Triggering at Hadron Colliders

Gustaaf Brooijmans

Columbia University
In the City of New York

Hadron Collider Summer School
Fermilab, August 2006

- Introduction
- Trigger architecture(s)
  - Hardware examples
- Commissioning
  - Hardware
  - Physics
- Trigger and analysis
- Online data quality

Introduction

The Basics

- Total cross-sections are large:
  - ~80 mb at $\sqrt{s}$ = 1.8 TeV
  - @10^{32}, that's 8 MHz!
- "Interesting" cross-sections (say W-> e) are much smaller:
  - $\mathcal{O}$(few nb)
  - @10^{32}, that's < 1 Hz
    - At 10^{34} at the LHC it becomes $\mathcal{O}$(10 Hz)!

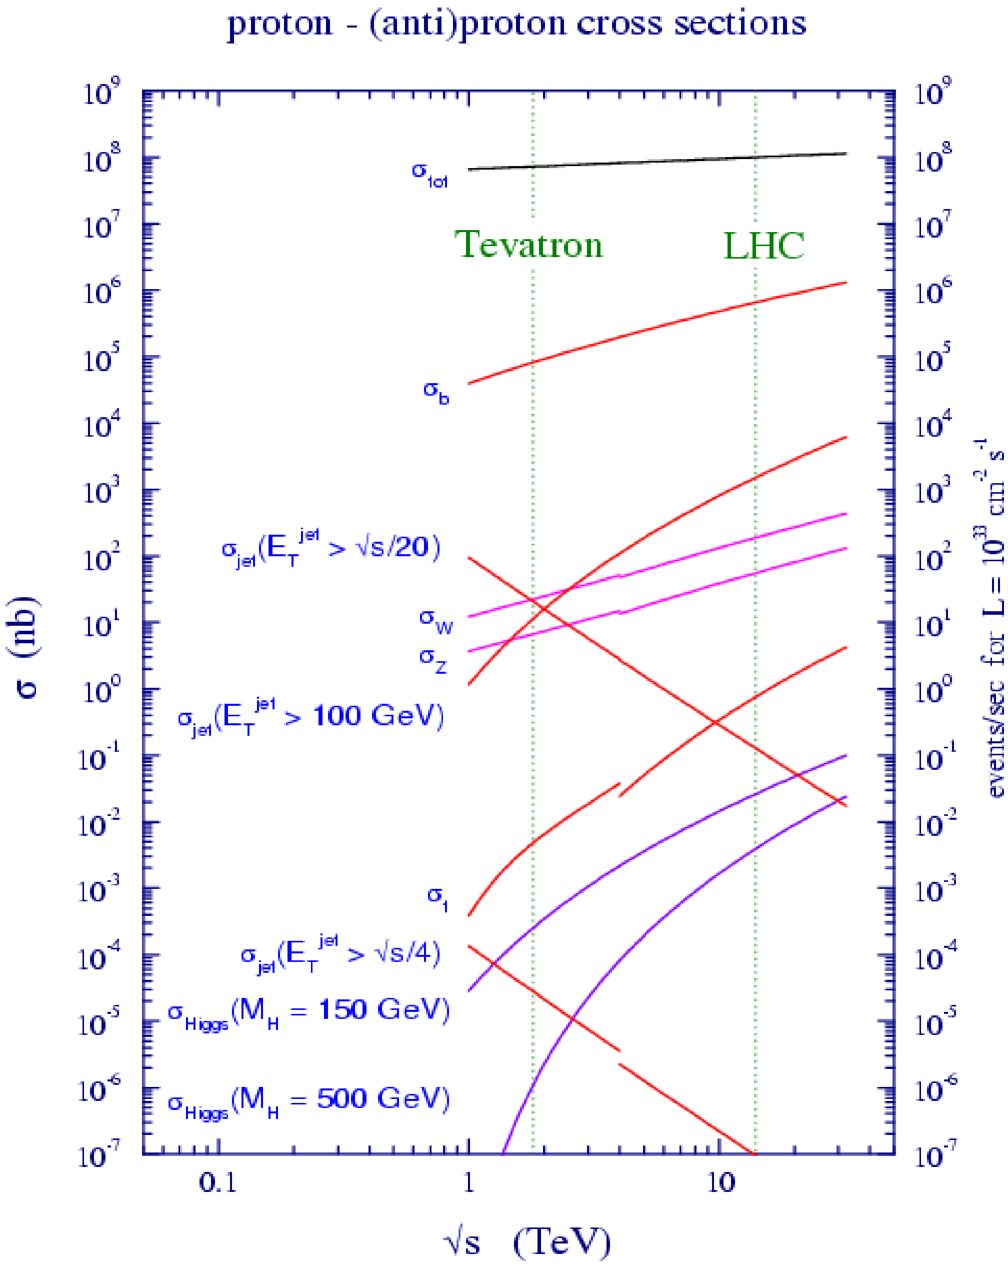

Basics II

- Trigger goal:
  - "To select interesting events for offline analysis"...
  - ... while minimizing deadtime!
- "Interesting" is a relative concept:
  - Depends on physics priorities (need for compromise in multi-purpose experiments)
  - Only interesting if event passes offline cuts!
  - Includes events needed to validate analysis
    - Determination of efficiencies
    - Control samples
    - ... (more later)

Basics III

- During decision-making process, data needs to be "stored"
  - Slower process ("latency") means "deeper memory"
  - There is a "traditional architecture" (CDF, DØ, ATLAS, CMS....)
  - Rapid evolution in technology opens door to new ideas however (BTeV, CKM, to a lesser extent LHCb?)
- But, all other things being equal, faster processing means less rejection and therefore more output bandwith (and storage and ...)

Trigger Architectures Hardware

Implementation

- Physics goal and technology dependent
- At hadron colliders, two types of experiments
  - Multi-purpose (CDF, DØ, ATLAS, CMS)
  - Dedicated (BTeV, LHCb, ALICE)
- Different technological epochs
  - CDF and DØ designs ~predate cheap Gb ethernet
  - Even LHC experiments use – by now – older technology
  - Always at the forefront during design, antiquated during construction...

Dataflow Arguments

- Tevatron: "precision" raw data ~200 kB/evt (zero suppressed and compressed)
  - L1 input if used that: > 3 Tbps
    - Need to slim and factorize for processing
    - But sometimes also duplicate....
  - To tape (100 Hz): ~20 MB/s
- LHC: ~1 MB/evts
  - L1 input if used that: > 300 Tbps
  - To tape (200 Hz): ~200 MB/s
- So, trigger is not just a physics argument

“Traditional” Architecture

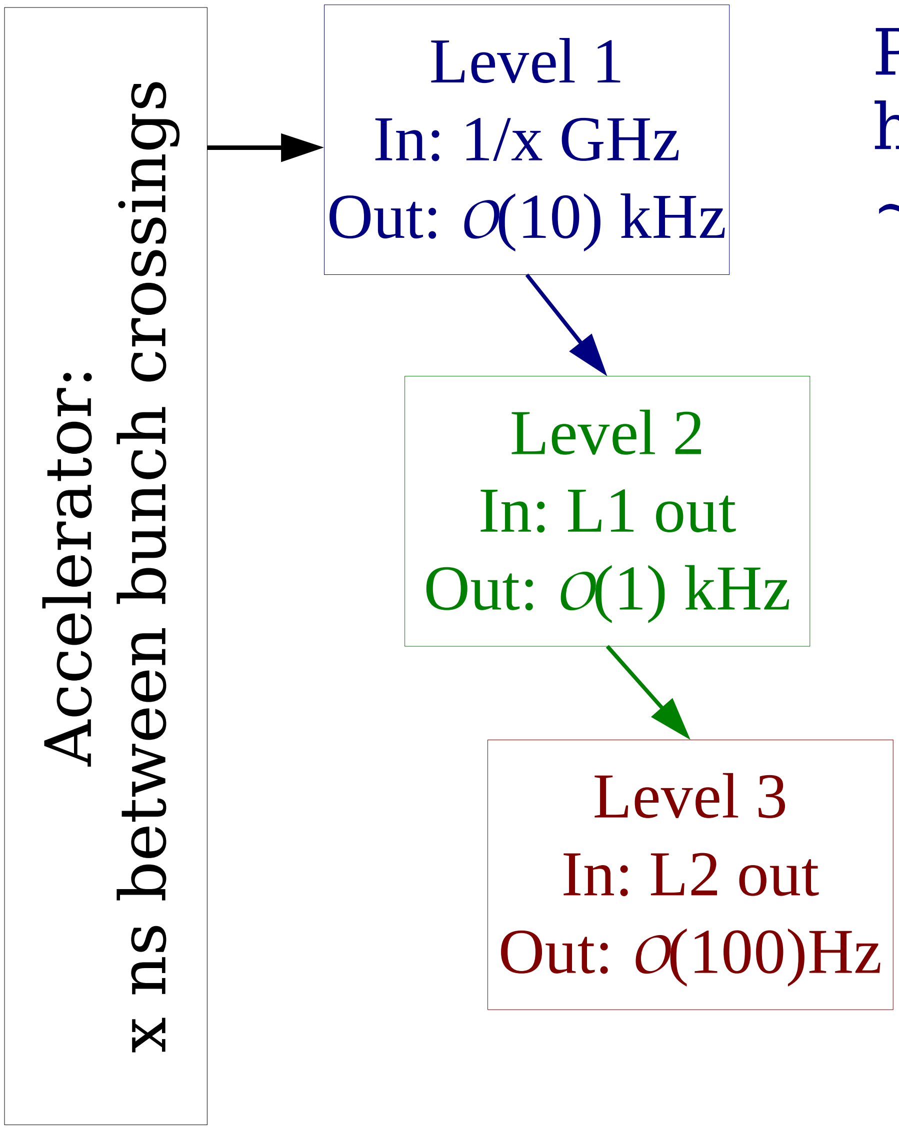

Pipelined (often deadtimeless), hardware only, coarse readout, ~few μs latency

Hardware/Software mix, L1 inputs, ~100 μs latency

CPU farm, access to full event information, $\mathcal{O}(1)$s/event

“Traditional” Elements

- Level 1 uses dedicated hardware, separate signals, per-subdetector “decision”
  - ASICs and FPGAs
- Level 2 uses dedicated hardware for “data preparation”, then CPUs for combination and decision
- Level 3 uses commercial CPUs
  - Difficulty is getting all of an event to a specific node, various approaches
    - “Concentrator(s)” -> bottleneck, single point of failure
    - “Fully distributed”

DØ

DØ Level 1 and 2

- Complex system uses tracking, calorimetry and muon system
  - … and matches between them!

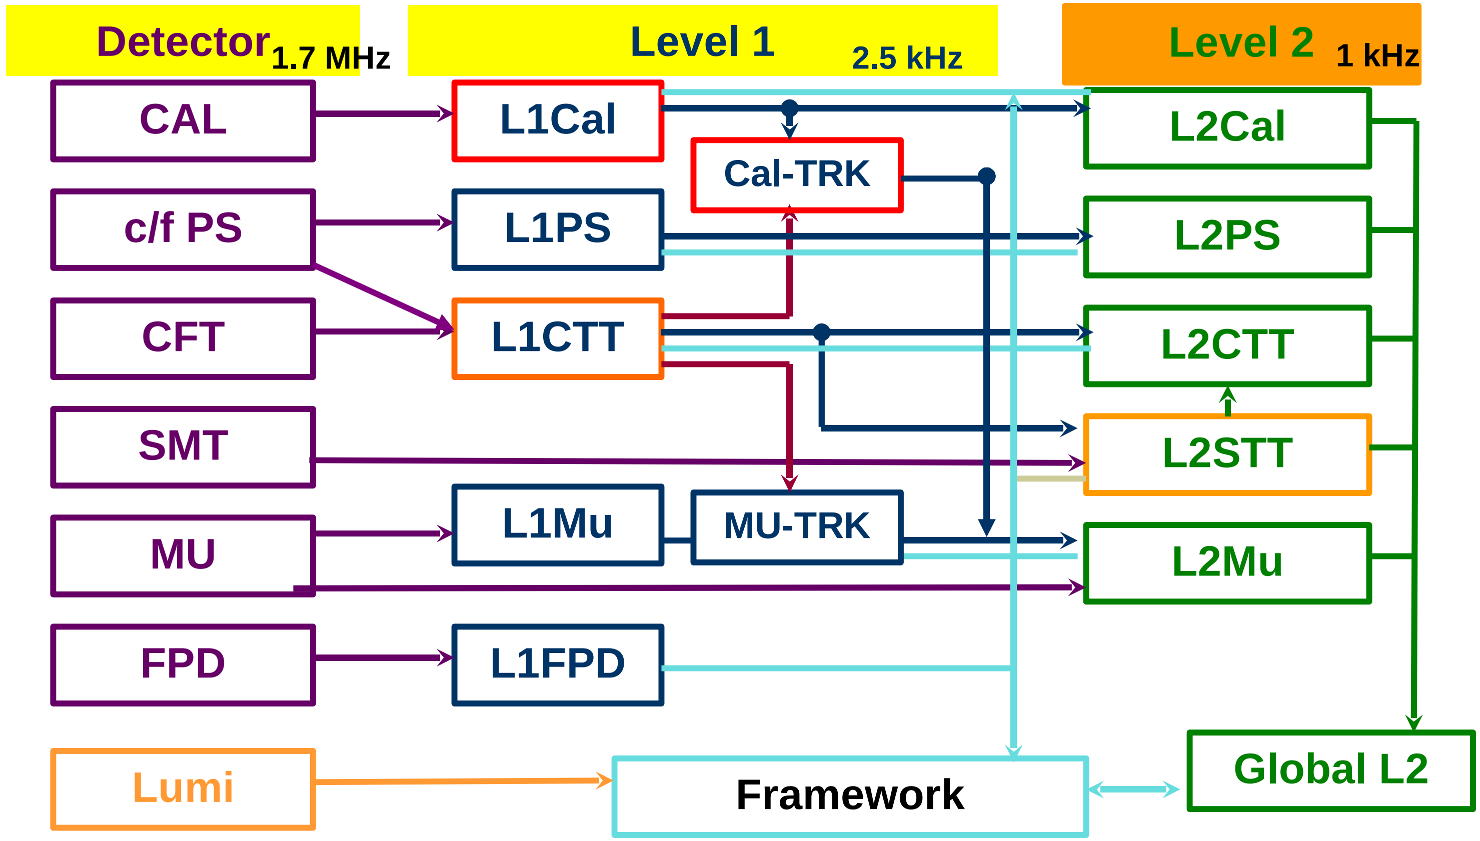

DØ Level 1 Elements

- Calorimeter: full readout has ~10 layers of cells of 0.1 x 0.1 in η x ϕ. The L1 trigger gets ~independent analog signals of 0.2 x 0.2 trigger "towers", both EM and hadronic energies.
- Track trigger (CTT): uses signals from individual scintillating fibers, compares with "lookup table" of preprogrammed track patterns
- Muon: uses both scintillator and wire chamber coincidences with various combinations possible, and track p^T estimates from match with CTT candidates

Hardware Example: L1CAL

- Inputs are energies in towers of 0.2 x 0.2 in $\eta \times \phi$.
- Ok for electrons – typical size is R < ~0.2
  - Origin of the choice of 0.2 x 0.2
  - Still lose electrons hitting far from tower center
- Not so good for jets – typical size is R ~ 0.6
  - Single tower threshold of 7 GeV is only fully efficient at ~50 GeV!
  - Sharpen turn-on substantially by clustering
    - "Poor man's": just require more towers above threshold
    - "Rich man's": develop clustering in FPGAs

“Sliding Windows”

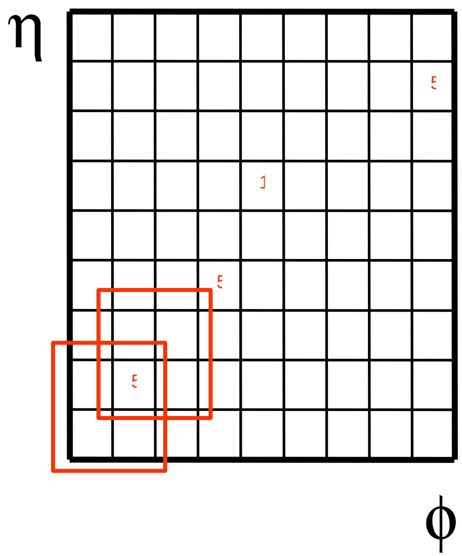

- Basically a search for local maxima by sliding a window on a grid of trigger towers
  - Many tuneable parameters
    - Size of window
    - Minimum separation between local maxima
    - Number of towers around maximum to consider in object
  - Close collaboration between physicists and engineers
  - Substantial dataflow issues (Tbps):
    - Neighboring towers can be “far away” due to physical cabling

- Added benefits:
  - Also recover electron inefficiencies
    - Isolation!
    - Hadronic fraction
  - More power for taus:
    - Isolated, narrow jet

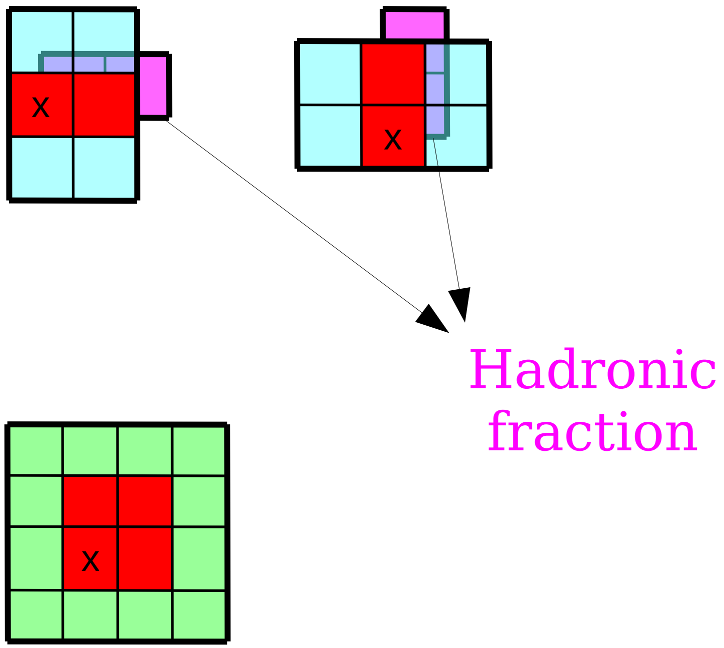

- Remember: in FPGAs, sums, comparisons easy; multiplication, division, “if – then” expensive

