SLHC Trigger & DAQ

Wesley H. Smith

U. Wisconsin - Madison

FNAL Forward Pixel SLHC Workshop

October 9, 2006

Outline:

SLHC Machine, Physics, Trigger & DAQ

Impact of Luminosity up to 10^{35}

Calorimeter, Muon & Tracking Triggers

DAQ requirements & upgrades

This talk is available on:

http://cmsdoc.cern.ch/cms/TRIDAS/tr/06/10/smith_slhc_fnal_oct06.pdf

LHC Trigger & DAQ Challenges

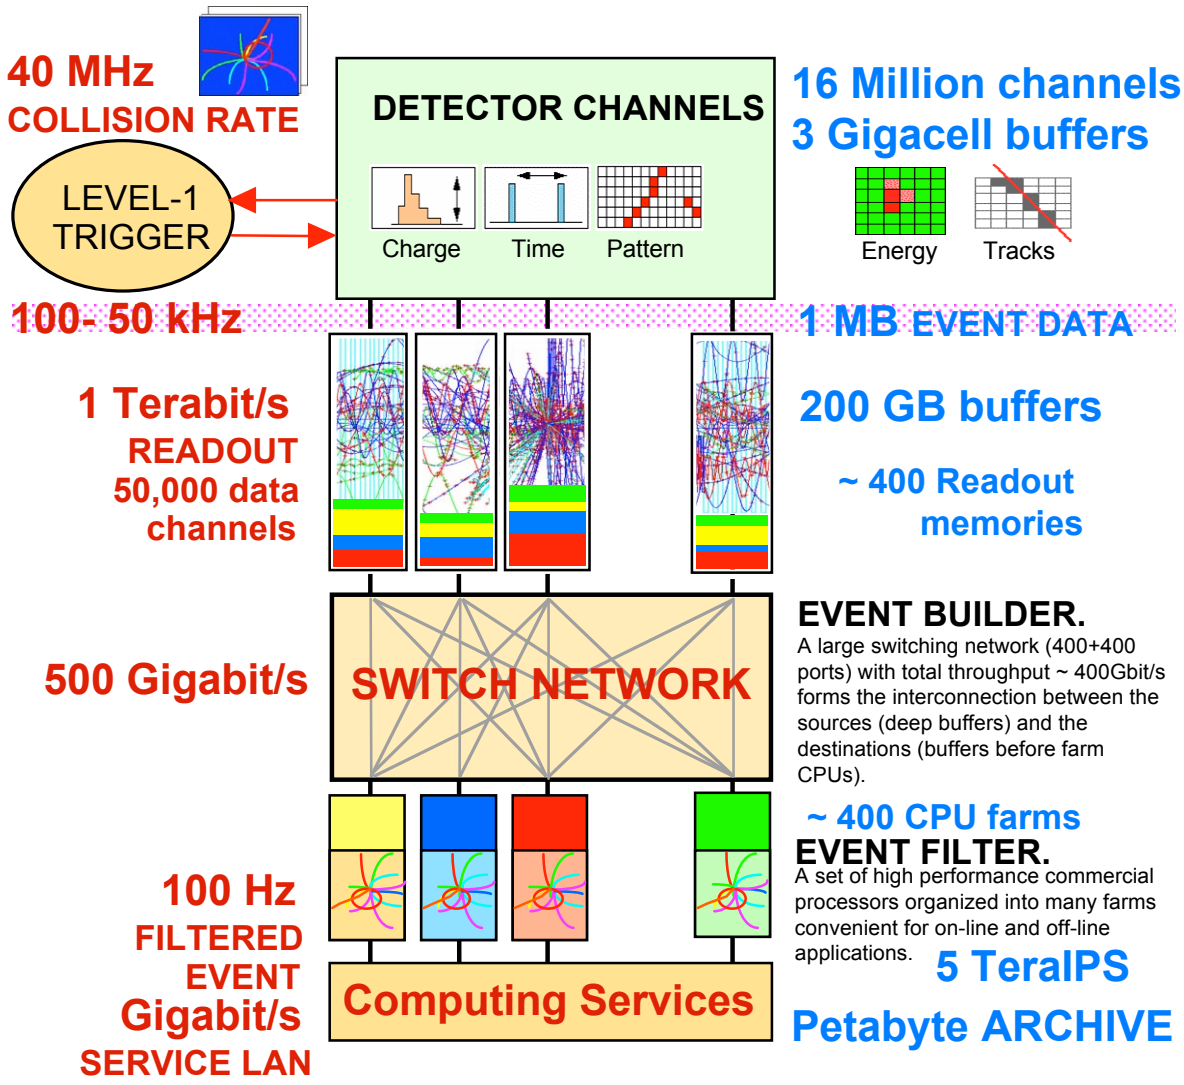

Challenges:

