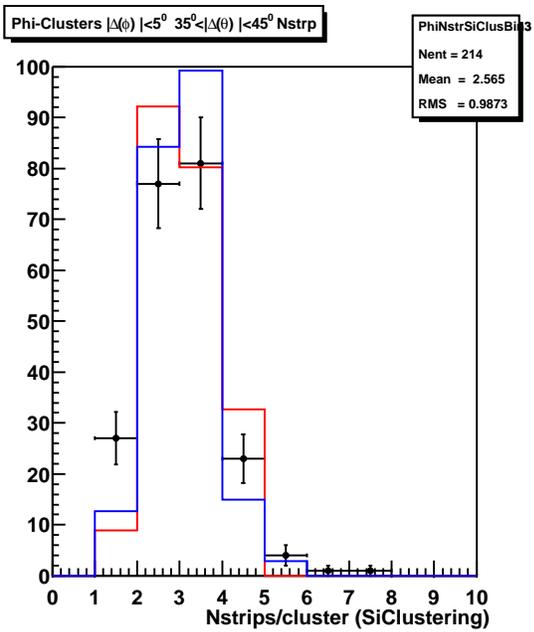
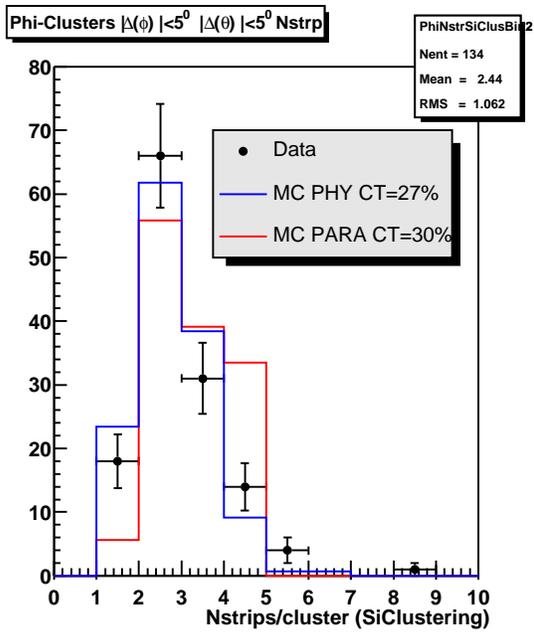
- Executive summary
  - cdfSim provides a working simulation of all subdetectors. Tuning and comparison with Run 2 data is primary focus.*
- Generators
  - ↪ Pythia 6.2 (Les Houches Accord #1 compliant)
  - ↪ Standalone generator interface *via* HEPEVT available
  - ↪ Progress on MLM  $V$ +jets, Grappa
- Unified random numbers (CLHEP) w/ save&restore
- First draft validation plots for cdfSim

## Simulation (cont'd)

- SVX
  - ↪ Charge deposition tuning in progress  
PLOT
  - ↪ Dead channels from Calib DB
  - ↪ Noise, misalignment under development
- COT: dE/dx implemented, tuning
- Calorimeter tuning
  - ↪ Testbeam tuning largely complete
  - ↪ Low  $p_t$  hadron tuning finished  
PLOT
- Muons: central OK, adding noise, IMU functioning
- Problems
  - ↪ cdfSim is slow (30 GHz-sec/ttbar event)
  - ↪ Pythia+production crashes (problem for farms)
  - ↪ Passive material 20% off from conversion maps?