Hardware Example: L1CTT

- Scintillating fiber tracker, read out using VLPCs

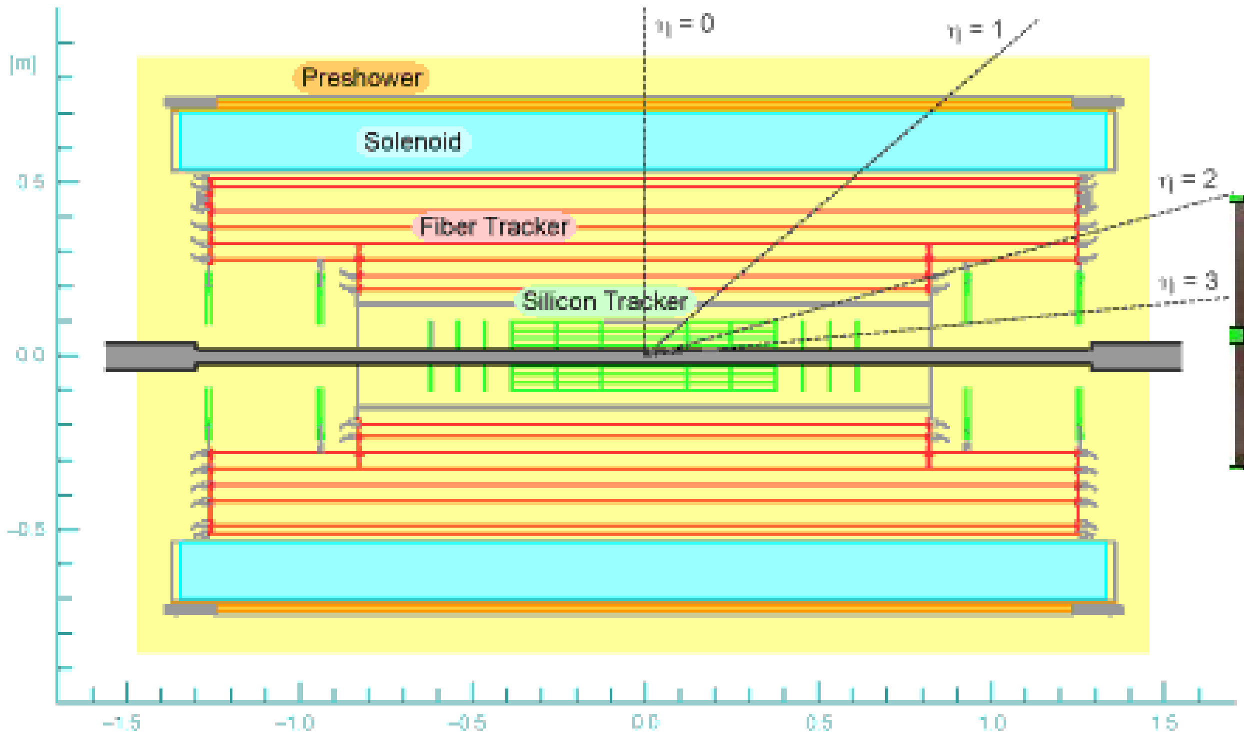

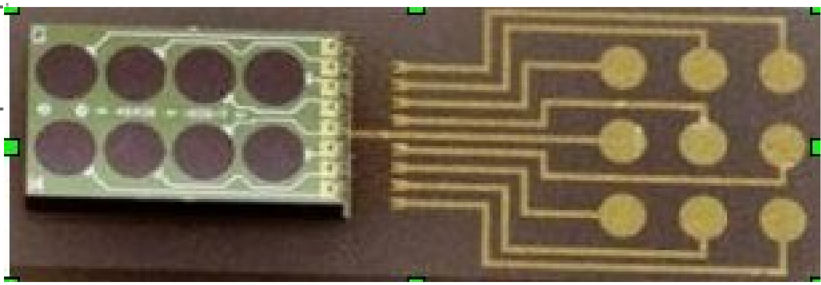

  - VPLC signals split, “trigger path” goes to discriminators
  - Compare hit pattern with pre-programmed track patterns for different p^T ranges

- Of course, preprogrammed track equations need to factor in “as-built” detector
  - Sensitive to alignment effects
- Beamspot tolerance is typically O(1 mm)
  - For Level 2 impact parameter, need to feed beamspot to the system
- Handle dead channels?
  - Loss in efficiency
  - “Turn on” dead channels

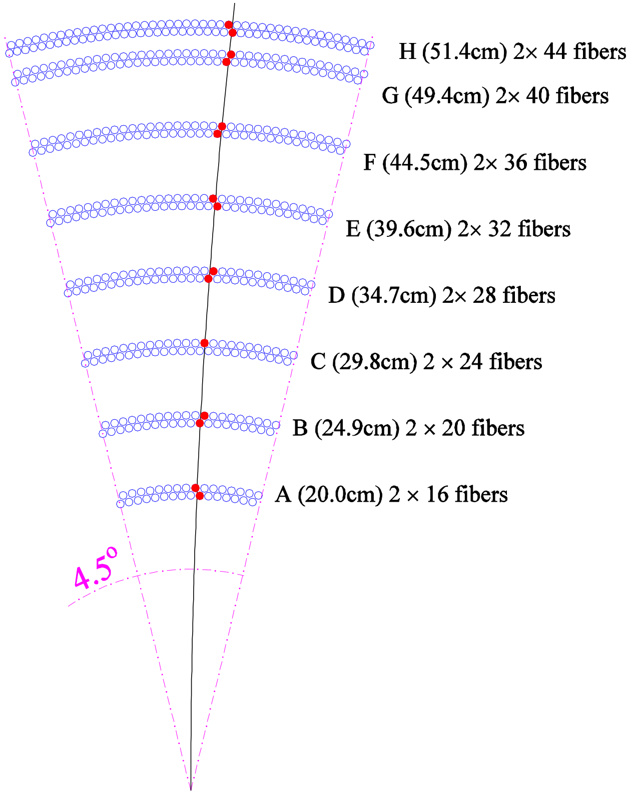

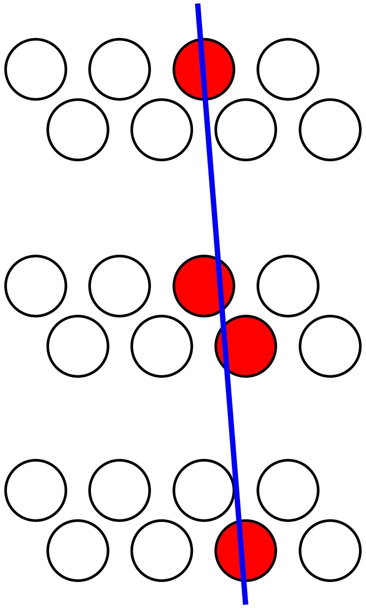

- Difficulty is usually getting the right data together
  - Grouping for detector readout usually not the grouping wanted for "reconstruction"
    - Not an issue offline, but means moving lots of data at trigger level
  - Boundaries between geographic regions particularly difficult
- Highly parallel activity (comparing many channels to preprogrammed patterns) is ideally suited for FPGAs
  - Modern Level 1 triggers rely heavily on FPGAs
  - Fast evolution of the technology opens new windows

DØ Level 3

- Commodity CPU farm with ~250 dual-CPU nodes
- "Fully Distributed" model:
  - Each readout crate has a Single Board Computer which reads the data over the VME bus and sends it via ethernet
  - For each L2 accept, a "routing master" decides which node will process an event (based on available buffers)
    - SBCs get "told" which node to send an event to, typically in packets of 10 events
  - Hardware "core" is a good quality, large bandwidth switch
  - Software core is the "routing master"

Logically

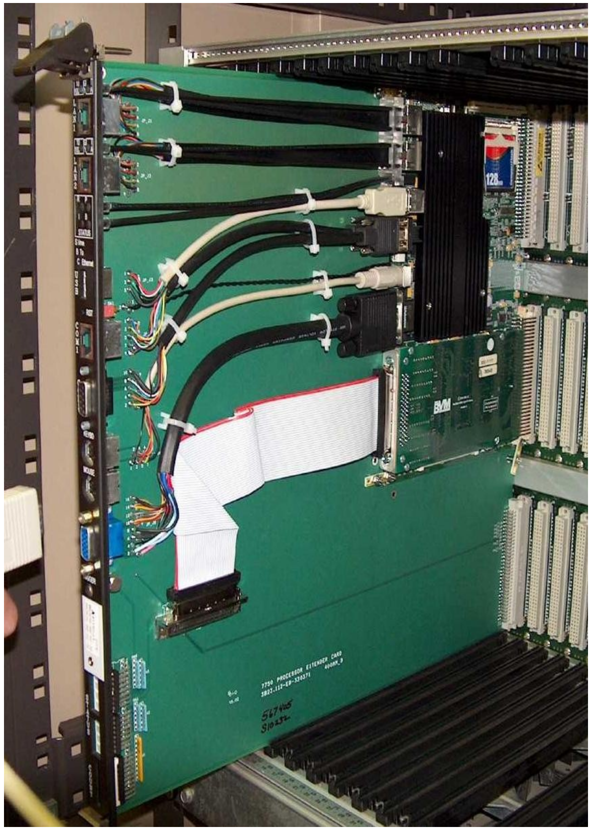

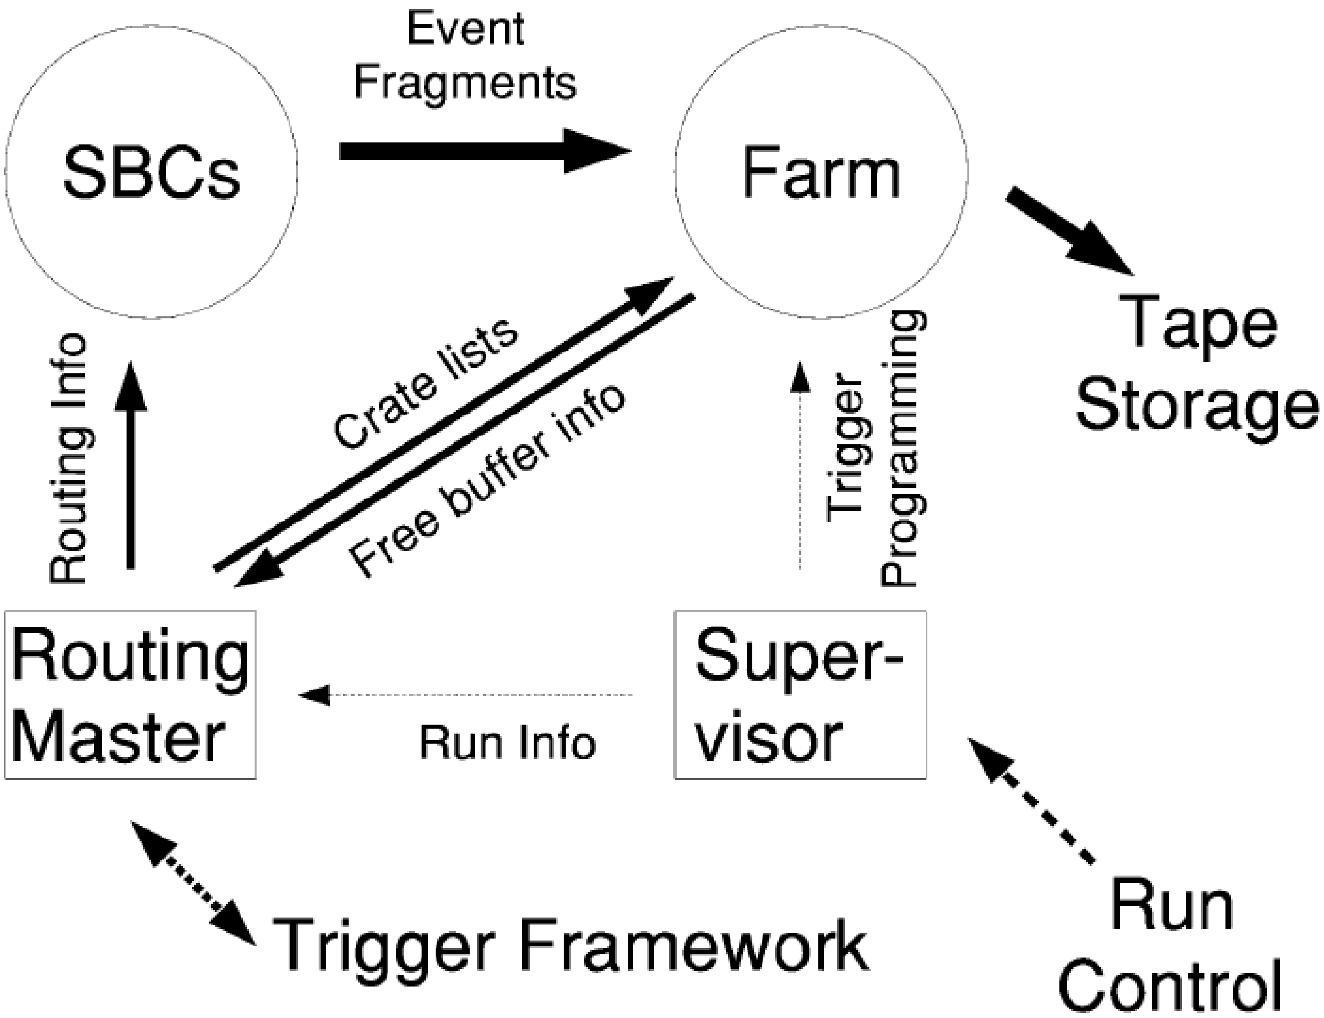
Event
Fragments
SBCs
Farm
Tape
Storage
Routing Info
Crate lists
Free buffer info
Trigger
Programming
Routing
Master
Run Info
Super-
visor
Trigger Framework
Run
Control

Physically

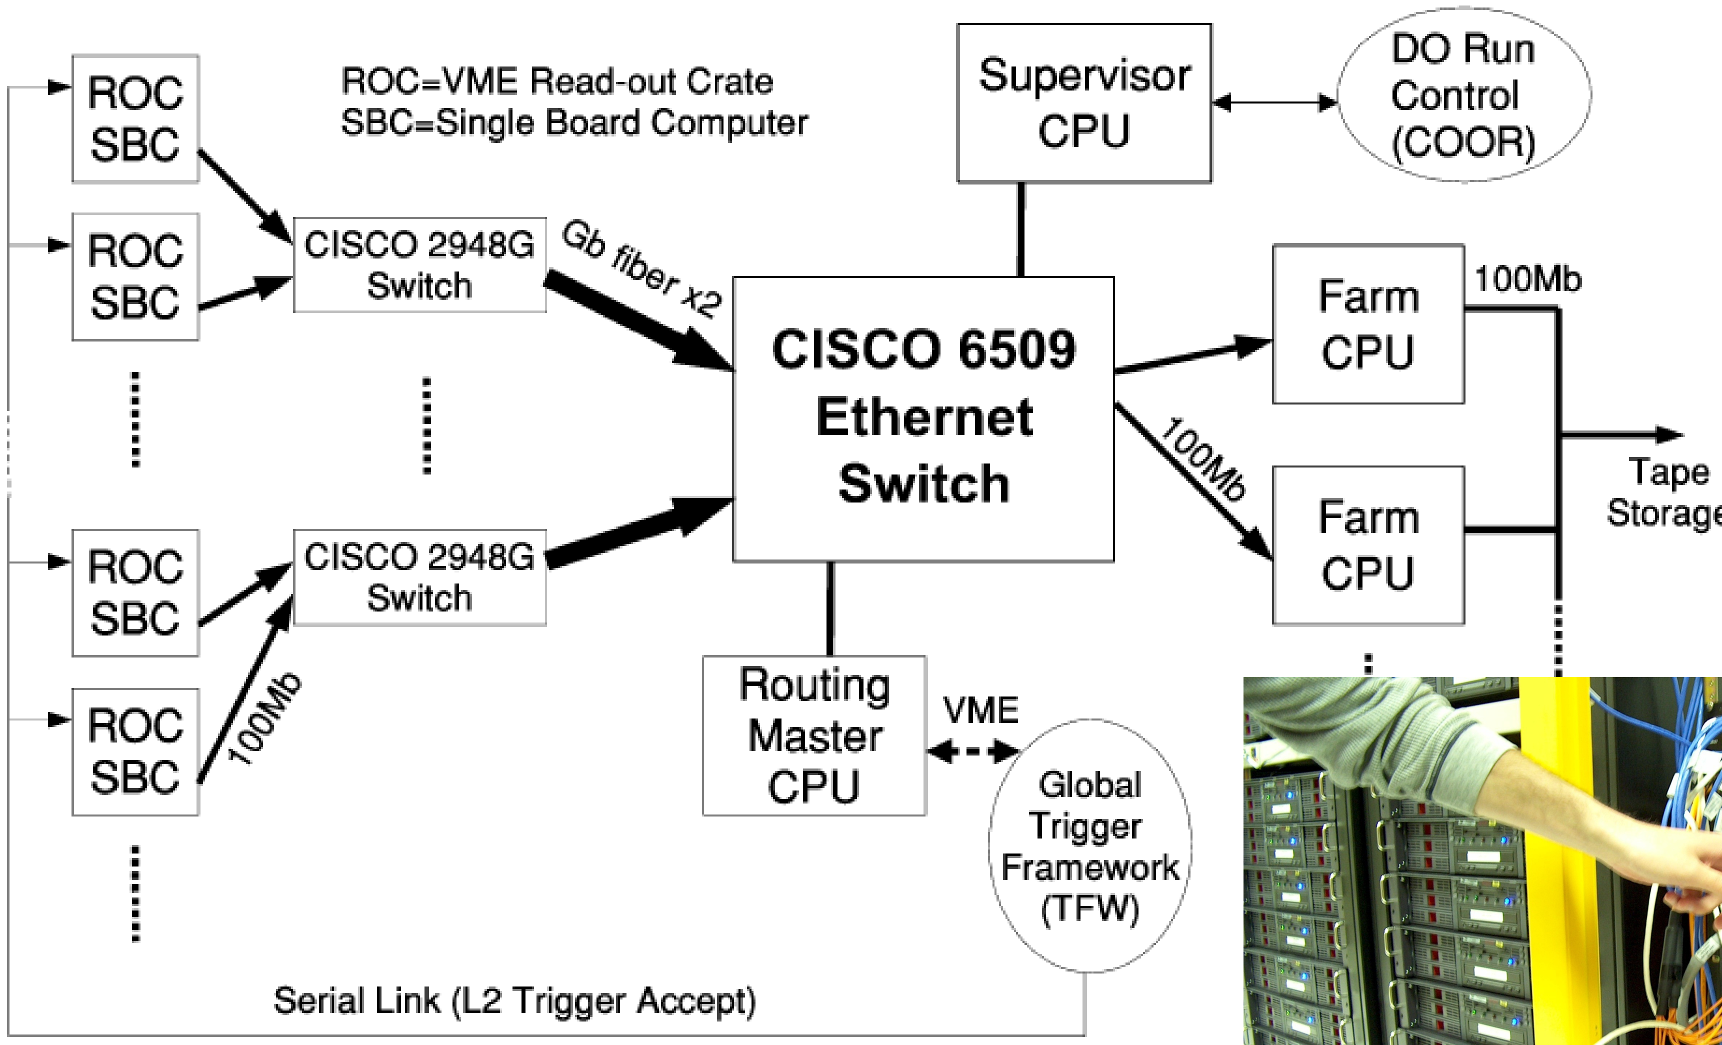
ROC=VME Read-out Crate
SBC=Single Board Computer
ROC SBC
ROC SBC
ROC SBC
ROC SBC
CISCO 2948G Switch
CISCO 2948G Switch
Gb fiber x2
100Mb
Supervisor CPU
DO Run Control (COOR)
CISCO 6509 Ethernet Switch
Farm CPU
Farm CPU
100Mb
100Mb
Tape Storage
Routing Master CPU
VME
Global Trigger Framework (TFW)
Serial Link (L2 Trigger Accept)