1 GHz of Input Interactions

Beam-crossing every 25 ns with ~17 interactions produces over 1 MB of data

Archival Storage at about 100 Hz of 1 MB events

Level 1 Trigger Operation

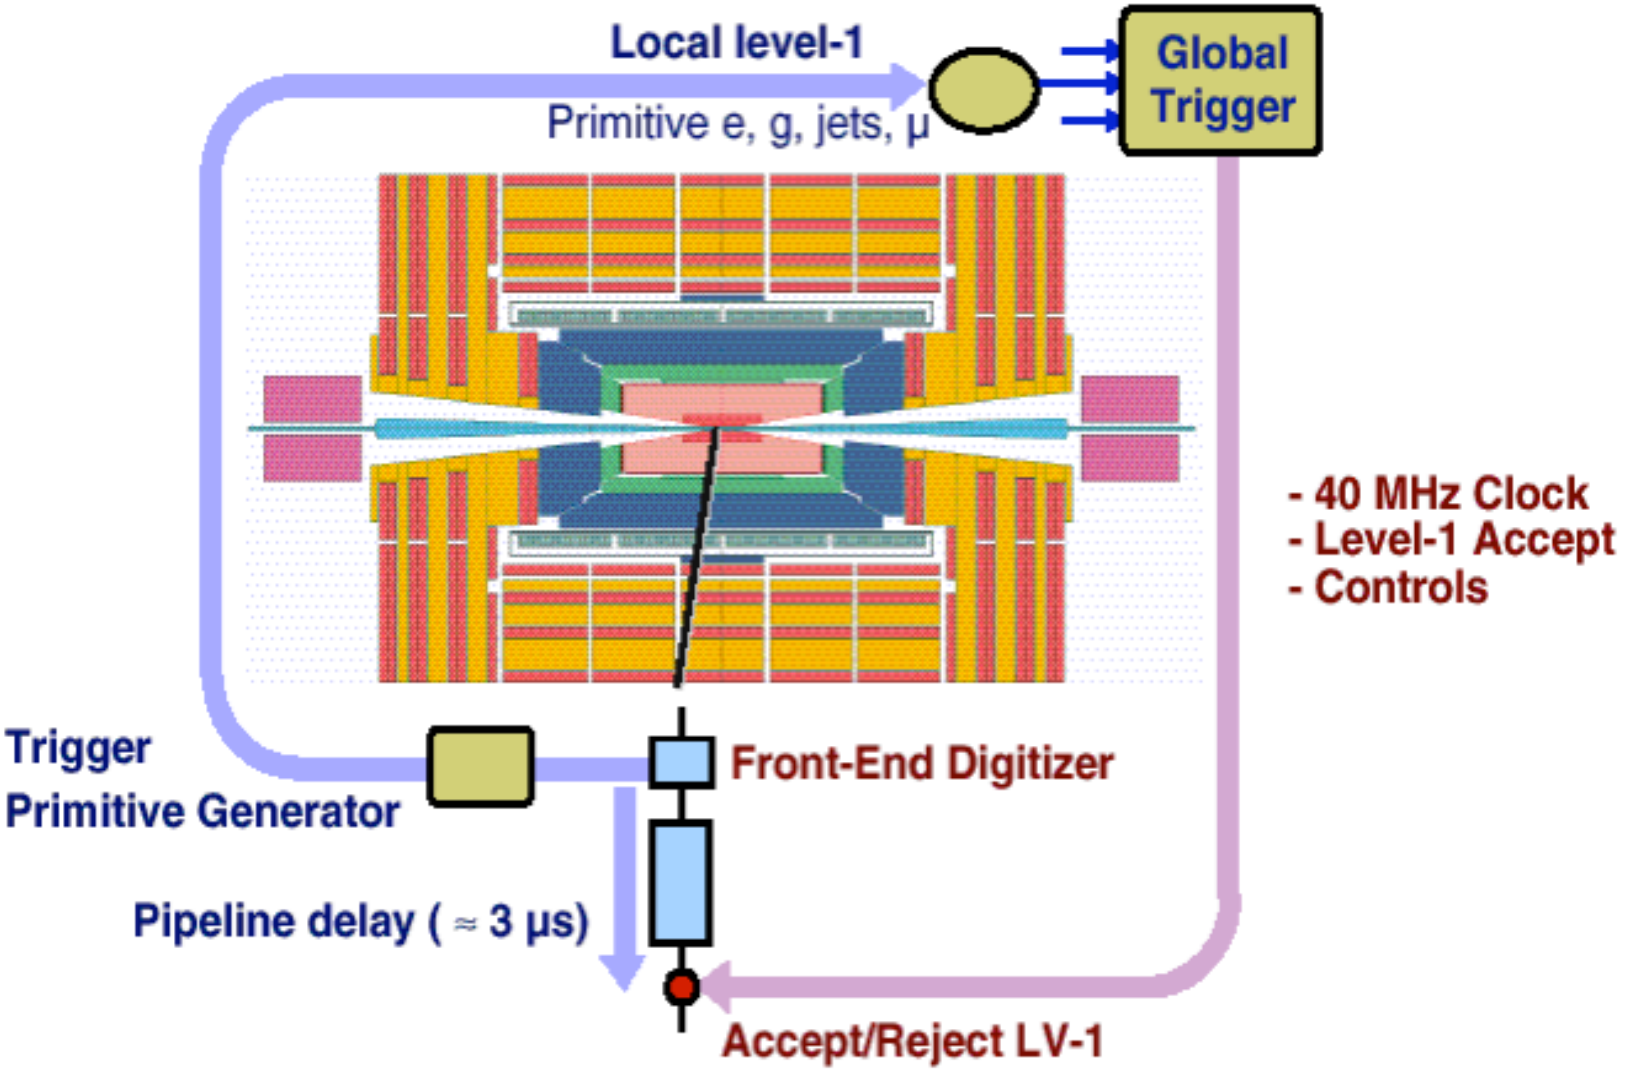

CMS Level-1 Trigger & DAQ

Overall Trigger & DAQ Architecture: 2 Levels:

Level-1 Trigger:

- 25 ns input
- 3.2 μs latency

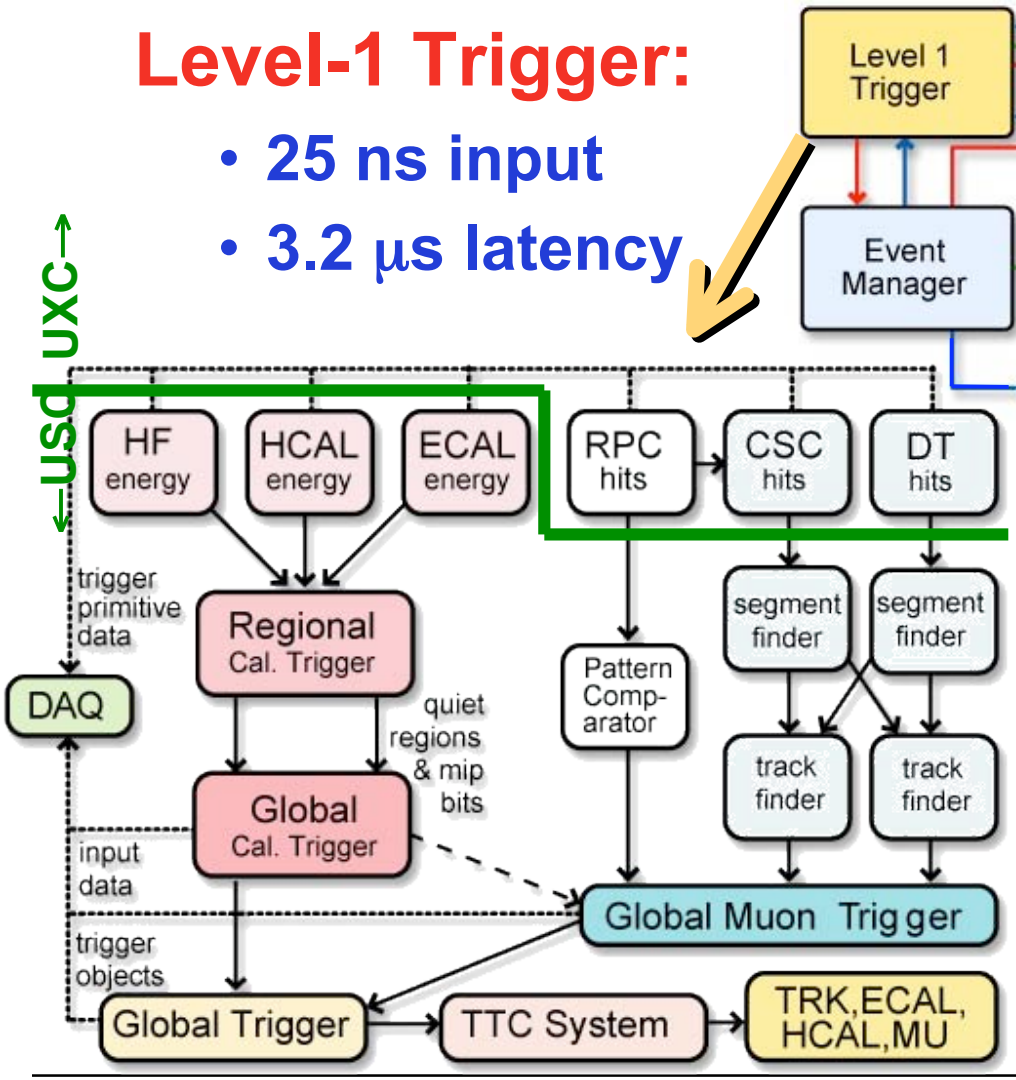

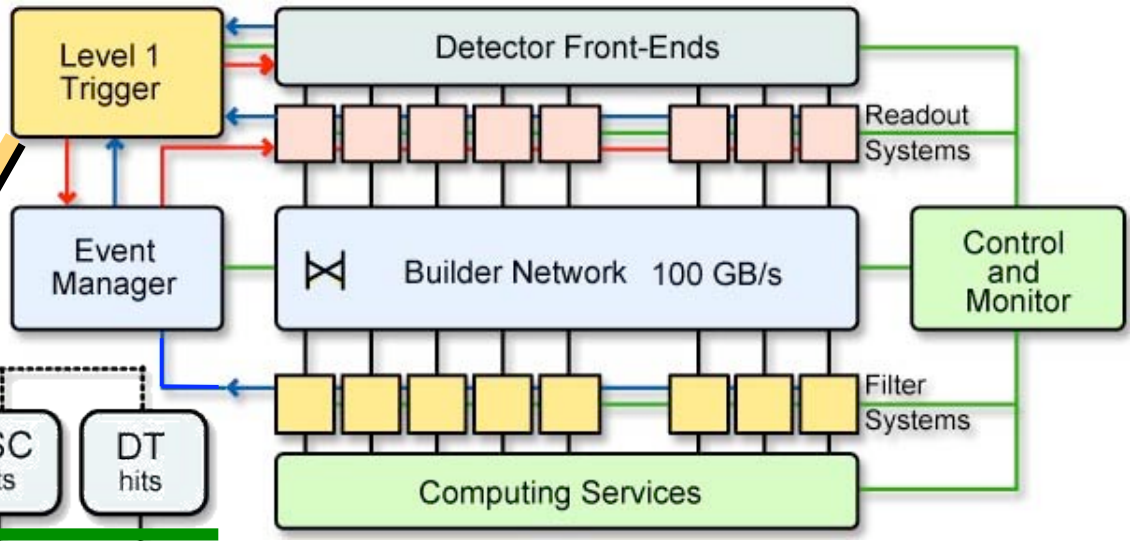

Interaction rate: 1 GHz

Bunch Crossing rate: 40 MHz

Level 1 Output: 100 kHz (50 initial)

Output to Storage: 100 Hz

Average Event Size: 1 MB

Data production 1 TB/day

Baseline (S)LHC Parameters

parameter	symbol	nominal LHC	ultimate LHC	shorter bunches	
					←SLHC
#bunches	n_b	2808	2808	5616	
protons/bunch	N_b [10^{11}]	1.15	1.7	1.7	
bunch spacing	Δt_{sep} [ns]	25	25	12.5	←25 ns → 12.5 ns
average current	I [A]	0.58	0.86	1.72	
norm. transv. emittance	ε_n [μm]	3.75	3.75	3.75	
longit. profile		Gaussian	Gaussian	Gaussian	
rms b. length	σ_z [cm]	7.55	7.55	3.78	
beta at IP1&IP5	β^* [m]	0.55	0.5	0.25	
crossing angle	θ_c [μrad]	285	315	445	
Piwinski parameter	$\theta_c\sigma_z/(\sigma^*$ 2)	0.64	0.75	0.75	
luminosity	L [10^{34} $cm^{-2}s^{-1}$]	1.0	2.3	9.2	←10^{34} → 10^{35}
events/ crossing		19	44	88	← pileup x 5
length luminous region (rms)	σ_{lum} [mm]	44.9	42.8	21.8	