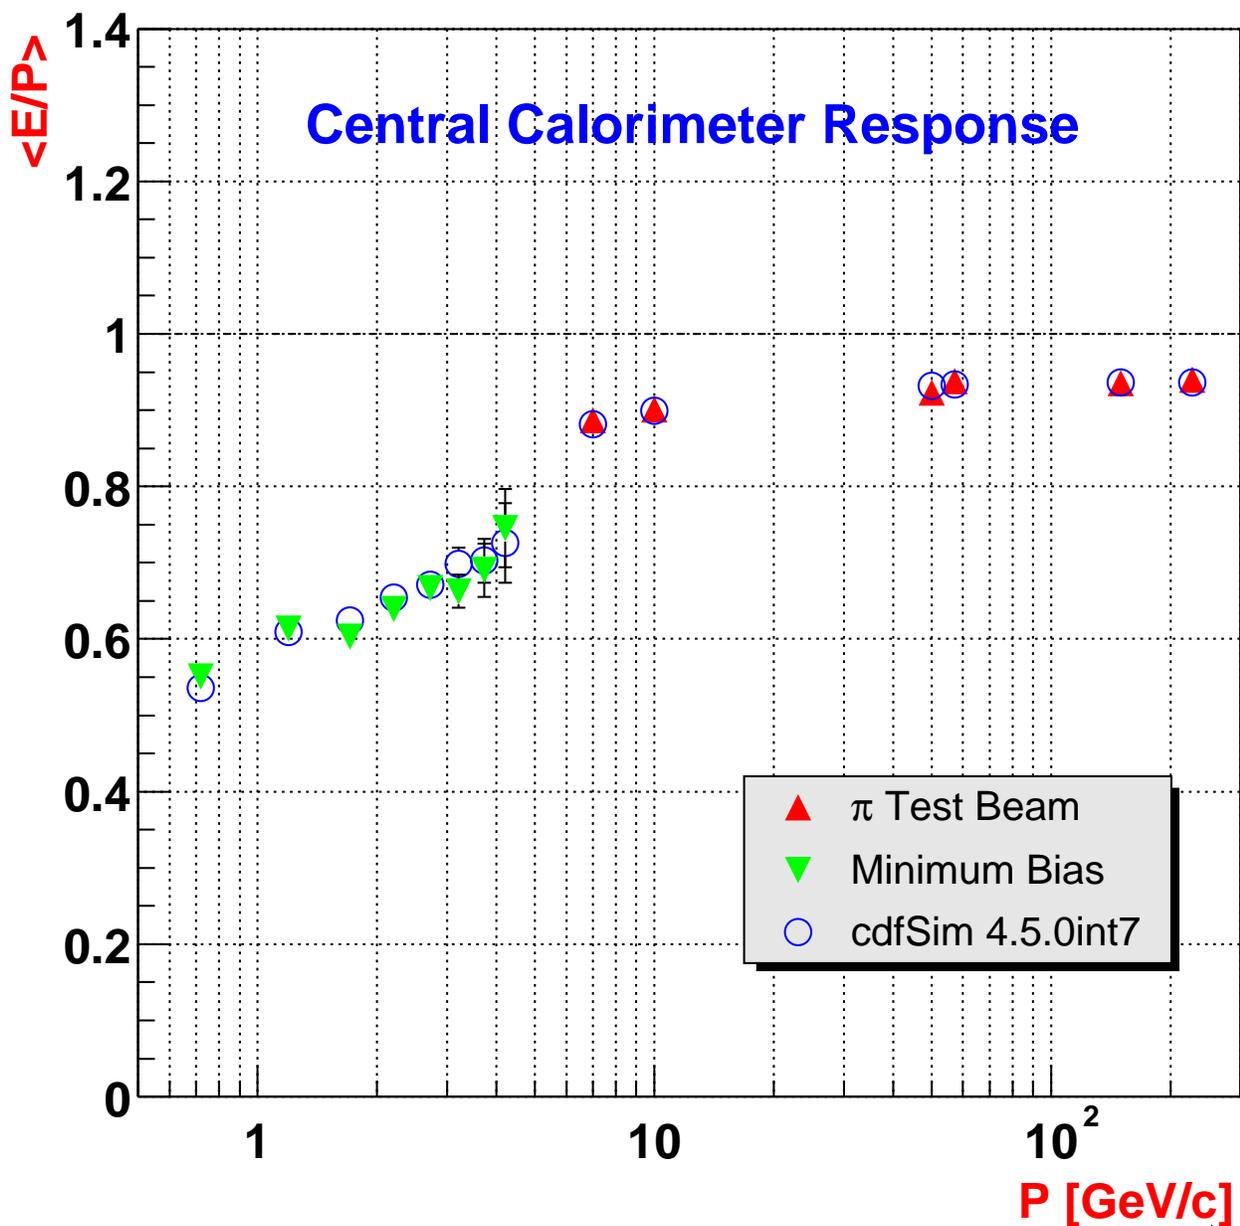
# Simulation (cont'd)

**DATA (DPS ON)137416,137484-137485\_bX\_al196fit\_I2 vs MC 1MU (PARA PHYSICAL CDM)**



*Tuning the physical and parametric charge deposition models*

## Simulation (cont'd)



Tuning Central Calorimeter to Min Bias and Testbeam  $\pi^\pm$

## Current and Pending Releases

### 4.5.2 Currently (just) running in production (reprocessing B, G, J)

- Tracking bug fixes applied to 4.5.0

### 4.6.0 Frozen May 20th, debugging

- Framework support for multi-branch root, dCache
- Tracking:
  - ↪ defTracks pruned, ordered collection
  - ↪ Significant Si geometry fixes
  - ↪ Improved OI speed, efficiency, fit
- June 15th is goal date for “farm ready” 4.6.0

### 4.6.x for tracking improvements

- More Si stereo geometry fixes
- defTracks cuts
- Schedule driven by this work (weeks?)

*Best case: next full production pass is 4.6.x*

## Permit Me a Few Reflections

*This has been an eventful two years for the offline*

- We discarded our central systems and **started anew**
  - ↪ Old model (1997) failed to provide adequate resources; was unsupportable; had major hardware problems.
  - ↪ Why has **new model** been so successful?
    - ★ Sufficient **operational cushion**
    - ★ **Network-based, decoupled** services
    - ★ Support from **CD and collaborators**
      - We have been most successful when both elements are present  
(Level-3, CS/L, now DH, CAF)
      - *It is an important lesson to remember*
- The *second* most awful task I had in this period was managing woefully inadequate resources of the legacy system
  - ↪ We have turned the corner on this  
(Enstore data access, farms reprocessing, soon CAF)

## Permit Me (cont'd)

- The *worst* part of this job has been functioning with inadequate personnel for infrastructure and integration
  - ↪ Software I&I continue to be schedule-drivers and determinants of user quality-of-service
  - ↪ Yet as a collaboration, we drastically under-support these areas
    - ★ Too much reliance on computing professionals for this work
    - ★ Lack of physicist involvement even breeds contempt for those who take on this work, which is very unfortunate
    - ★ These are technically challenging projects. CDF cannot do physics without successfully completing them.
- My advice (for whatever it is worth)
  - ↪ Invest more physicist effort in software I&I
  - ↪ Tighten the coupling between physics groups, whose analyses set the agenda, and software I&I

## Conclusions

- Enormous progress on DH, CAF (next talk) for central systems
- Reconstruction software is, largely, on the right track
- Many hard problems remain in our software infrastructure and in integration