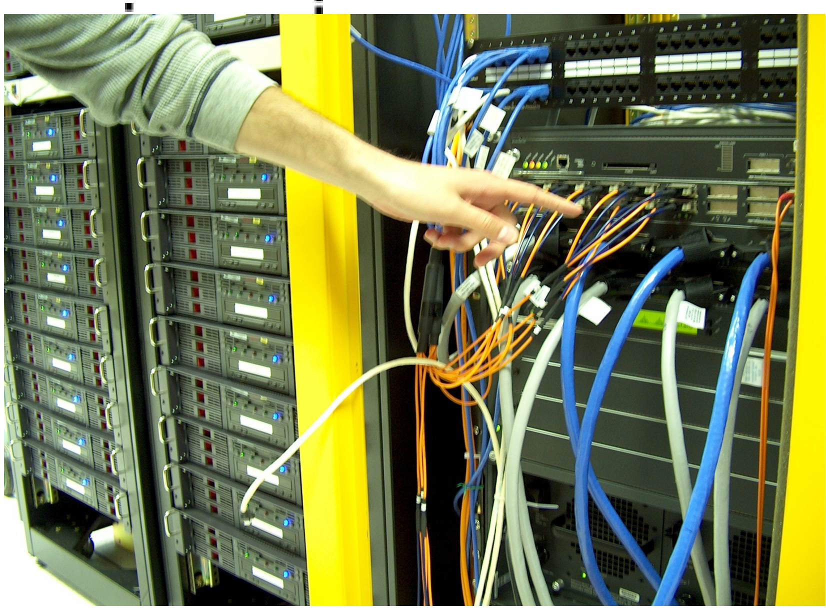

DØ Latencies and Buffering

- Level 1 latency determined by depth of SVX IIe pipeline (32, but running in 132 ns intervals -> ~4 μs)
  - That's the time to get the signals out, process them, make a decision, and send the L1 accept info back to the front-ends
- Level 1 accept rate determined by deadtime
  - Each L1 accept leads to loss of the other events in the pipeline, + need to transmit the data
    - For deadtime < 5%, max L1 accept rate is ~1.5 kHz
    - Level 2 needs to issue decisions at that rate
- Level 2 accept rate limited by VMEbus bandwidth (-> event size): ~850 Hz at high luminosity

- Level 3 accept rate limited by offline resources

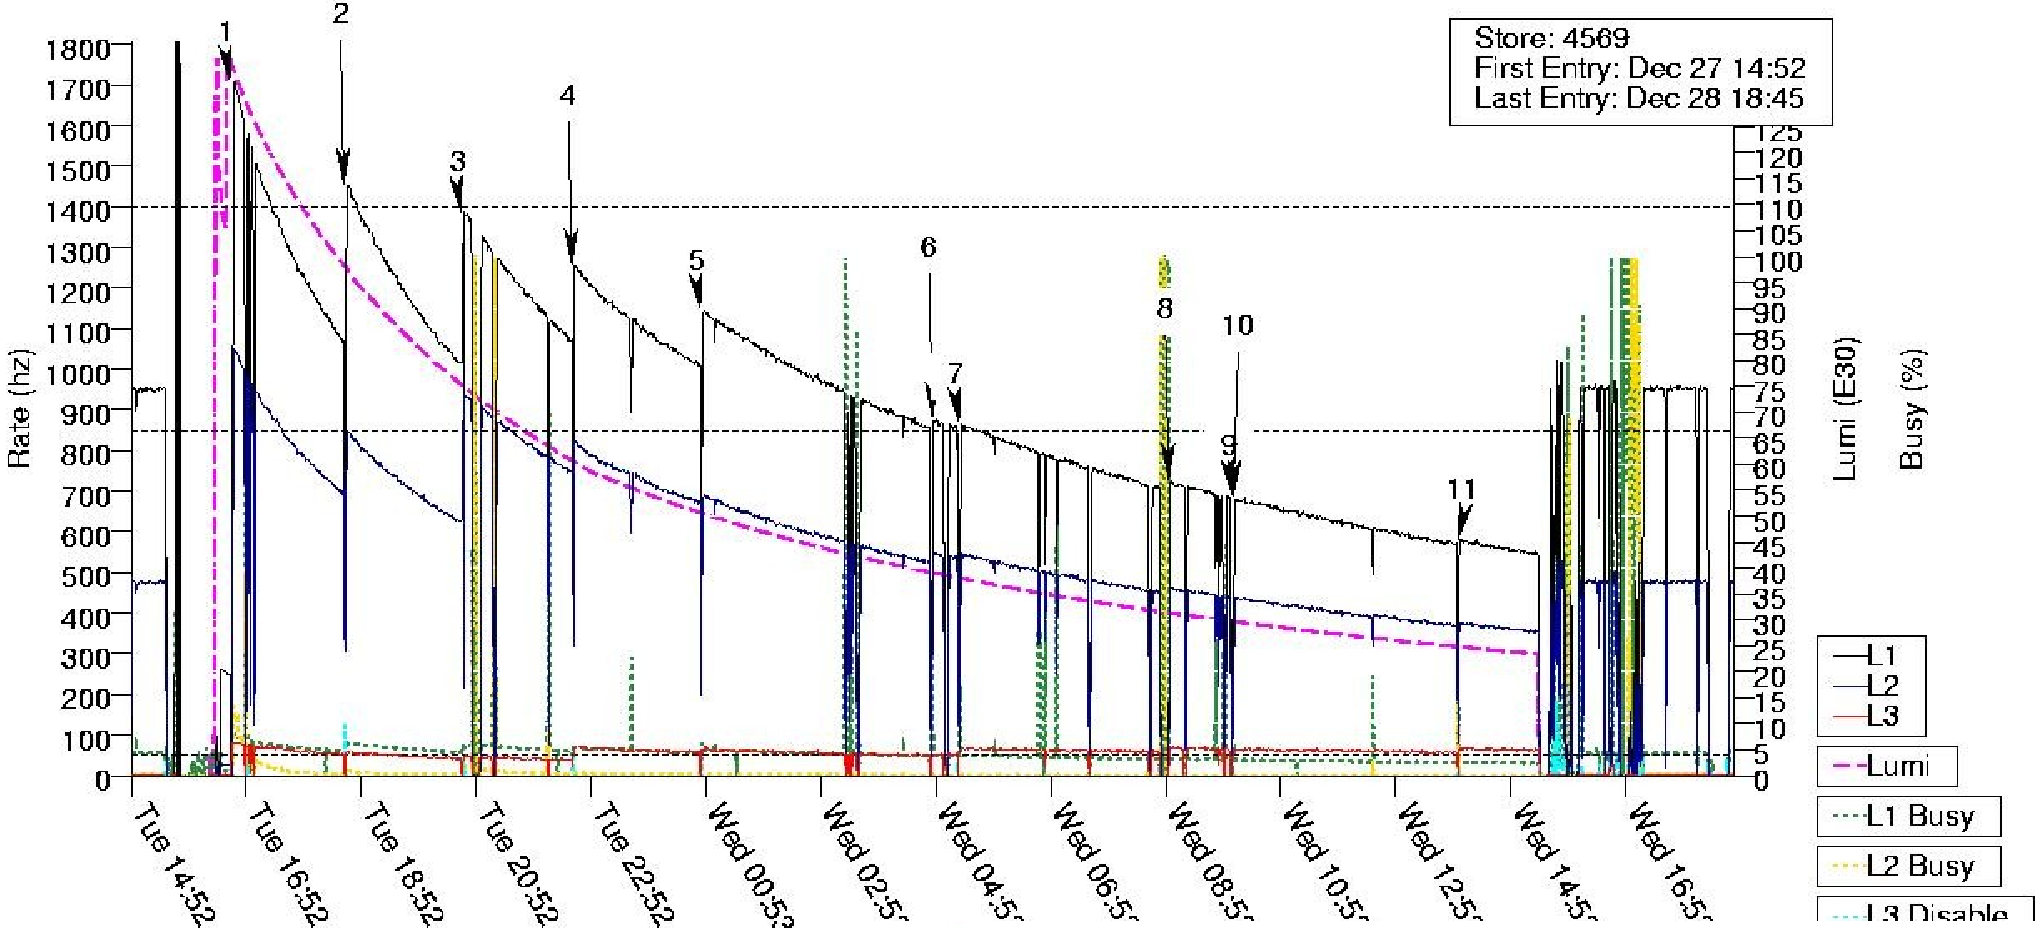

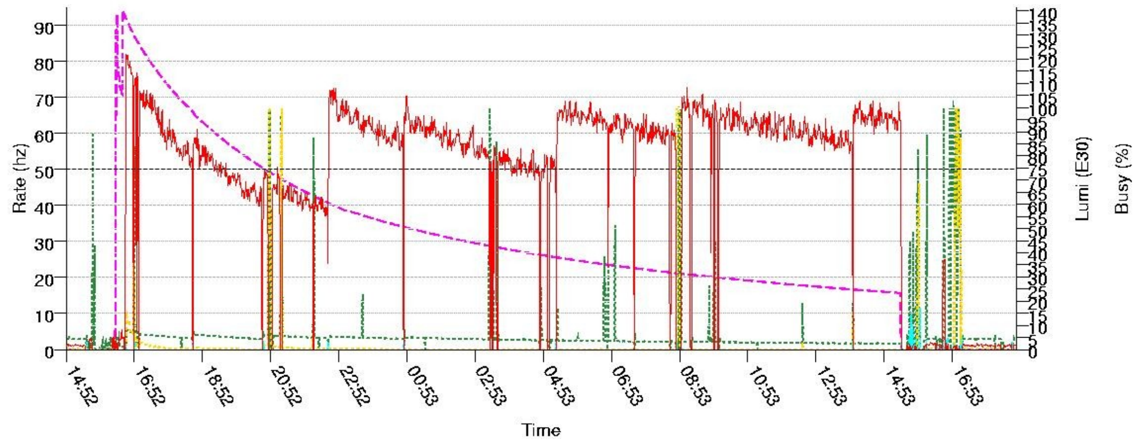

DØ Trigger Logic

- A level 1 trigger is a logical AND of multiple requirements
  - There are 256 possible “requirements” called “AND/OR terms”
    - The terms are hardwired (literally), and correspond to things like “2 EM towers above 6 GeV”
  - Up to 128 Level 1 triggers are allowed
  - Exact integrated luminosity can be determined for 8 groups of L1 triggers
    - Because need to keep track of deadtime

- Each Level 2 trigger hangs off a single Level 1 trigger
  - Generic constraint: understanding the system (and efficiencies) becomes difficult otherwise
  - Similarly, each Level 3 trigger hangs off a single Level 1 + Level 2 trigger
  - But, a single Level 1 trigger (max 128) maps to many Level 3 triggers (no hard max)
- Prescales only allowed at Level 1 (luminosity accounting) – leads to replication of Level 1 conditions

- Example of a Level 1 description:
  - TTK(2,5) x TTK(1,10) x TIS(1,5) x CER(1,C,6): at least two tracks with p^T > 5 GeV AND at least one track with p^T > 10 GeV AND at least one isolated track with p^T > 5 GeV AND at least one EM tower in the central calorimeter with E_T > 6 GeV
    - Not the simplest, but also not the most complex example
- Corresponding Level 2 trigger has:
  - L2CALTRK(1, 6, 5, TIS) x L2JETTRK(2, 5, 5, TTK): match between a 6 GeV EM tower and an isolated track with p^T > 5 GeV AND 2 jets with E_T > 5 GeV each matched to a track with p^T > 5 GeV
    - Note: Level 2 uses “trigger data”, not full data

- This L1/L2 condition then has multiple associated L3 conditions (not written out here), and the triggers are
  - E31T_SHT102TAU10
    - L3: Tight 10 GeV electron + 2 10 GeV taus (NN algo)
  - E31T_SHT15_M25
    - L3: Tight 15 GeV electron + 25 GeV of MET
  - E31T_SHT15_TK13
    - L3: Tight 15 GeV electron + 13 GeV track
  - E31T_2T5SH5
    - L3: 2 5 GeV electrons with loose shower shape and matched tracks
- Triggers for SUSY trileptons, W, Z, J/ψ, top, ...

CDF

CDF Trigger

- General architecture very similar to DØ
  - Three levels
  - Calorimeter, muon, and track triggers at L1
  - Level 2 is a mix of hardware and software
  - Level 3 is a PC farm (few 100 PCs)
  - Level 1 and 2 use "trigger data", Level 3 has full info
- Some important differences
  - SVT, using silicon info at Level 2 was part of the baseline design (came later in DØ):
    - Key in B-physics program
    - Led to a substantially different "rate architecture"

Silicon Vertex Trigger

- Take tracks found by Level 1 trigger
- Cluster "hits" in SVX
- Pattern recognition based on pre-programmed tracks
  - Manageable because work at coarser resolution than SVX
- Then use hits + track info to fit (linear approximation)

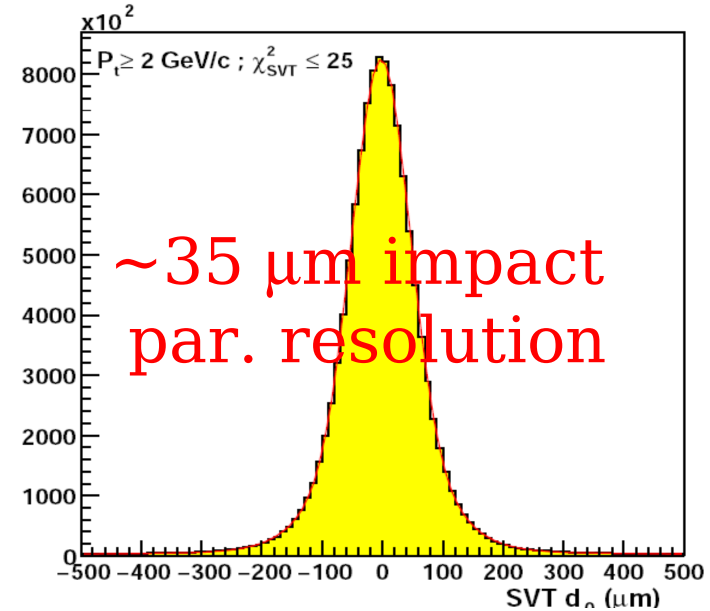

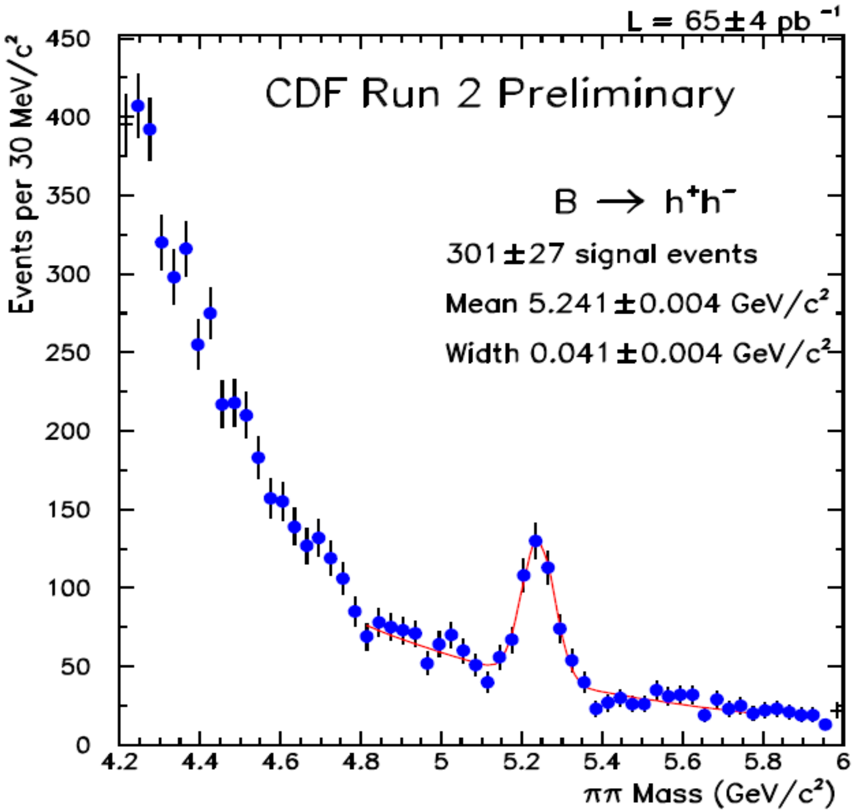

CDF Rates

- L1 accept up to 40 kHz – typical 20-35 kHz
  - SVX III pipeline not lost on L1 accept -> “deadtimeless”
  - SVX III pipeline is 42 cycles deep
- L2 accept rate up to 1 kHz (Run IIb) – 850 Hz achieved
  - Large fraction of L1 accepts are track triggers with large SVT rejection
- L3 accept rate 70-120 Hz
- Note: system does have deadtime, not from SVX III

CDF Features

- “Lumi Enables”: triggers turn on automatically below specified instantaneous luminosities
- Dynamic prescaling: prescales are varied automatically to keep rates within specific bands
  - Requires change of runs in DØ
- “Über prescaling”: L1 accept can only be issued if L2 buffer available
  - Used to fill up available bandwidth
- Multiplicity veto: for certain triggers, uses luminosity counters to veto events with many interactions to reduce fakes

ATLAS & CMS

- Some evolution from Tevatron:
  - Higher L1 bandwidth: 75-100 kHz
    - But 25 ns between bunch crossings vs 396 ns
  - Simpler L1: no track triggers
    - Expect 20 interactions per bunch crossing at design luminosity, and track triggers are highly sensitive to multiplicity...
  - Level 2 becomes a “front end” for Level 3 (now called the “event filter”), and together they form the High Level Trigger

“Region of Interest” and HLT

- “Level 2” choice at LHC different than Tevatron:
  - Tevatron: Level 2 relies heavily on L1 inputs, can refine decision somewhat
  - LHC: Level 2 gets “precision data”, but only for “Region of Interest”, i.e. around the object L1 triggered on
    - Expected to get large rejection in “short” decision time (~10 ms)
    - Only gets ~2% of data (for each L2 algorithm)
  - Event filter then looks at full event and gets ~1s
- Since both Level 2 and EF run on PC nodes, some flexibility in Level 2 performance
  - Provided you can get the data to the EF!