Detector Luminosity Effects

$H \rightarrow ZZ \rightarrow \mu\mu ee$, M_H= 300 GeV for different luminosities in CMS

10^{32} cm^{-2}s^{-1}

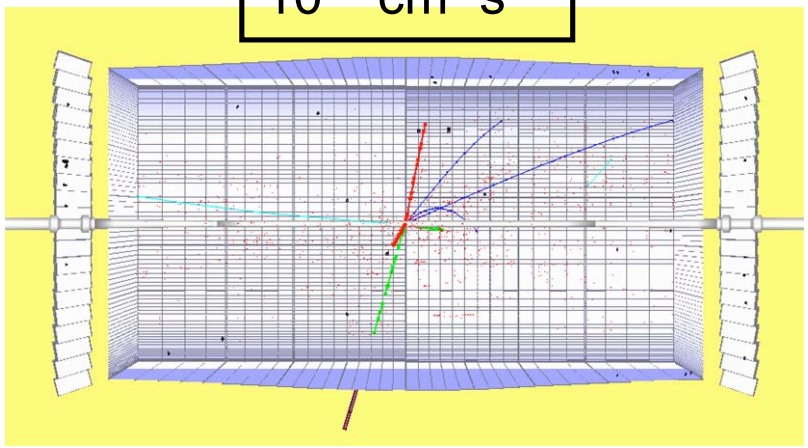

10^{33} cm^{-2}s^{-1}

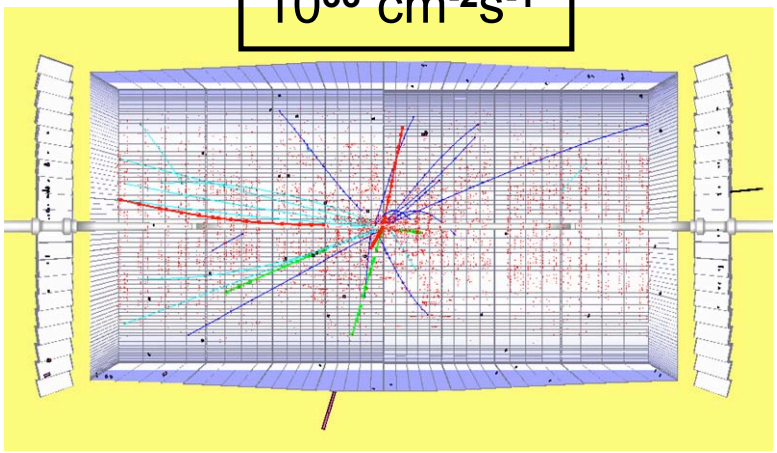

10^{34} cm^{-2}s^{-1}

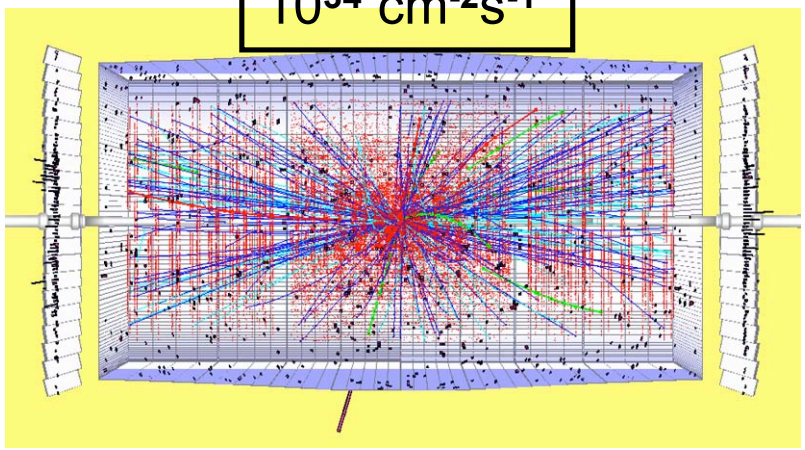

10^{35} cm^{-2}s^{-1}

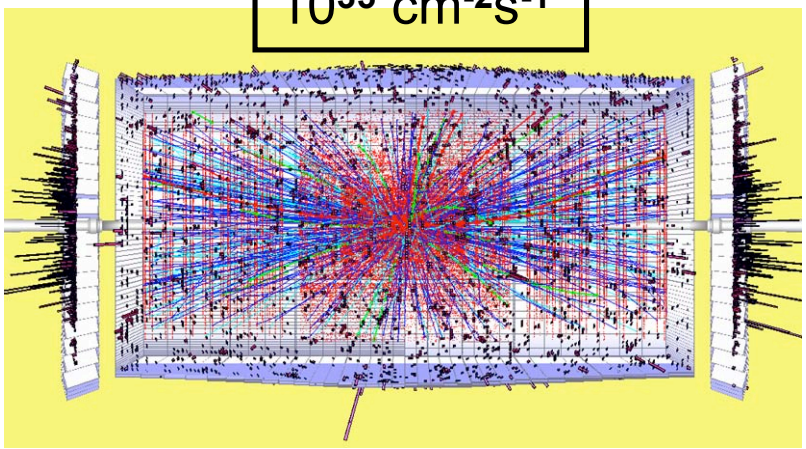

SLHC Level-1 Trigger @ 10^{35}

Occupancy

- **Degraded performance of algorithms**
  - Electrons: reduced rejection at fixed efficiency from isolation
  - Muons: increased background rates from accidental coincidences
- **Larger event size to be read out**
  - New Tracker: higher channel count & occupancy → large factor
  - Reduces the max level-1 rate for fixed bandwidth readout.

Trigger Rates

- **Try to hold max L1 rate at 100 kHz by increasing readout bandwidth**
  - Avoid rebuilding front end electronics/readouts where possible
    - **Limits: ⟨readout time⟩ (< 10 μs) and data size (total now 1 MB)**
  - Use buffers for increased latency for processing, not post-L1A
  - May need to increase L1 rate even with all improvements
    - **Greater burden on DAQ**
- **Implies raising E_T thresholds on electrons, photons, muons, jets and use of less inclusive triggers**
  - Need to compensate for larger interaction rate & degradation in algorithm performance due to occupancy

Radiation damage -- Increases for part of level-1 trigger located on detector

SLHC Trigger @ 12.5 ns

Choice of 80 MHz

- Reduce pile-up, improve algorithm performance, less data volume for detectors that identify 12.5 ns BX data
- Retain front-end electronics since 40 MHz sampling in phase
  - Not true for 10 ns or 15 ns bunch separation -- large cost