A Different Approach: BTeV

- Dedicated experiment to study B decays

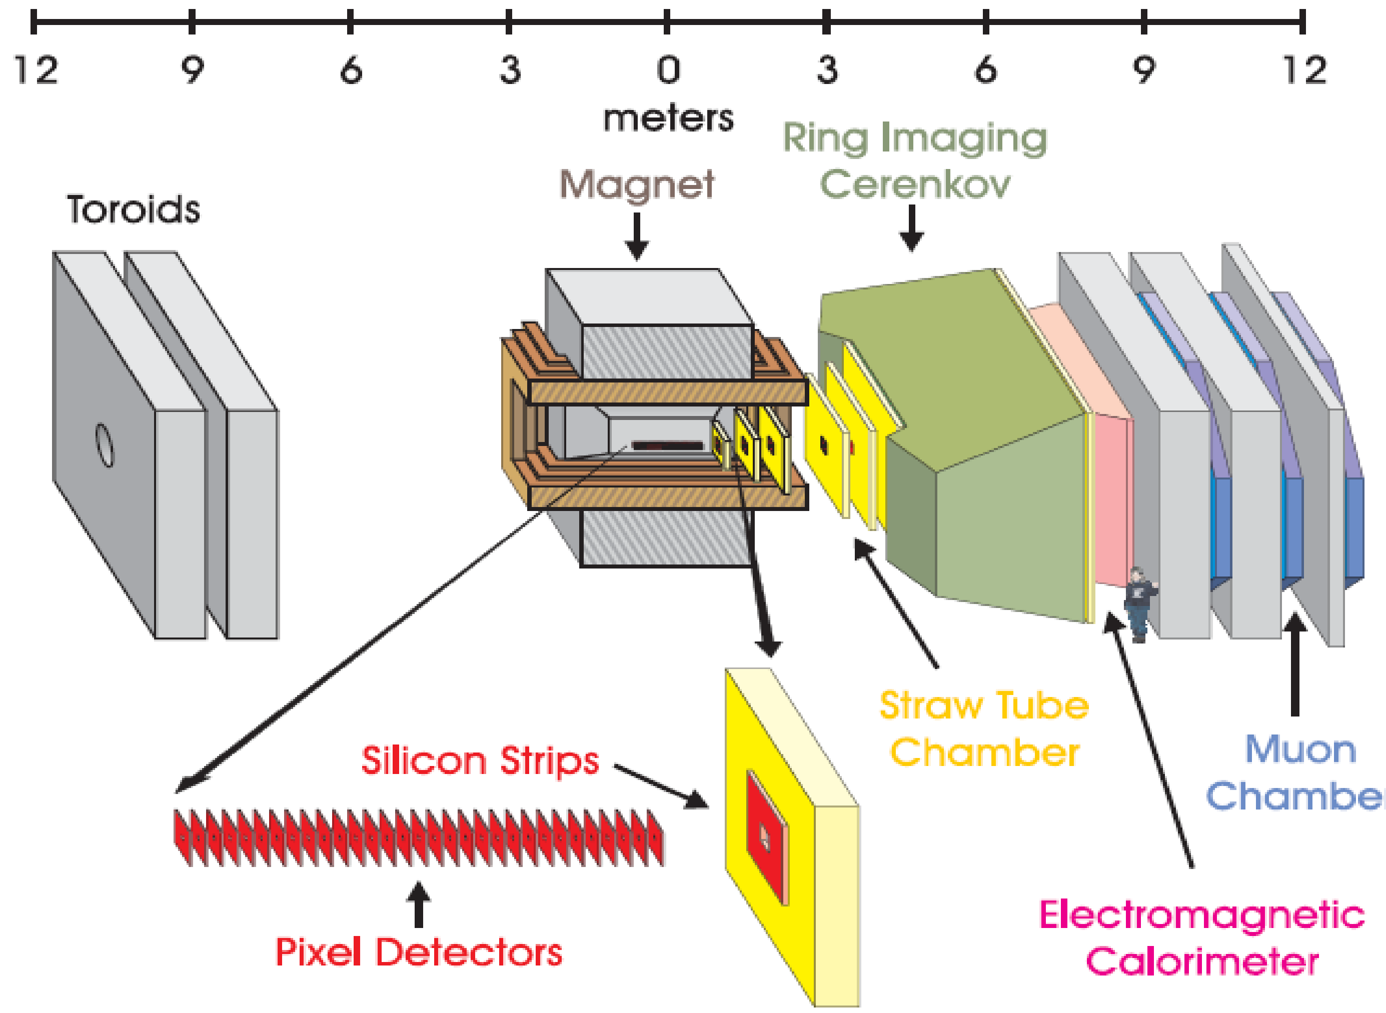

BTeV Trigger

- Identifying secondary vertices is key
- Developed Level 1 secondary vertex trigger
  - Exploits many layers of pixel detectors
    - Low occupancy leads to low fake rates
  - Still implies relatively complex computations
    - Pre-programming of patterns in FPGAs not practical for vertexing due to large number of possible patterns
    - Implies long latency
  - Long delays turned the design from “futuristic” to “difficult” to “feasible”
    - Will never be built though

Level 1 Pixel Trigger

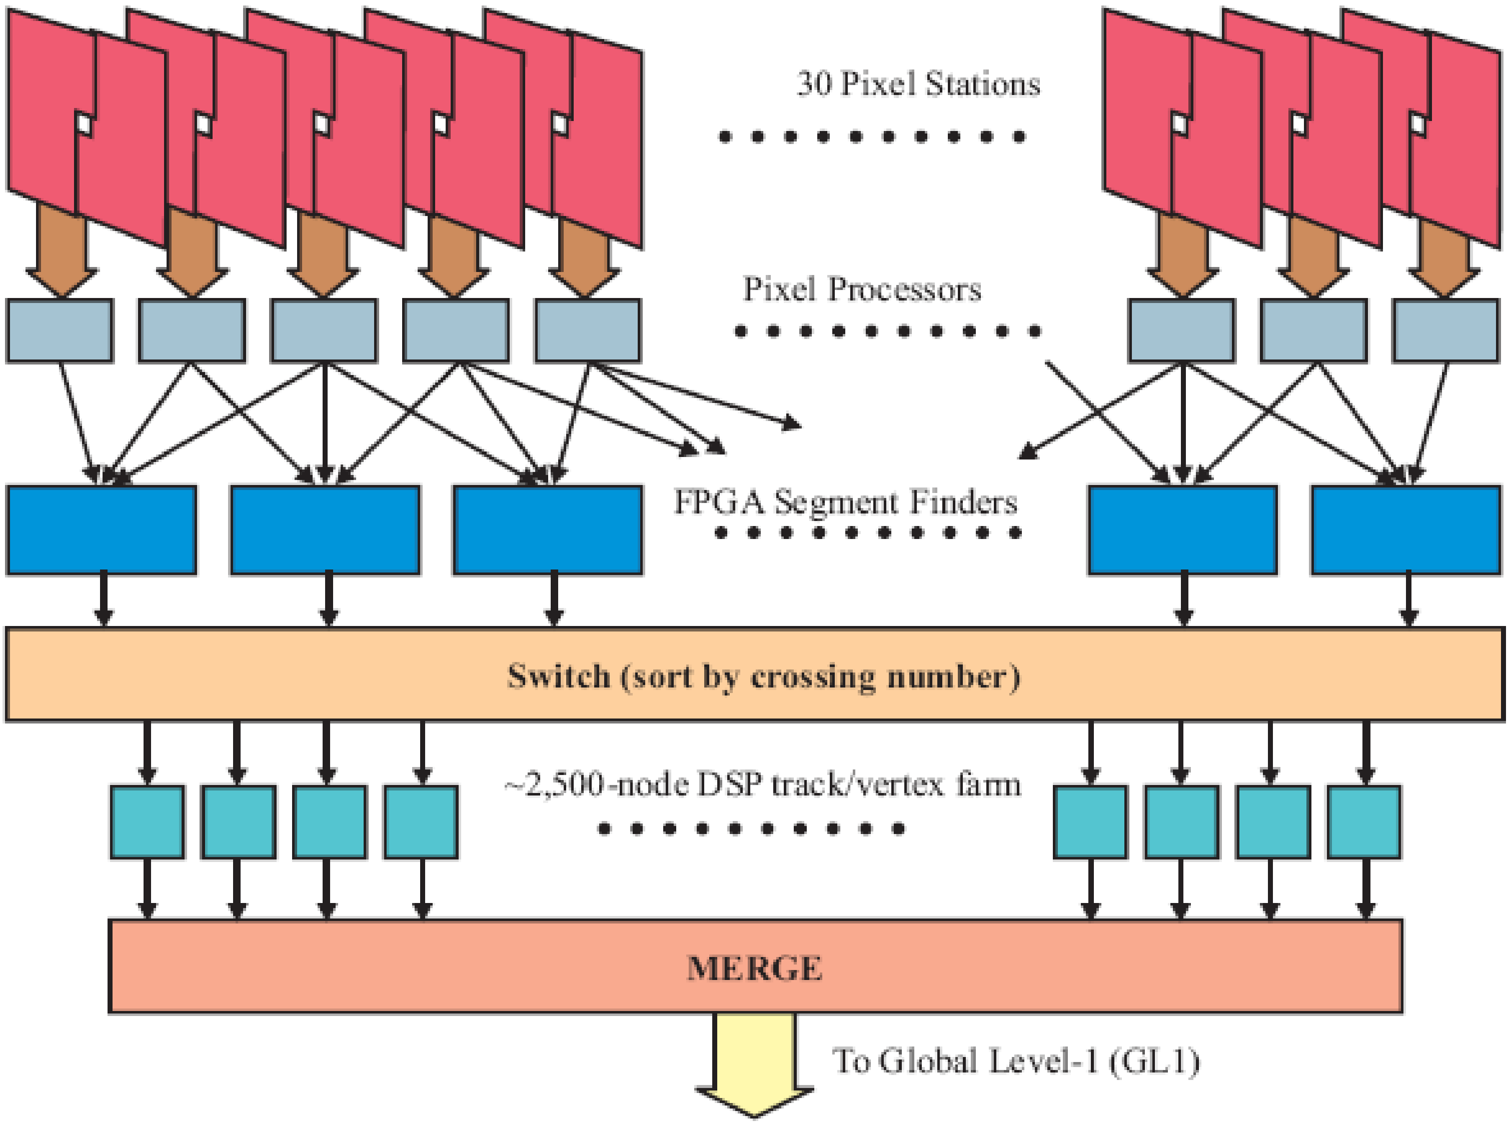

Overall Architecture

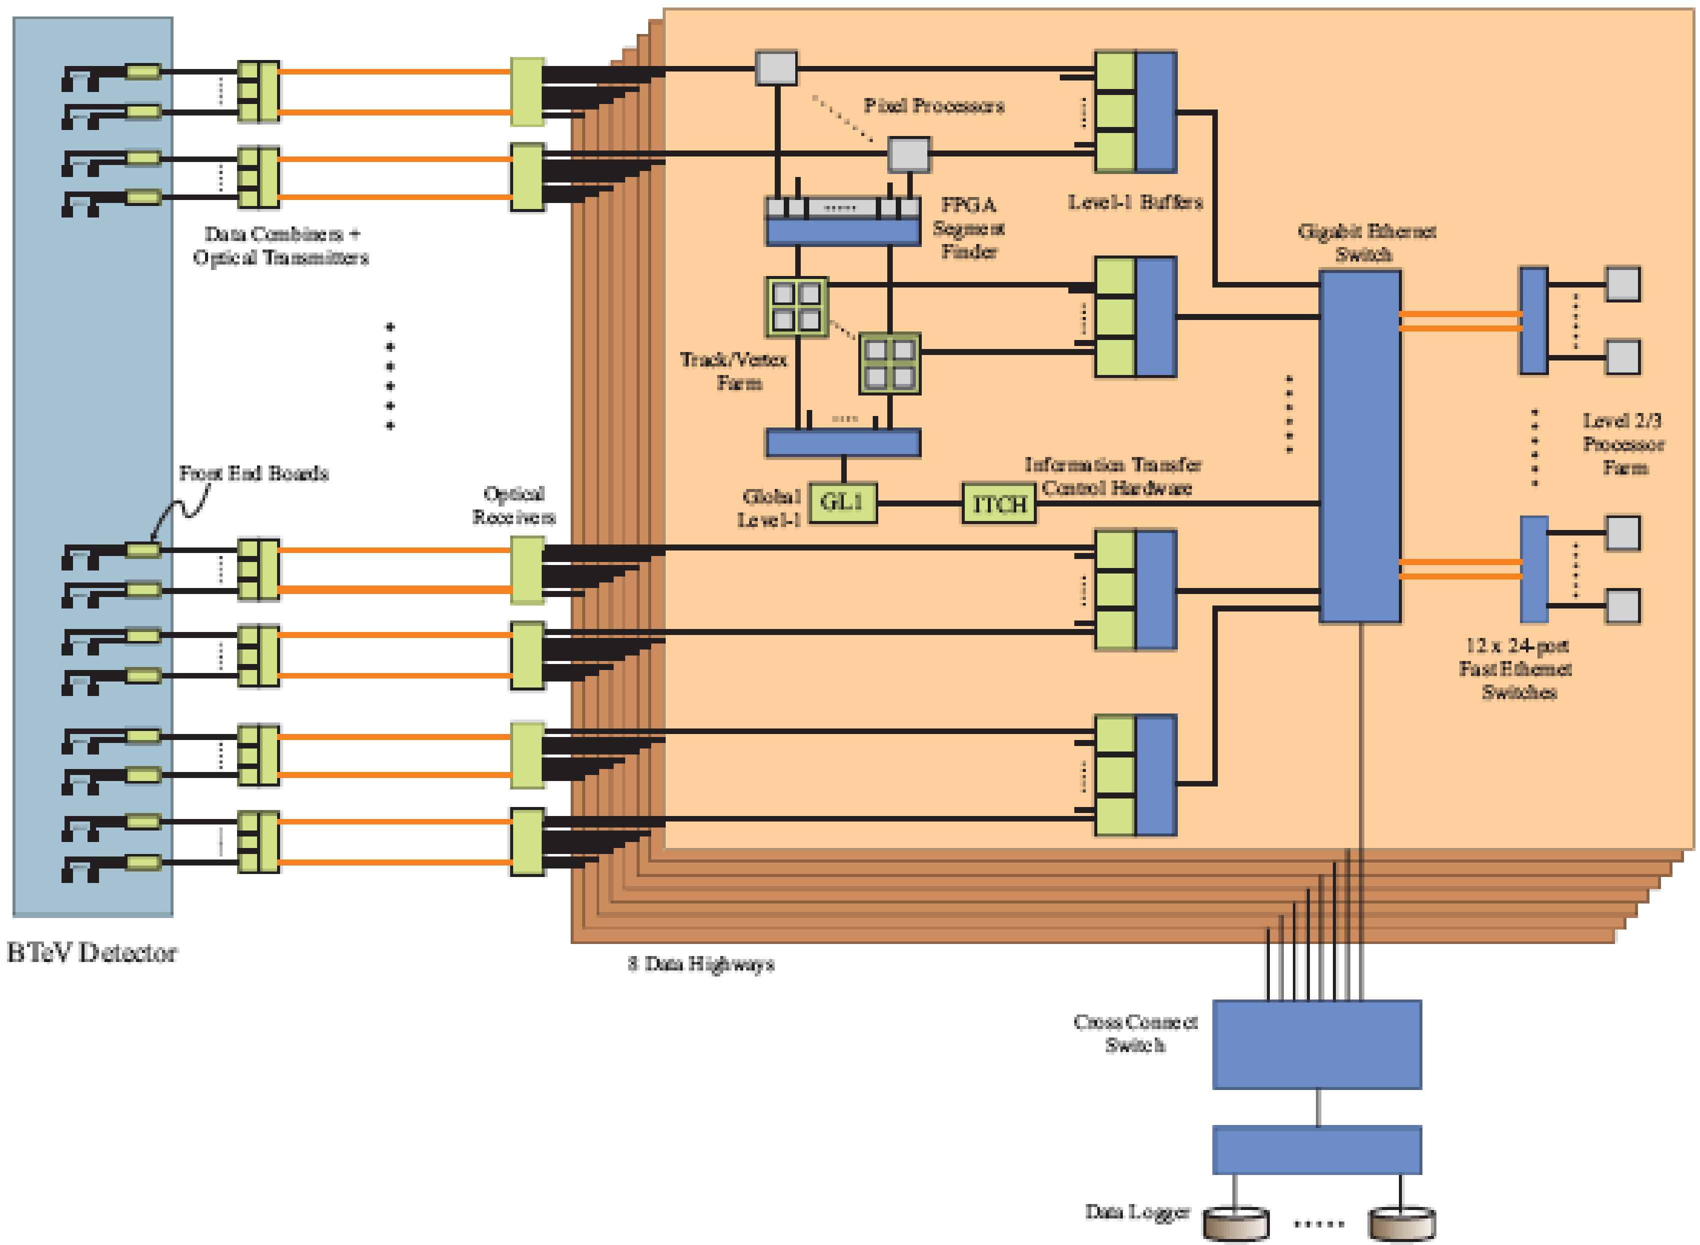

Rates and Latencies

- Data split in eight, sent to “highways” which each implement a full trigger system
  - 100 GB/s into each highway
- L1 buffers implemented in *commodity SDRAM*
  - Allows 1 second L1 latency!
    - More than enough time for L1 pixel trigger
  - SDRAM managed by an FPGA
- L1 accept rate ~40 kHz, L2/3 accept rate ~2 kHz
  - 200 MB/s to tape
- L1 muon trigger to measure L1 pixel efficiency

Commissioning

“Hardware” Commissioning

- The trigger is the nervous system of the experiment. It's very complex, relatively fragile, and bad behavior can be very debilitating. It's also where you discover big problems (hot cells, etc. leading to unacceptable rates)
  - OTOH, it's very difficult to detect problems at the < 1% level (see later...)
- The trigger is the one system where subdetectors can have a large impact on each other
  - Pathological behavior that doesn't affect one system will bring down another
  - Teststands and testbeams do not reproduce reality

Getting Started

- A lot can in principle be done without beam
  - Read out “noise”
    - Events are either small, or huge (no zero suppression)
  - Cosmics
    - Our detectors are designed for events that happen at specific times
  - Testpatterns
- None of these are substantially better than teststands
  - Major benefit is checking out combined control software

Some Lessons from Run II

- You can't fully debug a trigger and readout stage until the downstream system can take the full rate
  - No matter how sophisticated the simulated triggers in “the lab”, the real thing will find a pattern that leads to problems (due to race conditions, buffer management issues, ...)
  - Corollary: if you increase the rates in steps, you need to verify data integrity at each step, in addition to finding and fixing crashes/hangs
  - This is, in general, done *by choice* (deferral of purchase of PC farms)

- Diagnostic tools diagnostic tools diagnostic tools! You can never have too many diagnostic tools!
  - Dataflow GUIs are among the most valuable tools
    - See where the data is stuck
    - See buffer occupancy, both instantaneous and averaged
    - Making a good one takes some thought...
  - Need the capability (for experts) to examine the data at all interfaces
    - Yes: hex dumps
  - Dump status registers of any type of hardware
  - Hard to guess what the most “interesting” problems will be -> code needs to be clear and documented so that others can adapt it

- Hardware does funny things
  - Designer usually can't anticipate wacky conditions that can be generated by a real detector
    - Therefore can't simulate them
  - Interaction with other systems then leads to race conditions, and the *impossible* happens
  - People also forget about stuff inserted for debugging
  - So, never underestimate the hardware's ability to do *"interesting"* things
- Also remember that many of today's experts will have another job when LHC beams collide....

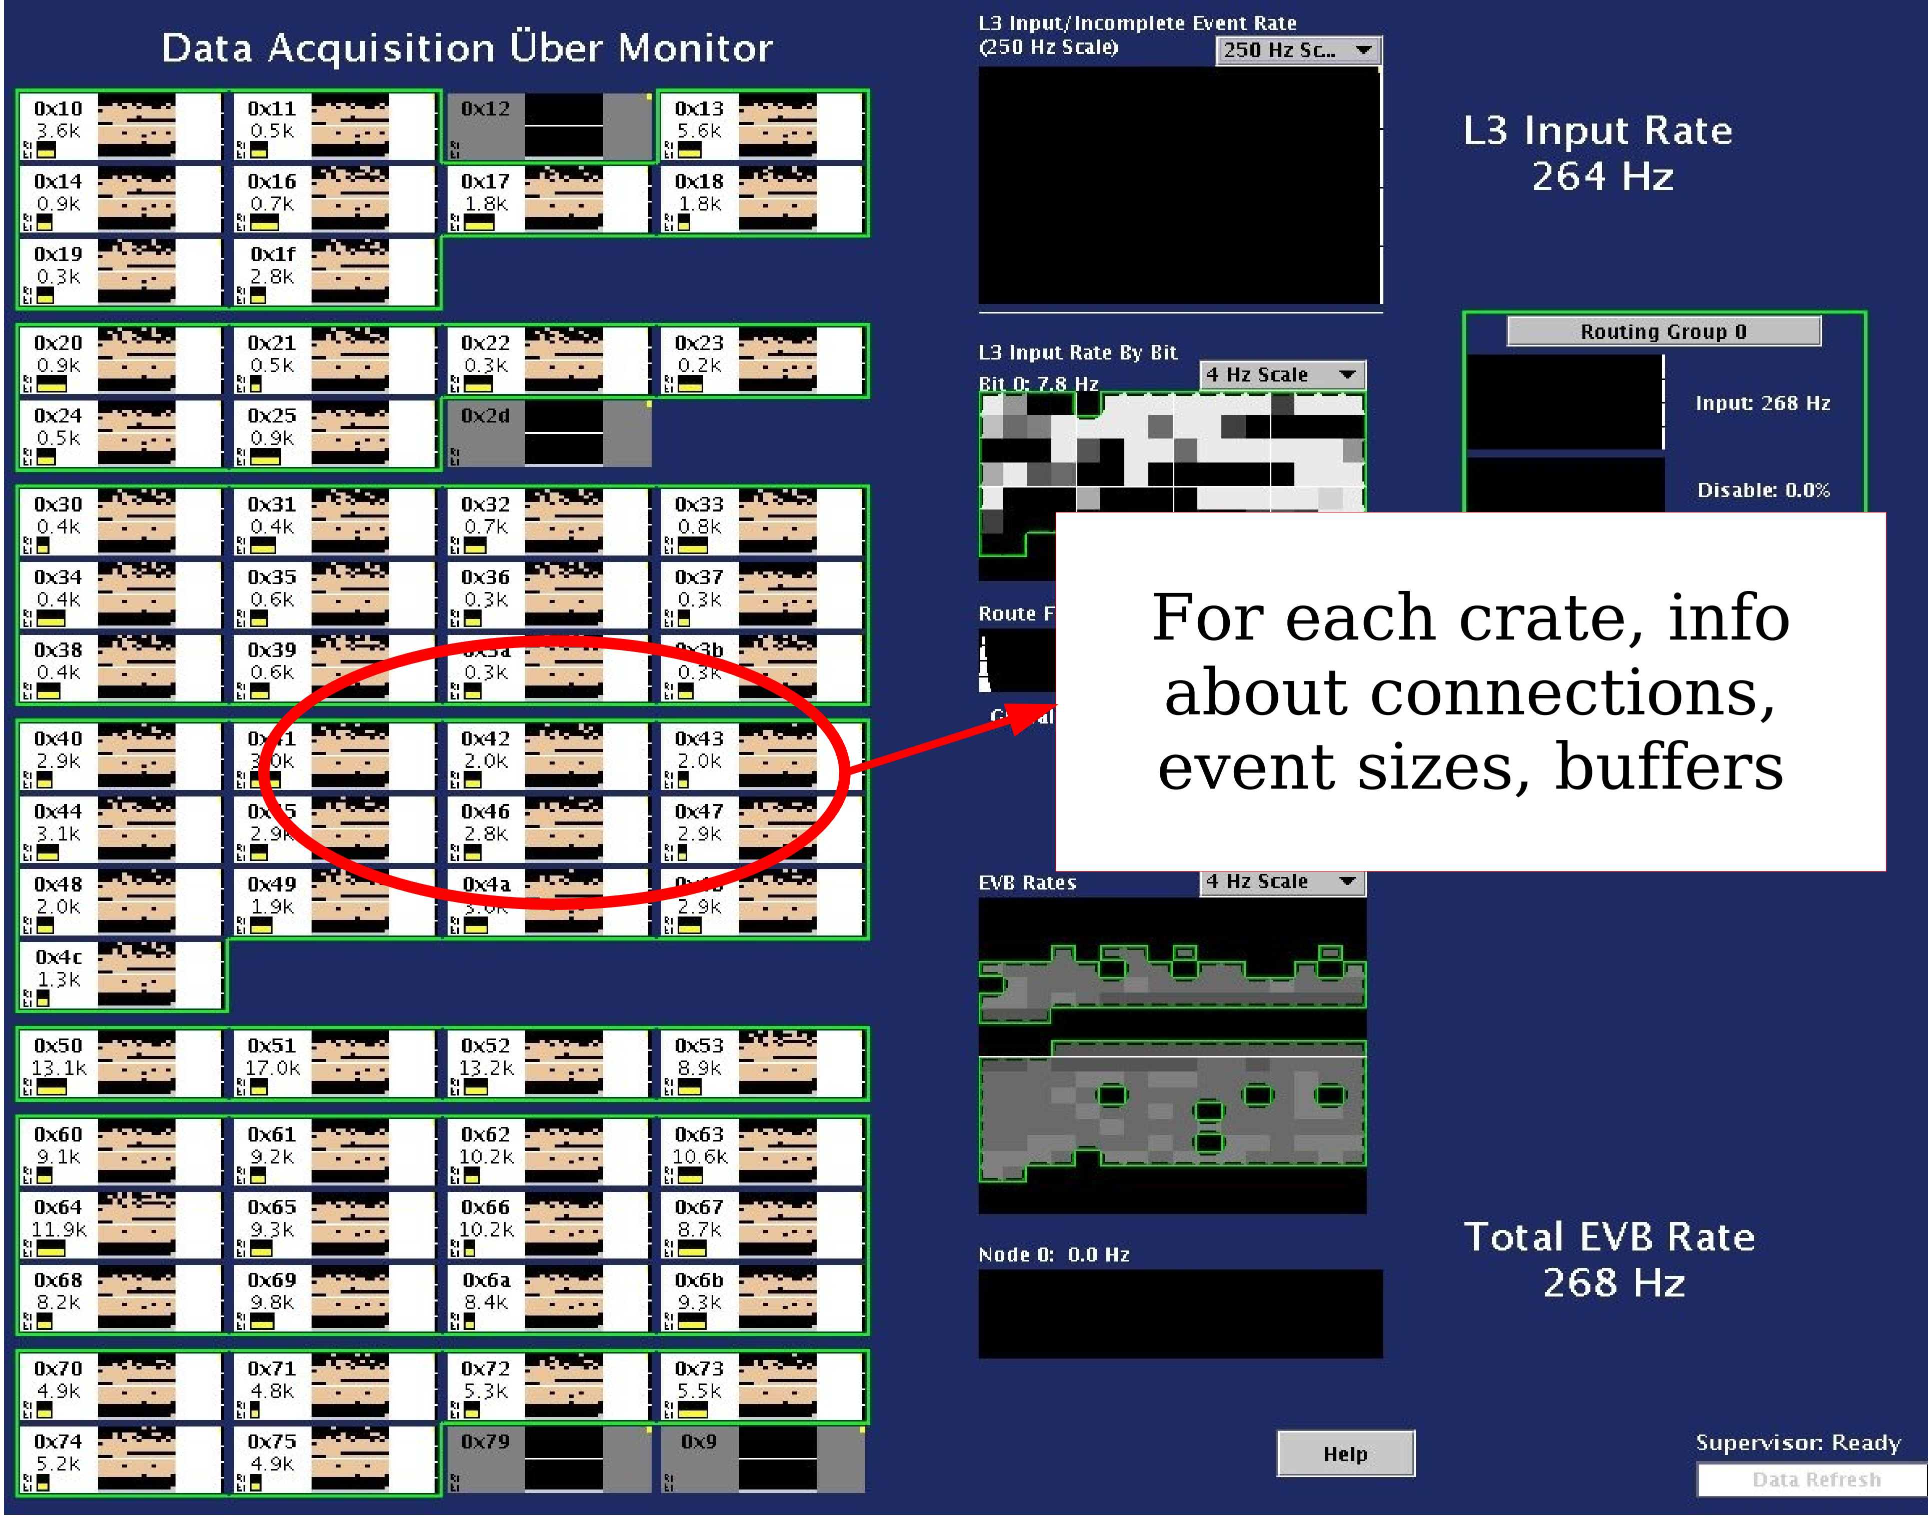

For each crate, info about connections, event sizes, buffers

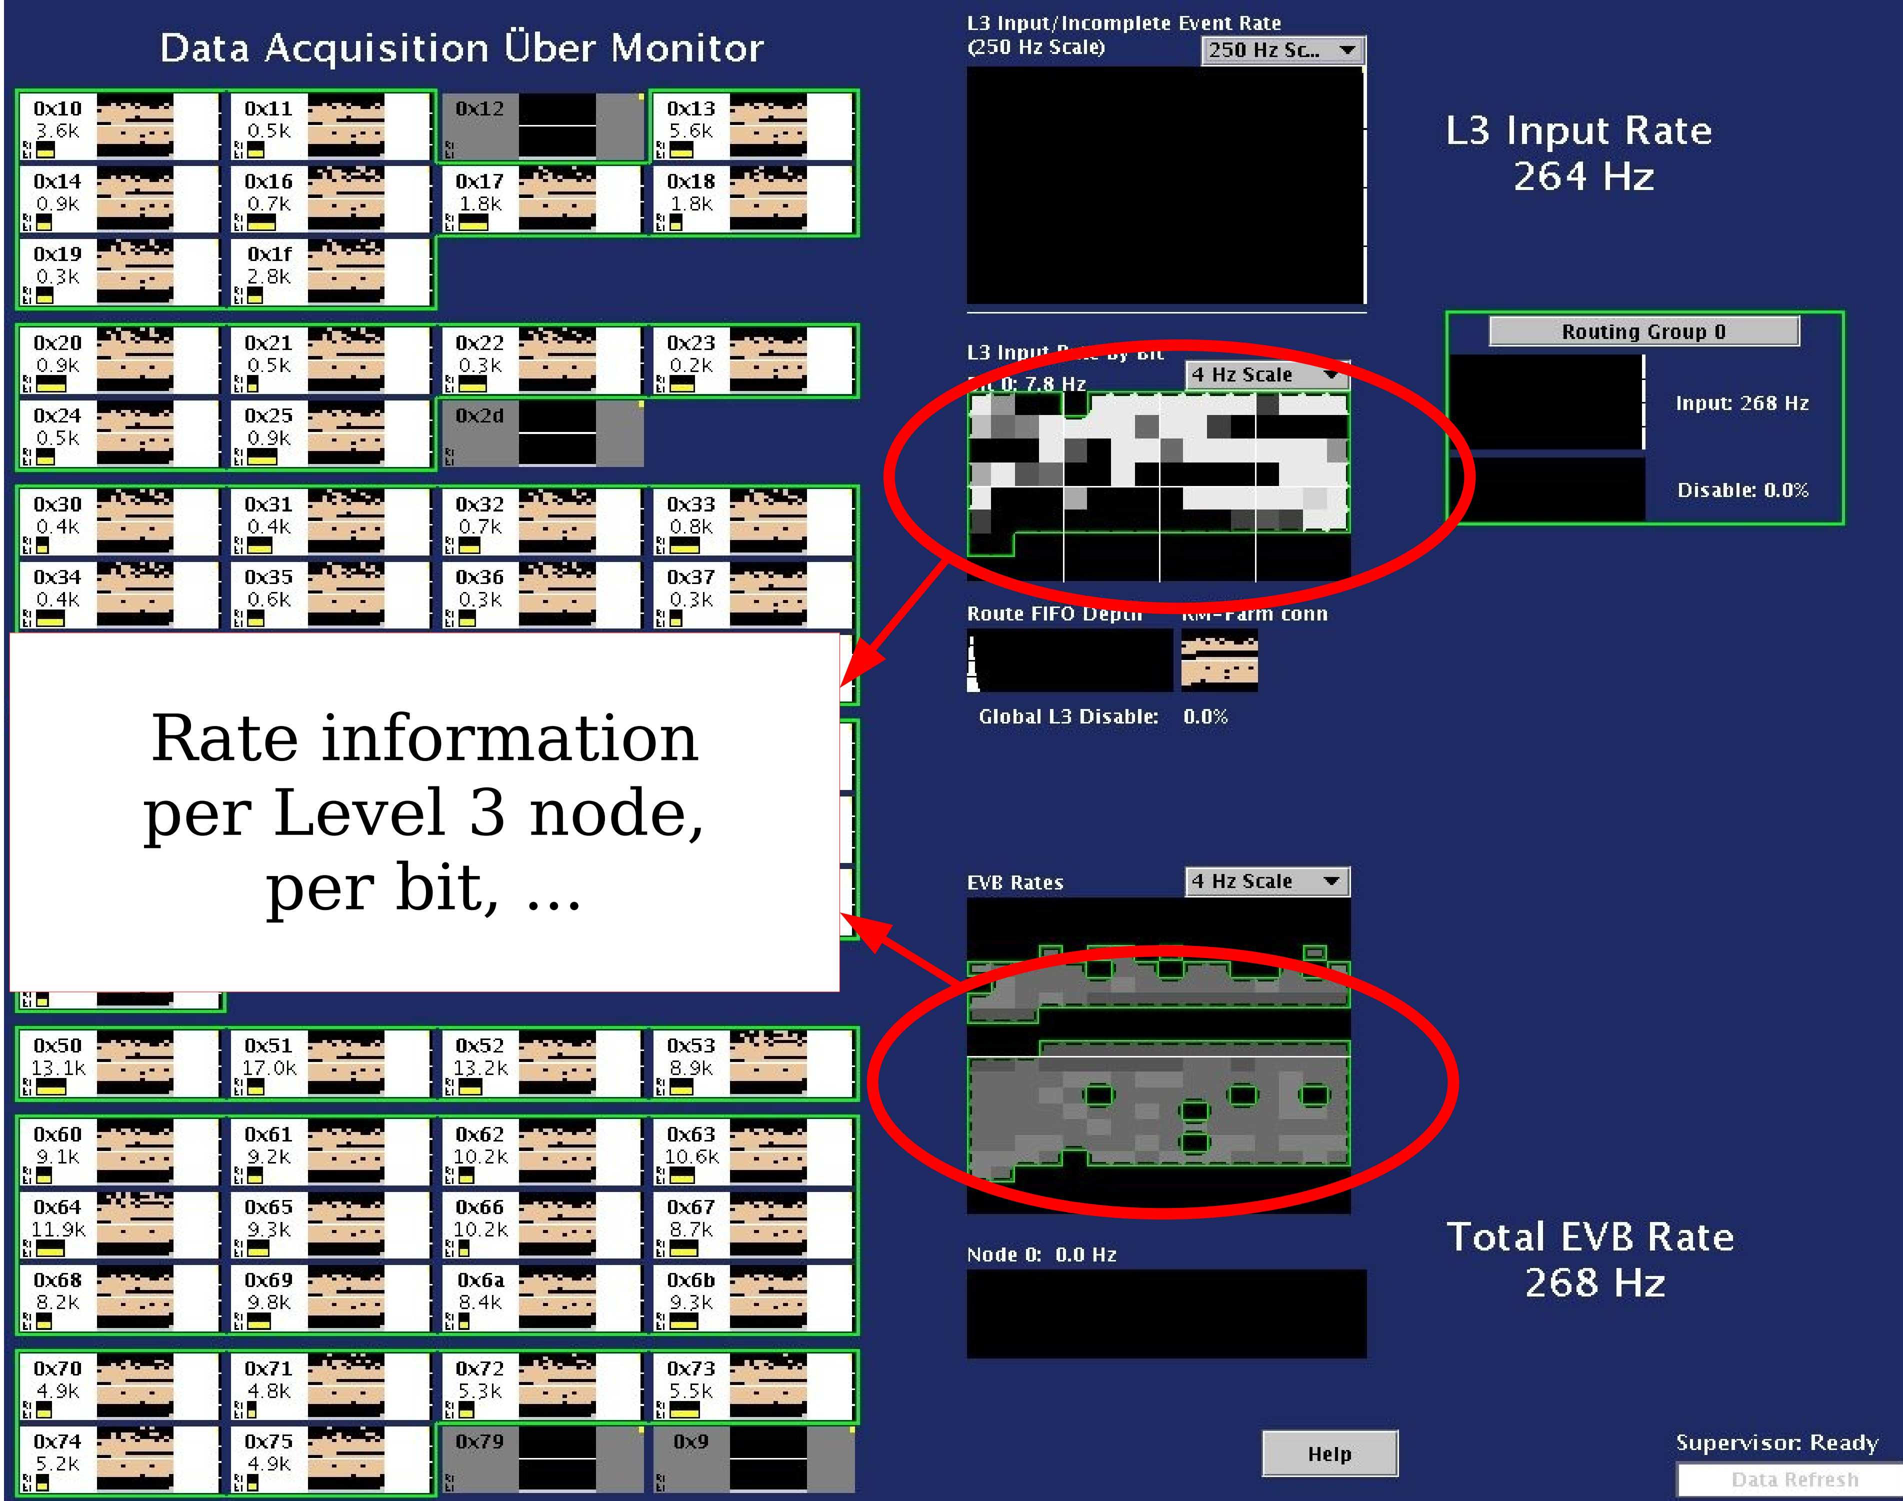

Rate information per Level 3 node, per bit, ...

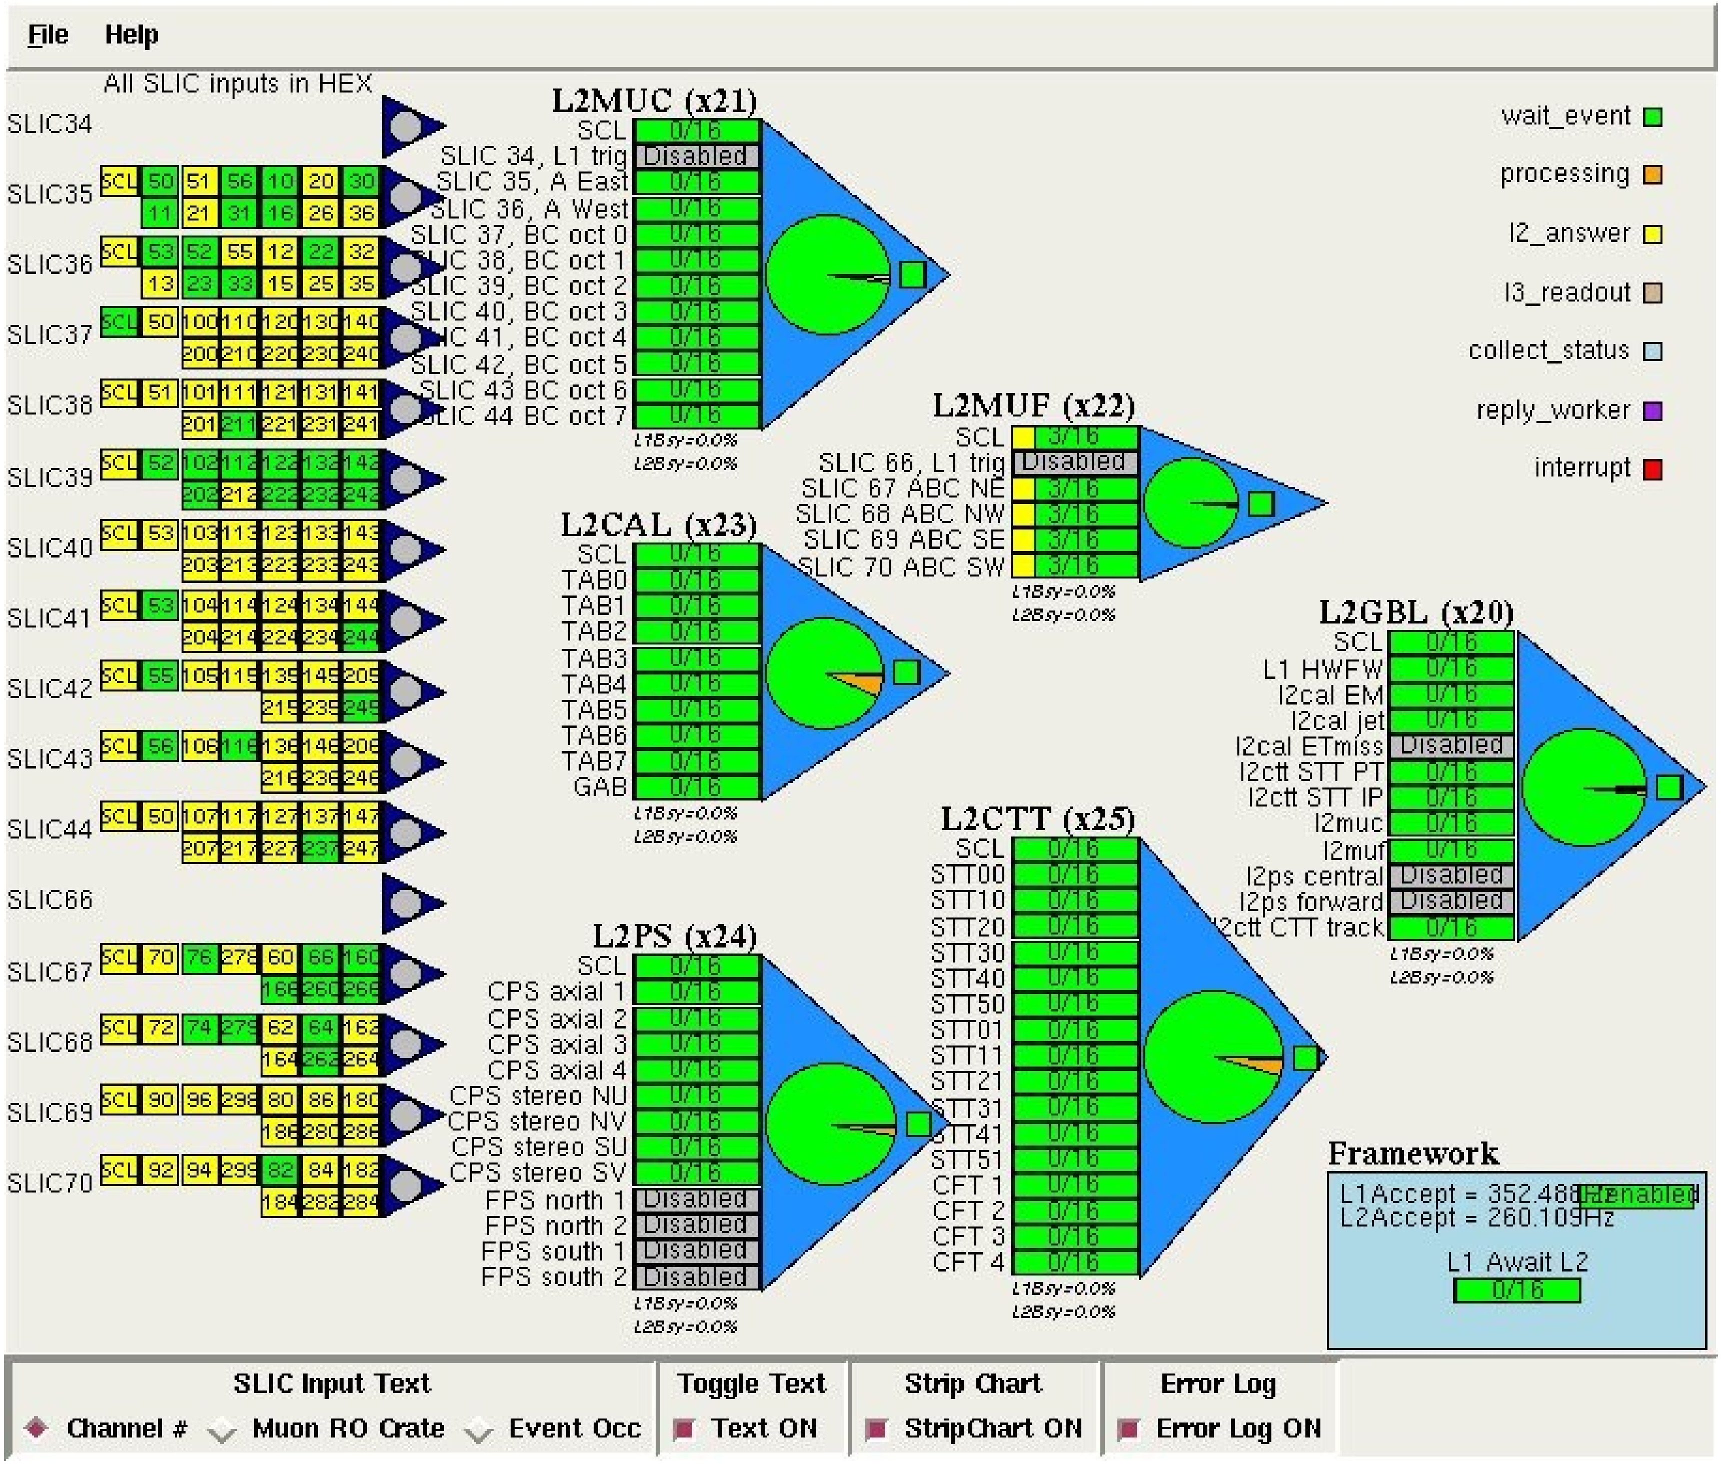
File
Help
All SLIC inputs in HEX
SLIC34
SLIC35
SLIC36
SLIC37
SLIC38
SLIC39
SLIC40
SLIC41
SLIC42
SLIC43
SLIC44
SLIC66
SLIC67
SLIC68
SLIC69
SLIC70
L2MUC (x21)
SCL
SLIC 34, L1 trig
SLIC 35, A East
SLIC 36, A West
SLIC 37, BC oct 0
SLIC 38, BC oct 1
SLIC 39, BC oct 2
SLIC 40, BC oct 3
SLIC 41, BC oct 4
SLIC 42, BC oct 5
SLIC 43 BC oct 6
SLIC 44 BC oct 7
L2MUF (x22)
SLIC 66, L1 trig
SLIC 67 ABC NE
SLIC 68 ABC NW
SLIC 69 ABC SE
SLIC 70 ABC SW
L2CAL (x23)
TAB0
TAB1
TAB2
TAB3
TAB4
TAB5
TAB6
TAB7
GAB
L2PS (x24)
CPS axial 1
CPS axial 2
CPS axial 3
CPS axial 4
CPS stereo NU
CPS stereo NV
CPS stereo SU
CPS stereo SV
FPS north 1
FPS north 2
FPS south 1
FPS south 2
L2CTT (x25)
STT00
STT10
STT20
STT30
STT40
STT50
STT01
STT11
STT21
STT31
STT41
STT51
CFT 1
CFT 2
CFT 3
CFT 4
L2GBL (x20)
L1 HWFW
l2cal EM
l2cal jet
l2cal ETmiss
l2ctt STT PT
l2ctt STT IP
l2muc
l2muf
l2ps central
l2ps forward
l2ctt CTT track
wait_event
processing
l2_answer
l3_readout
collect_status
reply_worker
interrupt
Framework
L1Accept = 352.488Hz
L2Accept = 260.109Hz
L1 Await L2
SLIC Input Text
Channel #
Muon RO Crate
Event Occ
Toggle Text
Text ON
Strip Chart
StripChart ON
Error Log
Error Log ON

Physics Commissioning

- Two major aspects:
  - Calibration
    - Very similar to physics calibration... but not quite
  - Developing a trigger list
    - Difficult process: development of a new list takes ~6 months at DØ (and from what I heard a similar amount of time at CDF), and the new list usually barely runs at the luminosity it was designed for
    - Partially due to the large number of available features, and partially due to the difficulty in accepting a loss in efficiency
- Remember: you never see most events!

Calibration at Level 1

- Example: calorimeter
  – Constants are downloaded infrequently
- In principle, a simple problem:
  – Determine pedestals from “noise runs”
    - Ah, but what exactly does that mean? Pedestal = “number of ADC counts without signal”
    - What about pileup? Underlying event?
    - At the trigger level, in principle would prefer to factor pileup into pedestals... but then they depend on luminosity!
  – Determine gains by comparing with offline
    - Of course, that means “offline” is “calibrated”

Calibration at Level 3

- In principle, can use ~offline calibration
  - Make sure it's valid for that time
  - Having hundreds of processes access DB simultaneously is problematic, so need to distribute a “file”
  - Ensure all nodes always have the right version
    - This is true for any file: filtering code, geometry, etc.
    - In principle, versioning through the trigger list is probably the safest solution
  - And then, of course, verify

Trigger List Development

- Complex task:
  - Optimize efficiency within a certain rate budget
    - Implies being able to estimate rates
  - Many signatures, particularly in multi-purpose experiments
  - Enormous flexibility, especially at higher levels
- Current lists:
  - DØ: ~600 triggers, including monitoring triggers
  - CDF: ~180 triggers, idem

Rate Estimates

- Rates are very sensitive to events that are not recorded
  - Ideally, would like to have 10s of seconds of unbiased accelerator data
    - Not practical: at LHC: 40 MHz x 10s / 200 Hz = 2 10^6 s, or 1-2 months of exclusive datataking
  - Take “enhanced bias” data: use lowest thresholds for each of the Level 1 objects, apply prescales at Level 3 (but still useful to run the algorithms)
    - Still need a lot of bandwidth
    - No need to reconstruct – only trigger objects needed offline

Rate Projections

- Unfortunately, can't take all enhanced bias at low luminosity
  - And even at high lumi, you're typically designing a new list for even higher luminosities
- Many trigger objects have non-linear rates due to increased occupancy, so two options
  - Fit the rate vs lumi curve
    - Extrapolation with large uncertainty
  - Re-weigh events as a function of the number of primary vertices
    - Implies running reconstruction

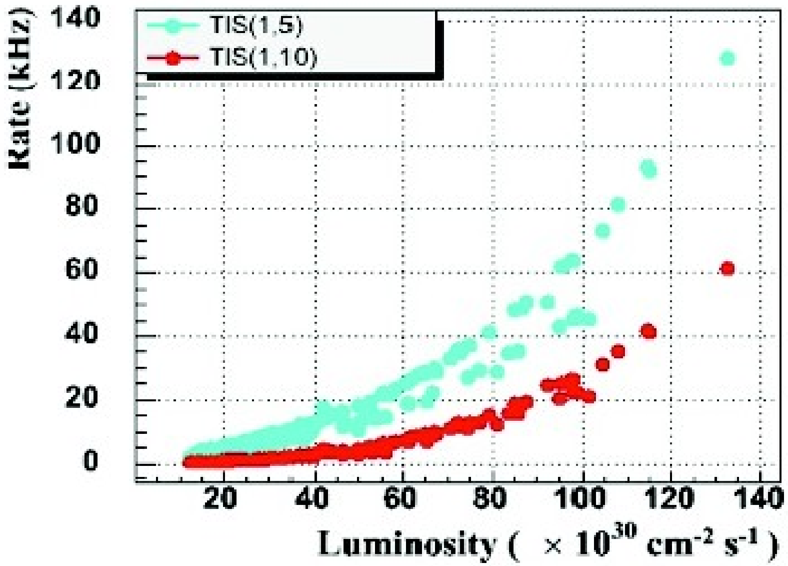

Initial Efficiency Estimates

- Trigger objects from simulation useful tool for initial efficiency estimates
  - MC usually does a fair job at reproducing p^T distribution of signals
    - Ok, maybe not for jets in W/Z+jets, but the jets shouldn't be crucial in your trigger strategy there
  - OTOH, MC is usually not so good at reproducing variables that depend on occupancy, like isolation, “hadronic veto”, missing E_T
    - Often, these involve “absence of signal”
- More on determining efficiencies *a posteriori* soon

Trigger Simulation

- Two tasks:
  - Determining trigger efficiency for a particular signal during the design phase
    - This can, in fact, be done to a large extent by having "trigger objects" written in the simulated data
    - Exception is development of new algorithms – arguably an "expert" task anyway
  - Verifying the trigger decision
    - Critical: given the signal in the detector, did the trigger issue the expected decision
      - Particularly important to find problems in firmware
    - Also uncover unexpected correlations in trigger list, and optimize order of filters

- So, *two* tools:
  - As part of reconstruction, write trigger objects in the data
    - Main use cases: trigger development (existing algos), comparison of trigger and precision readout
    - Fairly simple: for real data just extract, for simulation apply trigger algorithm
  - Trigger simulation
    - Main use cases: trigger verification, algorithmic development
    - Programmable: need to be able to feed in an “online” trigger list
    - “Users” need to be able to modify that list....
      - Or have one produced
    - Detailed simulation of firmware

Trigger List Contents

- Natural to group triggers by final state:
  - Single muon/electron/photon
  - Di-muon/electron/photon
  - “EM” + muon
  - Lepton/photon + jet(s)
  - Monojet + MET
  - Multijet
  - “Impact parameter”
  - ...

Each Final State

- Each group consists of many triggers
- For “single objects”:
  - Multiple p^T thresholds, tighter quality criteria for the lower ones
- For “mixed triggers” (e.g. lepton+jets):
  - Play with number of objects
  - Different mixes of thresholds
    - Generally justified by the physics: in top for example, total event E_T is above certain threshold

Sums

- Tempting to trigger on "Sums", like MET
  - Highly non-linear with luminosity
- Almost always better to use individual objects
  - Acolinear jets
  - "Missing H_T"
- Or cross-correlate different detectors with independent resolution/noise:
  - Angular and magnitude match between MET and "Missing p^T"
    - Tracking is expensive in CPU

Inside Groups

- Two very different categories:
  - "Prescalable": the physics case does not need to get *all* events
    - Some B-physics topics
    - Multijet at low/moderate p^T
    - "Monitoring triggers" - note that one analysis' monitor trigger is the other's physics trigger
  - "Unprescalable"
    - Searches for and studies of rare processes
      - Not the same at LHC and Tevatron, e.g. top, W

Putting It Together

- A working strategy is to start with "unprescalables" and cap rate at x% of max (excluding calibration)
  - x ~ 70-80 (typically)
- Conflict arises because goal is not met
  - Reduce rate by increasing thresholds and/or tightening quality criteria
  - But who should sacrifice efficiency?
    - Difficult decision, particularly in multi-purpose experiments
- Then add in "prescalables"

Compromises

- Some physics in "prescalables" easier at low lumi
  - Exclusive B decays
  - Diffractive physics
- Trade bandwidth:
  - In principle, "unprescalables" only fill half the bandwidth at half the max lumi
  - Fill in the other half with events for analyses that particularly like "cleaner events"
- Of course, rate-to-tape need not be the same at all luminosities
  - But check with offline people...

Trigger and Analysis

Trigger and Analysis

- Trigger reduces the data rate from MHz to $\mathcal{O}(100)$Hz
  - So the trigger does 99.99% of the physics analysis, and you better understand the biases it introduces
- Question 1: trigger efficiency with respect to what?
  - Absolute? Difficult, not necessarily useful
  - Usually, w.r.t. offline reconstruction efficiency
    - Disadvantage: moving target, especially in the early days

Trigger Efficiency

- To determine trigger efficiency, really need to determine trigger inefficiency
  - Means determining which events you *didn't* get
  - First tool is monitoring triggers
    - Typically use same trigger objects with lower thresholds or quality criteria
    - Big caveat is that these are heavily correlated with primary trigger (same object, so same acceptance, etc.)
  - Diverse trigger menu
    - Get events that passed orthogonal triggers but failed yours
      - E.g. muon triggers for jet efficiency, and biases!
    - Logistics! (Depends on streaming model)

Streaming and Trigger Efficiency

- Two existing streaming models:
  - Based on trigger decision:
    - Implies events fired by orthogonal triggers are in different streams
    - Need a way to go through streams without unpacking events so that relevant ones can be found quickly
  - Based on offline reconstructed objects:
    - Means no online streaming (or randomized online streaming):
      - No offline (re)processing priorities possible
    - Trivial to get at events that failed main trigger(s) but have good reconstructed object

Functional Form

- At perfect resolution, the trigger efficiency as a function of a certain parameter is a step function
- But detectors aren't perfect
  - "Step" is convoluted with Gaussian
- Integral of a gaussian is called the "Error Function"

$$\mathrm{erf}(z) \equiv \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2}\, dt.$$

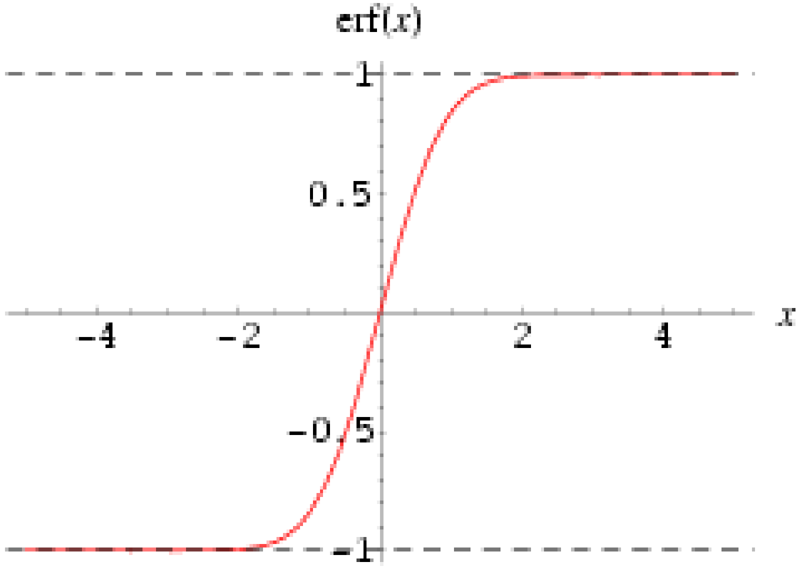

In Practice

- Most used is efficiency vs p^T:
  - Plot is usually called "turn-on" curve
  - "Turn-on" point is usually where efficiency reaches ~plateau (sometimes midpoint)
  - Many analyses only use data above turn-on, due to severe systematics below
  - To get rate, need to convolute with exponentially dropping QCD spectrum:
    - Most events are at low end

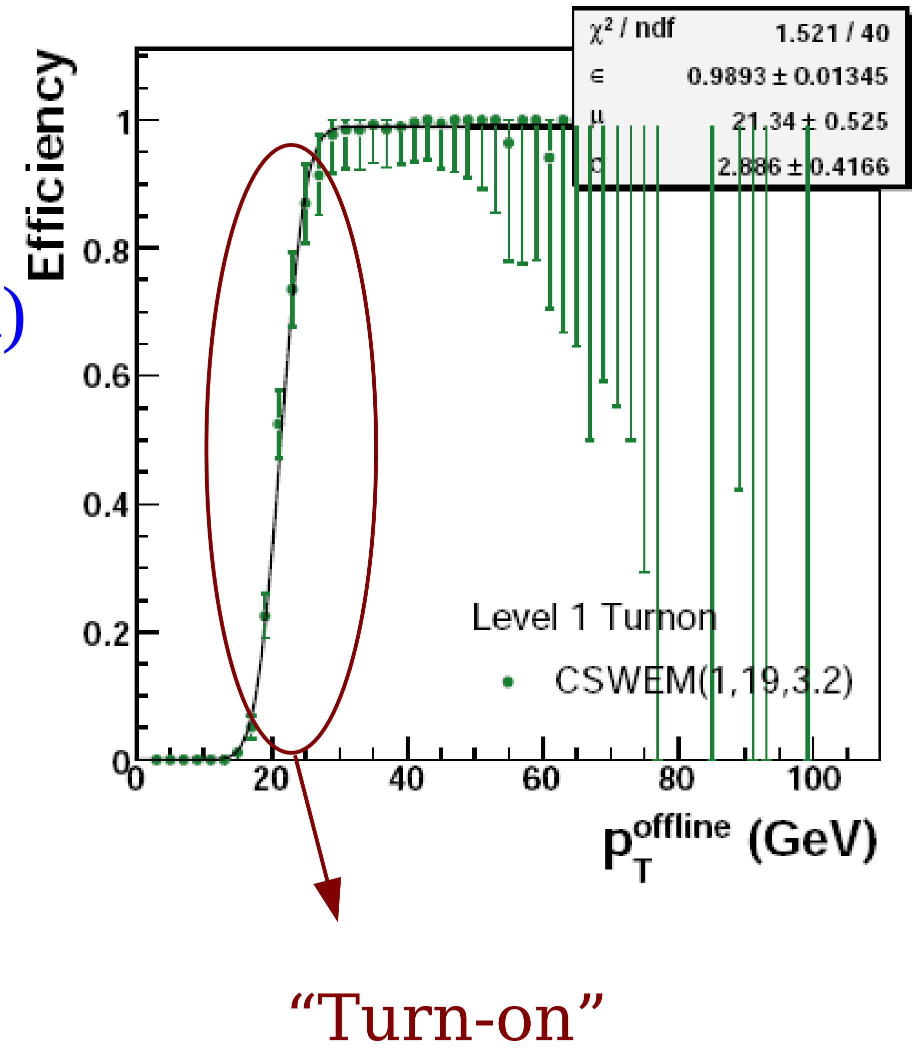

Electron Trigger Efficiencies

- First problem is to get a clean sample of real electrons:
  - Most medium p^T objects that satisfy good calorimetric criteria (EM fraction, isolation, shower shape) are *jets*, so you get

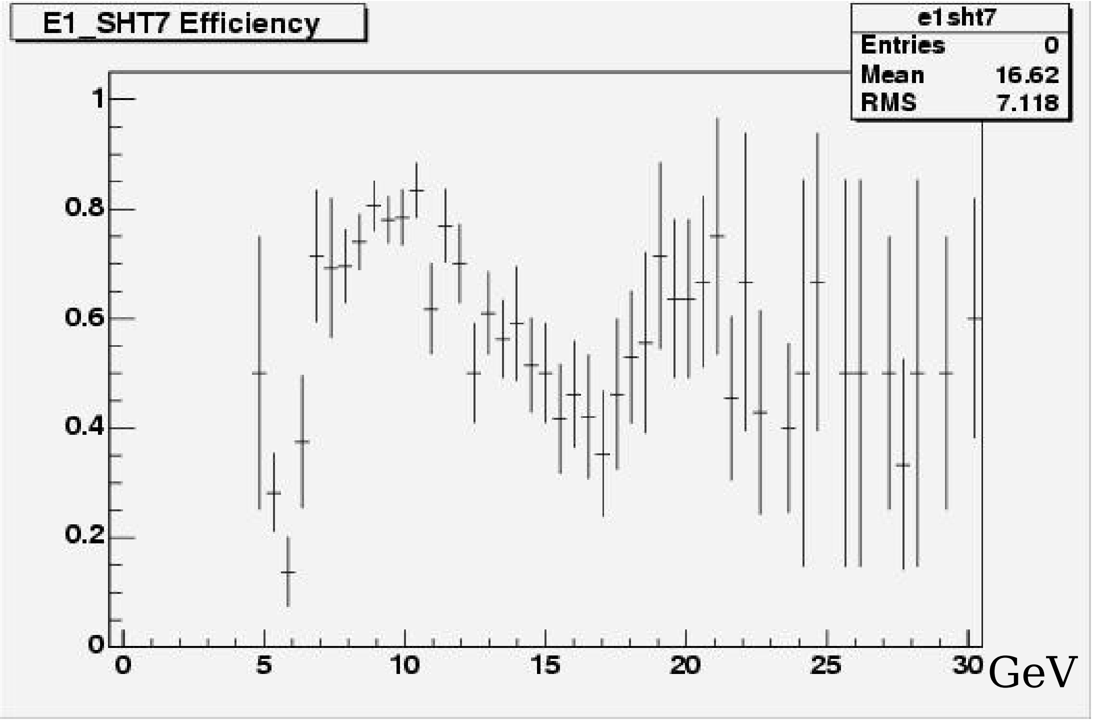

Particularly painful here because no track requirement

Selecting Electrons

- A good source of true electrons at Tevatron is Z, with benefit that it has *two*!
- "Tag and Probe":
  - Select events with two good offline electrons, look at invariant mass and select Z's
  - Tag electron is matched to trigger object in single electron triggers
  - Derive trigger efficiency from fraction of time second electron is also reconstructed as an electron by the trigger

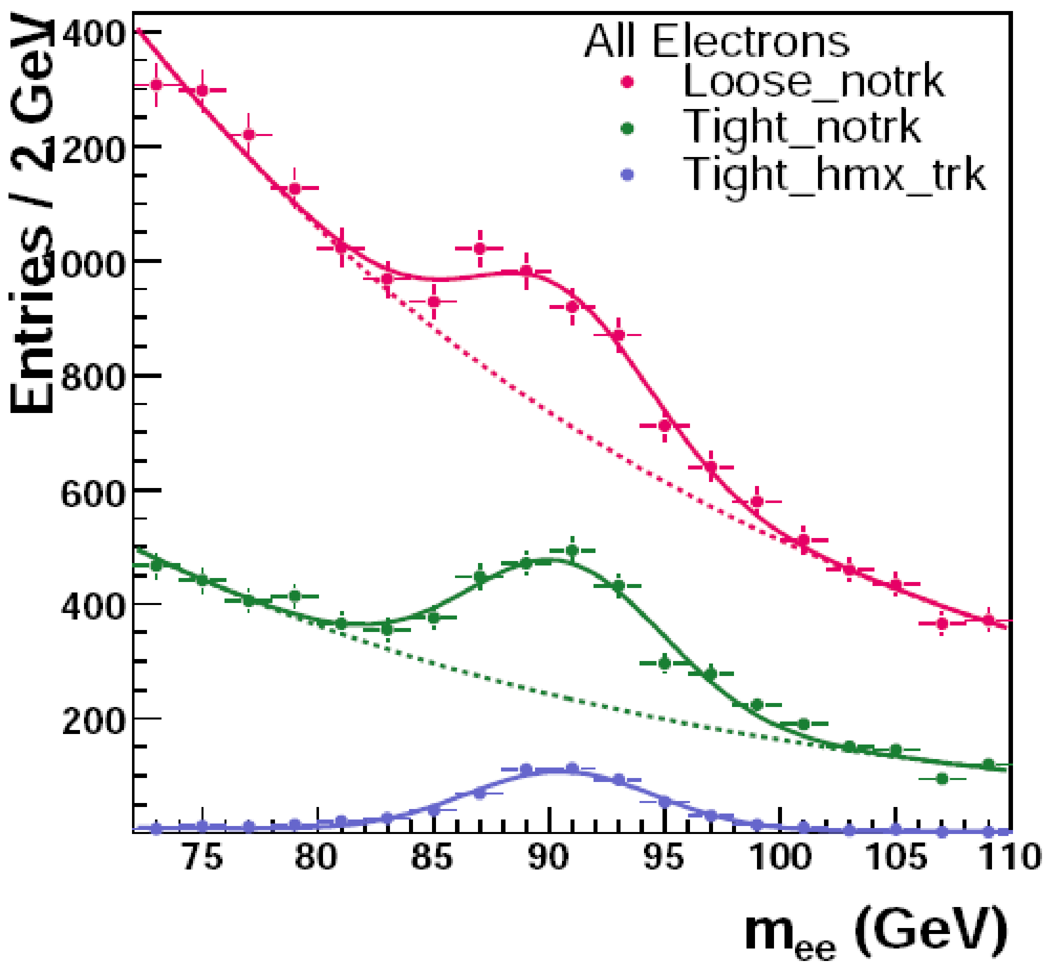

Results

- Determination of efficiency becomes "straightforward"
- Then verify
  - linearity (threshold vs turn-on)
  - turn-on vs offline requirements (i.e. contamination by jets)
  - ....

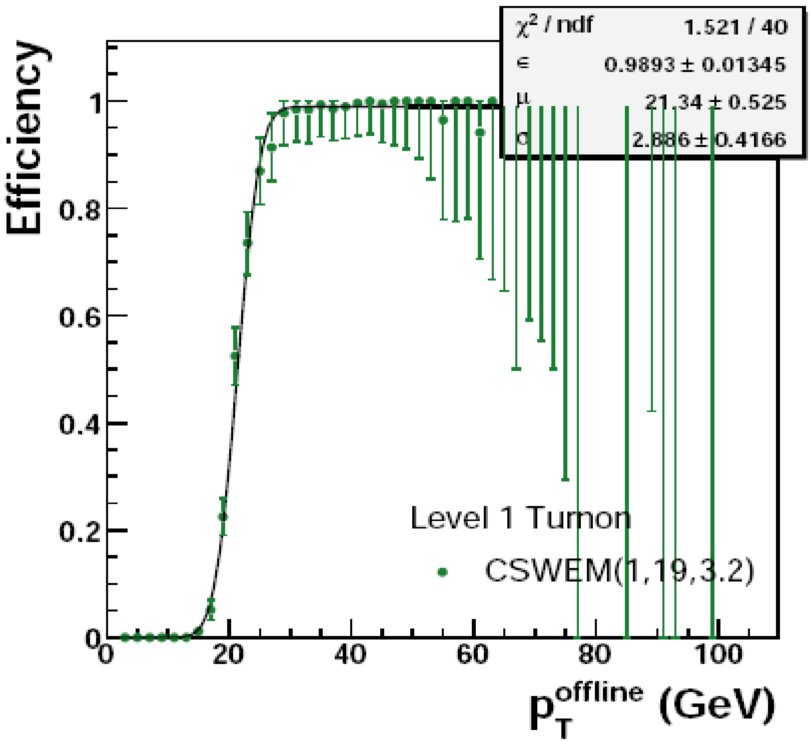

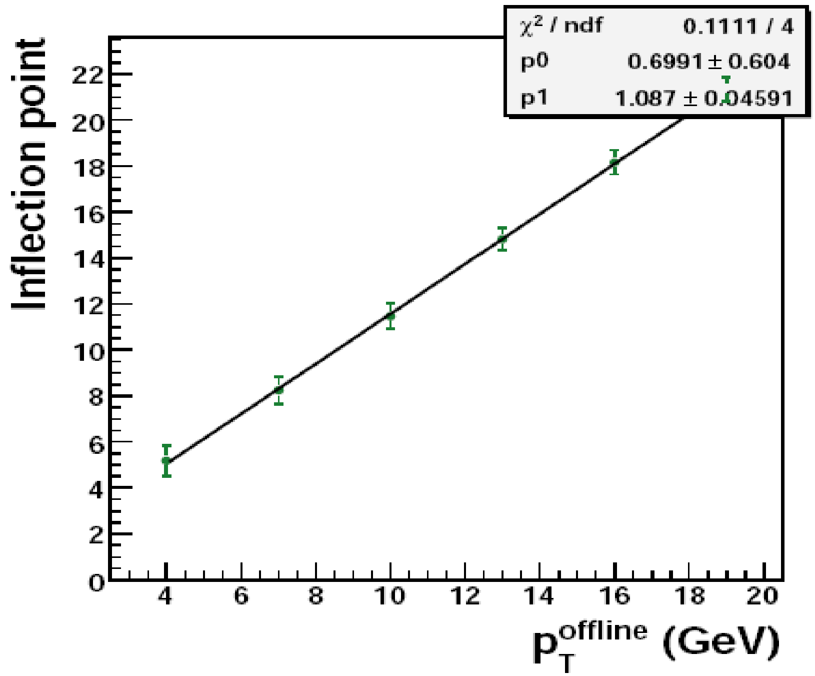

Trouble

- Z's don't yield many low p^T electrons
- For "loose" triggers, ok because it's about identifying very electromagnetic clusters
  - Use loose offline requirements
  - Doesn't really work for more sophisticated things like "isolated electrons"
  - Get wacky curve, or no stats
- At LHC, may not be an issue given high Z σ

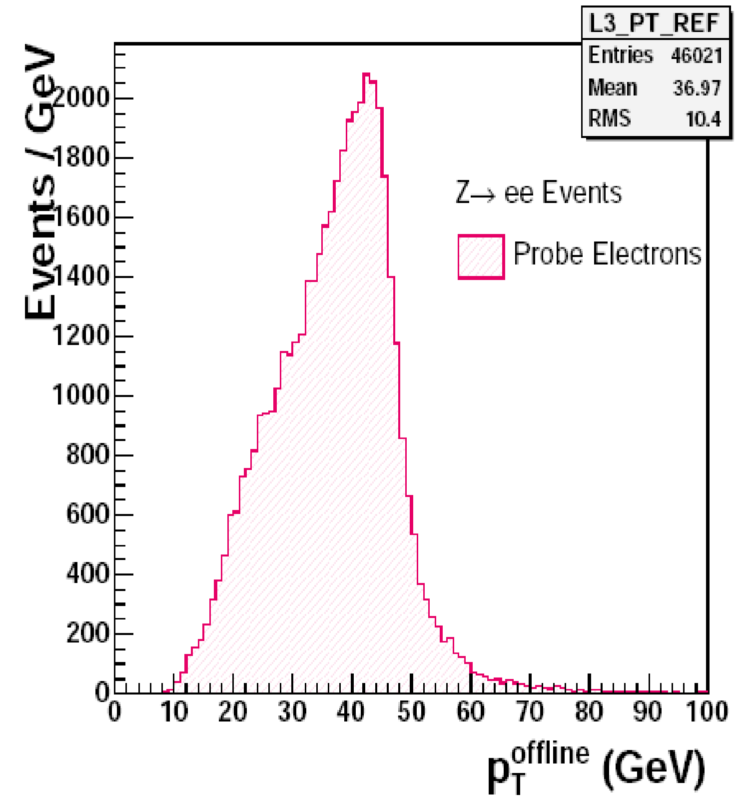

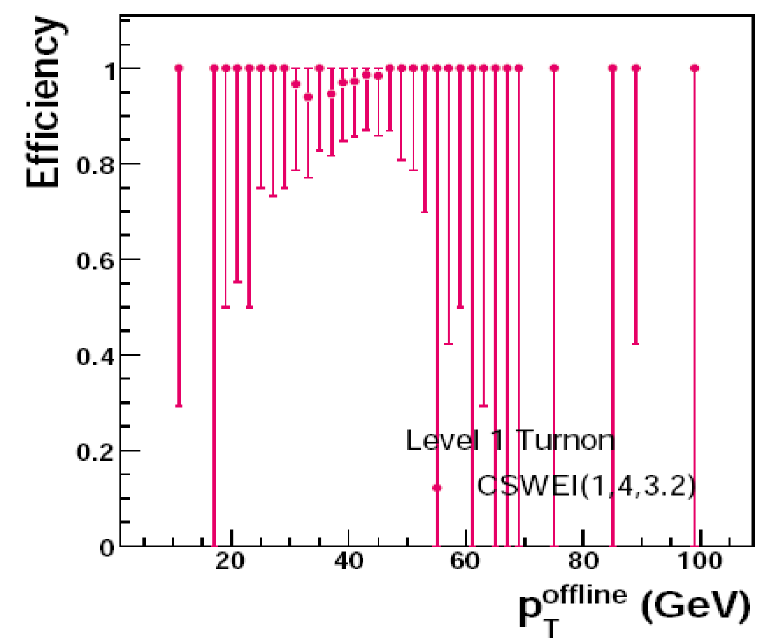

Muons

- Have J/ψ in addition to Z, so cover broader range easily
  - Thresholds tend to be substantially lower: once it's past the calorimeter, QCD rate much lower
    - Of course, assumes that beam-associated backgrounds are under control
    - If not, single muon triggers may be hard to maintain since p^T resolution at Level 1 typically not very good
    - Beam tunnel shielding critical
- Still substantial analysis work in extracting unbiased trigger efficiency...

Jets

- If you can't trigger on jets at a hadron collider, well, you should probably...
- Work is in setting thresholds and prescales to get reasonable samples at all E_T values

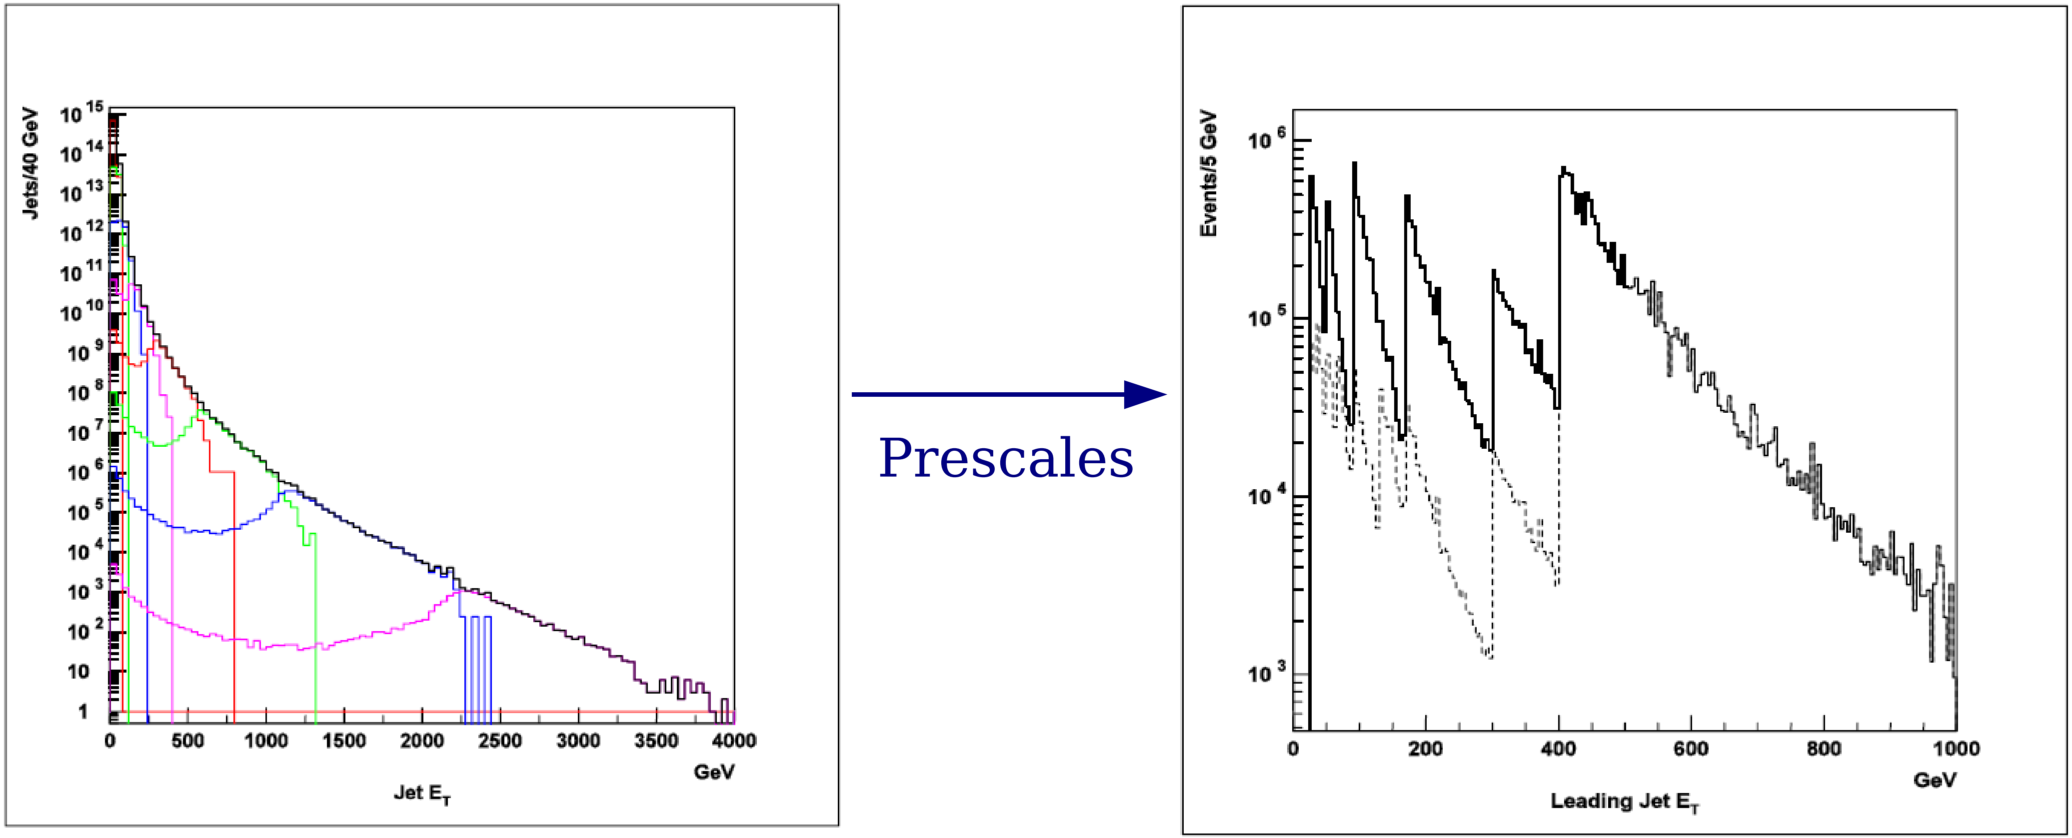

Impact of Calibration

- Calibration sharpens turn-on curve
  - Substantially reduces “garbage events”
- But... rates can change substantially
  - After all, most events at low end
  - Depending on direction of correction, rates can go up!
    - Then need to readjust thresholds

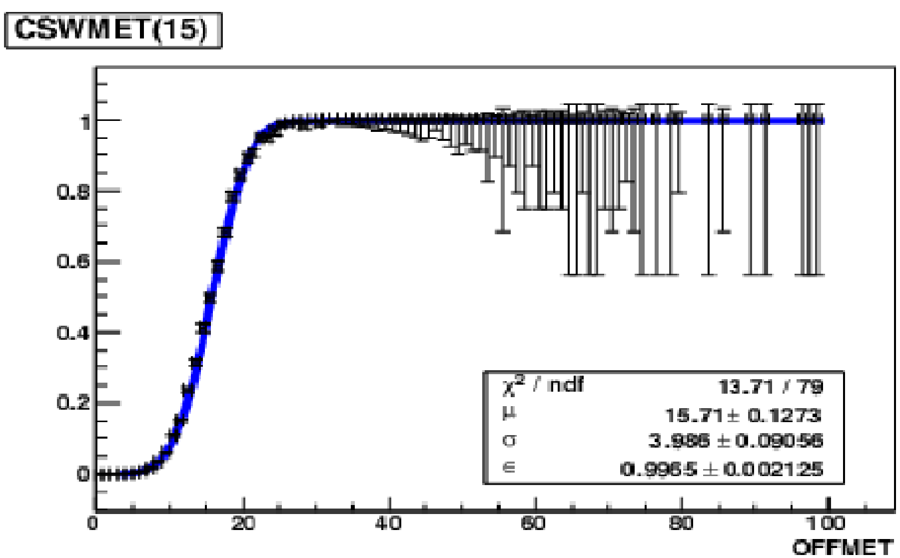

Rates
+~20%

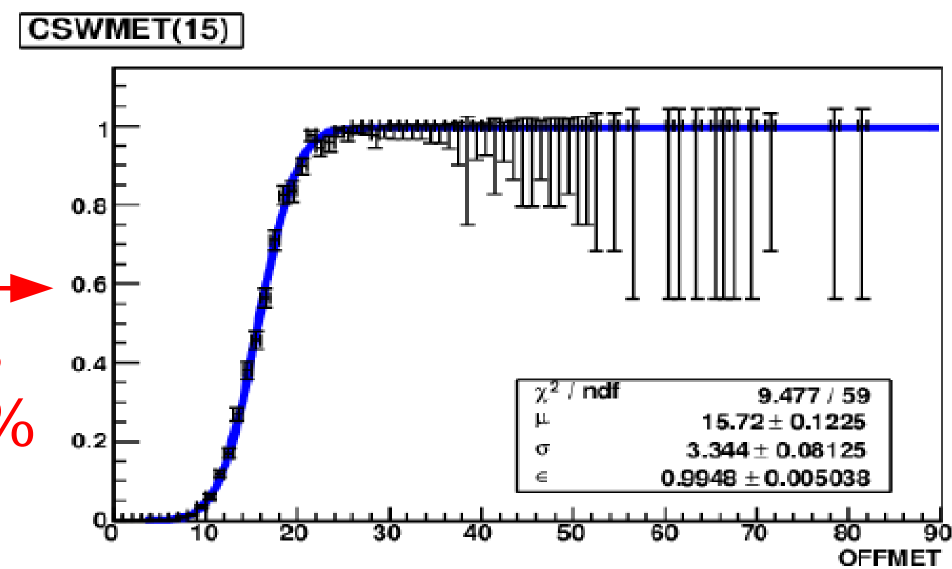

Online Data Quality

Online Data Quality

- Basic data quality monitoring consists in
  - Event display
    - Don't underestimate its value!
  - Occupancy plots
    - Not just geometrical, also timing, number of trigger objects, etc.
  - Full reconstruction
    - Can only be done for a small sample (reconstruction farms have thousands of nodes)
  - Slightly smarter:
    - Calorimeter occupancy for events with large MET, ...

Trigger and Data Quality

- In general, trigger is the first place to discover serious problems, since the system simply breaks down:
  - Hot channels leading to excessive rates
  - Dataflow issues
  - ...
- If excessive rates, correct diagnostic difficult
  - The system should be "throttling" itself using disables
    - Need to be able to see "true" firing rate
- If dataflow, in principle easy

- More subtle problems sometimes also visible
  - Rate "oscillations" primary example
    - Often not easy to track down
  - Occasional spikes
    - Almost impossible to track down
- Many interesting things to look at
  - Normalized trigger rate vs bunch number
  - ...
- More x-checks = more problems found

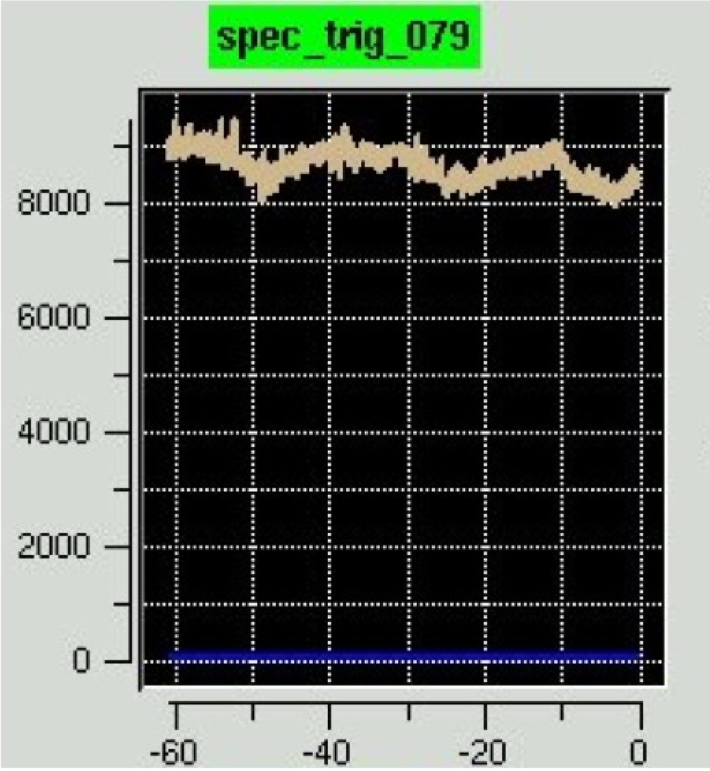

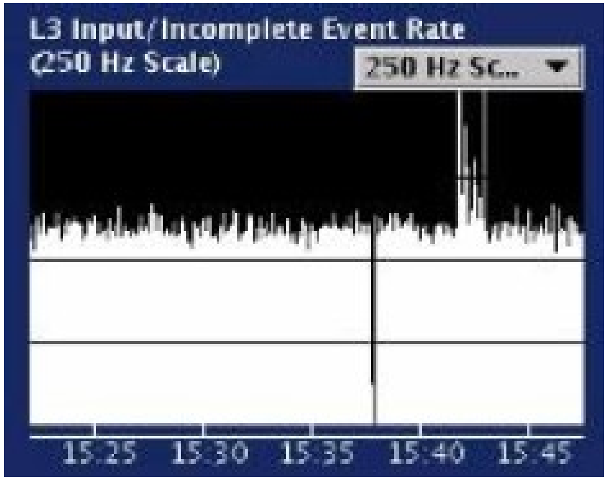

- But ultimately, most of the bad things that happen at < 1% rate are almost impossible to detect online
  - As long as you don't know which pattern to look for
    - Keep track of TGV schedules, popular TV programs, multitude of cron jobs set up by people, ...
  - Continuous feedback from analysis a necessity
    - The *really* subtle stuff may take years to find
- Doesn't mean you can't take good data starting on day 2
  - But detailed understanding takes lots of effort, and therefore time

Implications

- Since mostly only rather obvious problems can be found online, the rest will come offline
  - Feed back and fix the problem
  - If software solution possible (regularly true), need to “correct” the data *a posteriori*
- Note that while these become less frequent, don't completely go away
  - Detector “improvements” during downtimes
  - Some problems exposed by higher luminosities or rates
  - Global warming

Looking Forward

LHC Trigger Tables

- Don't really exist yet...
  - Studies being made, drafts with different complexity
  - In ATLAS, focus on tools (CMS, I don't know)
  - TDR tables exist
- Useful exercise for discussion:
  - Do rates make sense (back-of-an-envelope)?
  - What's clearly missing?
  - What seems strange?
  - Priorities?
  - Remember, it's about physics!

Some Guiding Numbers

- Cross-sections (1 $cm^{-2} = 10^{-24}$ barn):
  - top pairs ~800 pb, Z -> ee ~ 2 nb, W -> eν ~ 20 nb, QCD 40+ GeV ~ 75 μb, QCD 100+ GeV ~100 nb, Z' (1 TeV- SM-like) -> ee ~ 600 fb
- Bunch-crossing rate 15 x Tevatron
  - L1 accept ~50 kHz (vs 1-40 @ Tevatron)
    - Calorimeter and Muon only
  - L3 accept ~200 Hz (vs 100 @ Tevatron)
    - So ultimately, need ~7x better rejection
- Note: at DØ, about 85% of the "physics stream" is used in analysis one way or another

CMS L1

Numbers are for 2×10^{33}

Trigger	L1 Threshold (GeV)	L1 Rate (kHz)	Cumulative L1 Rate (kHz)
Inclusive $e\gamma$	22	3.9 ± 0.3	3.9 ± 0.3
Double $e\gamma$	11	1.0 ± 0.1	4.6 ± 0.3
Inclusive μ	14	2.5 ± 0.2	7.1 ± 0.3
Double μ	3	4.0 ± 0.3	11.0 ± 0.4
Inclusive τ	100	2.2 ± 0.2	12.9 ± 0.5
Double τ	60	3.0 ± 0.2	14.9 ± 0.5
1-,2-,3-,4-jets	150,100,70,50	2.2 ± 0.2	15.8 ± 0.5
H_{T}	275	2.0 ± 0.2	16.2 ± 0.5
$E_{\mathrm{T}}^{\mathrm{miss}}$	60	0.4 ± 0.1	16.3 ± 0.5
$H_{\mathrm{T}} + E_{\mathrm{T}}^{\mathrm{miss}}$	200, 40	1.1 ± 0.1	16.6 ± 0.5
jet + $E_{\mathrm{T}}^{\mathrm{miss}}$	100, 40	1.1 ± 0.1	16.7 ± 0.5
$\tau + E_{\mathrm{T}}^{\mathrm{miss}}$	60, 40	2.7 ± 0.2	18.8 ± 0.5
$\mu + E_{\mathrm{T}}^{\mathrm{miss}}$	5, 30	0.3 ± 0.1	19.0 ± 0.6
$e\gamma + E_{\mathrm{T}}^{\mathrm{miss}}$	15, 30	0.5 ± 0.1	19.1 ± 0.6
μ + jet	7, 100	0.2 ± 0.1	19.1 ± 0.6
$e\gamma$ + jet	15, 100	0.6 ± 0.1	19.2 ± 0.6
$\mu + \tau$	7, 40	1.2 ± 0.1	19.8 ± 0.6
$e\gamma + \tau$	15, 60	2.6 ± 0.2	20.5 ± 0.6
$e\gamma + \mu$	15, 7	0.2 ± 0.1	20.5 ± 0.6
Prescaled	See Tables 5-10		22.3 ± 0.6
Total L1 Rate			22.3 ± 0.6

CMS HLT

Numbers are for 2 x 10^{33}

Trigger	L1 bits used	L1 Prescale	HLT Threshold (GeV)	HLT Rate (Hz)
Inclusive e	2	1	26	23.5 ± 6.7
e-e	3	1	12, 12	1.0 ± 0.1
Relaxed e-e	4	1	19, 19	1.3 ± 0.1
Inclusive γ	2	1	80	3.1 ± 0.2
γ-γ	3	1	30, 20	1.6 ± 0.7
Relaxed γ-γ	4	1	30, 20	1.2 ± 0.6
Inclusive μ	0	1	19	25.8 ± 0.8
Relaxed μ	0	1	37	11.9 ± 0.5
μ-μ	1	1	7, 7	4.8 ± 0.4
Relaxed μ-μ	1	1	10, 10	8.6 ± 0.6
$\tau + E_T^{miss}$	10	1	65 (E_T^{miss})	0.5 ± 0.1
Pixel τ-τ	10, 13	1	—	4.1 ± 1.1
Tracker τ-τ	10, 13	1	—	6.0 ± 1.1
$\tau + e$	26	1	52, 16	< 1.0
$\tau + \mu$	0	1	40, 15	< 1.0
b-jet (leading jet)	36, 37, 38, 39	1	350, 150, 55 (see text)	10.3 ± 0.3
b-jet (2^{nd} leading jet)	36, 37, 38, 39	1	350, 150, 55 (see text)	8.7 ± 0.3
Single-jet	36	1	400	4.8 ± 0.0
Double-jet	36, 37	1	350	3.9 ± 0.0
Triple-jet	36, 37, 38	1	195	1.1 ± 0.0
Quadruple-jet	36, 37, 38, 39	1	80	8.9 ± 0.2
E_T^{miss}	32	1	91	2.5 ± 0.2
jet + E_T^{miss}	32	1	180, 80	3.2 ± 0.1
acoplanar 2 jets	36, 37	1	200, 200	0.2 ± 0.0
acoplanar jet + E_T^{miss}	32	1	100, 80	0.1 ± 0.0
2 jets + E_T^{miss}	32	1	155, 80	1.6 ± 0.0
3 jets + E_T^{miss}	32	1	85, 80	0.9 ± 0.1
4 jets + E_T^{miss}	32	1	35, 80	1.7 ± 0.2
Diffractive	See Ref. [10]	1	40, 40	< 1.0
$H_T + E_T^{miss}$	31	1	350, 80	5.6 ± 0.2
$H_T + e$	31	1	350, 20	0.4 ± 0.1
Inclusive γ	2	400	23	0.3 ± 0.0
γ-γ	3	20	12, 12	2.5 ± 1.4
Relaxed γ-γ	4	20	19, 19	0.1 ± 0.0
Single-jet	33	10	250	5.2 ± 0.0
Single-jet	34	1 000	120	1.6 ± 0.0
Single-jet	35	100 000	60	0.4 ± 0.0
Total HLT rate				119.3 ± 7.2

ATLAS: Signatures

Selection signature	Examples of physics coverage
e25i	$W \to e\nu$, $Z \to ee$, top production, $H \to WW^{(*)}/ZZ^{(*)}$, W',Z'
2e15i	$Z \to ee$, $H \to WW^{(*)}/ZZ^{(*)}$
μ20i	$W \to \mu\nu$, $Z \to \mu\mu$, top production, $H \to WW^{(*)}/ZZ^{(*)}$, W',Z'
2μ10	$Z \to \mu\mu$, $H \to WW^{(*)}/ZZ^{(*)}$
γ60i	direct photon production, $H \to \gamma\gamma$
2γ20i	$H \to \gamma\gamma$
j400	QCD, SUSY, new resonances
2j350	QCD, SUSY, new resonances
3j165	QCD, SUSY
4j110	QCD, SUSY
τ60i	charged Higgs
μ10 + e15i	$H \to WW^{(*)}/ZZ^{(*)}$, SUSY
τ35i + xE45	$qqH(\tau\tau)$, $W \to \tau\nu$, $Z \to \tau\tau$, SUSY at large tan β
j70 + xE70	SUSY
xE200	new phenomena
E1000	new phenomena
jE1000	new phenomena
2μ6 + $\mu^+\mu^-$ + mass cuts	rare b-hadron decays ($B \to \mu\mu X$) and $B \to J/\psi$ (ψ')X

ATLAS: L1 Rates @ 2 x 10^{33}

LVL1 signature	Rate (kHz)
EM25I	12.0
2EM15I	4.0
MU20	0.8
2MU6	0.2
J200	0.2
3J90	0.2
4J65	0.2
J60+XE60	0.4
TAU25I+XE30	2.0
MU10+EM15I	0.1
Others (prescaled, exclusive, monitor, calibration)	5.0
Total	~25.0

ATLAS: HLT Rates @ 2 x 10^{33}

HLT signature	Rate (Hz)
e25i	40
2e15i	< 1
γ60i	25
2γ20i	2
μ20i	40
2μ10	10
j400	10
3j165	10
4j110	10
j70+xE70	20
τ35i+xE45	5
2μ6 with vertex, decay-length and mass cuts (J/ψ, ψ', B)	10
Others (prescaled, exclusive, monitor, calibration)	20
Total	**~200**

Concluding Remarks

- Modern detectors are way to complicated for every individual to know everything in the smallest details, but:
  - You should know how the data flows through your experiment
  - What the trigger algorithms do to data used in your analysis
  - What the strengths of your experiment are
    - Which means knowing about other experiments
  - What typical detector resolutions are
  - ...

- Trigger systems for hadron collider experiments are diverse, complex, technologically advanced, but wonderful things...

Now, why wasn't this school held in the Virgin Islands?

- Many thanks to people for inputs (often without their explicit knowledge):
  - Many DØ people whose plots I copied
    - Marc Hohlfeld, Samuel Calvet, Bertrand Martin, Remi Mommsen
  - Tom Lecompte
  - Dave Waters
  - All the people that designed/put together all these trigger systems