- Be prepared for LHC Machine group electron-cloud solution
- Retain ability to time-in experiment
  - Beam structure vital to time alignment
- Higher frequencies ~ continuous beam

Rebuild level-1 processors to use data "sampled" at 80 MHz

- Already ATLAS & CMS have internal processing up to 160 MHz and higher in a few cases
- Use 40 MHz sampled front-end data to produce trigger primitives with 12.5 ns resolution
  - e.g. cal. time res. < 25 ns, pulse time already from multiple samples
- Save some latency by running all trigger systems at 80 MHz I/O
- Technology exists to handle increased bandwidth

SLHC Trigger Requirements

High-P_T discovery physics

- Not a big rate problem since high thresholds

Completion of LHC physics program

- Example: precise measurements of Higgs sector
- Require low thresholds on leptons/photons/jets
  - Use more exclusive triggers since final states will be known

Control & Calibration triggers

- W, Z, Top events
- Low threshold but prescaled

SLHC Level-1 Trigger Menu

ATLAS/CMS Studies in hep-ph/0204087:

- inclusive single muon p_T > 30 GeV (rate ~ 25 kHz)
- inclusive isolated e/γ E_T > 55 GeV (rate ~ 20 kHz)
- isolated e/γ pair E_T > 30 GeV (rate ~ 5 kHz)
  - or 2 different thresholds (i.e. 45 & 25 GeV)
- muon pair p_T > 20 GeV (rate ~ few kHz?)
- jet E_T > 150 GeV.AND.E_T(miss) > 80 GeV (rate ~ 1-2 kHz)
- inclusive jet trigger E_T > 350 GeV (rate ~ 1 kHz)
- inclusive E_T(miss) > 150 GeV (rate ~1 kHz);
- multi-jet trigger with thresholds determined by the affordable rate

Trig. Primitives: CMS Calorimeter

HF:Quartz Fiber: Possibly replaced

- Very fast - gives good BX ID
- Modify logic to provide finer-grain information
  - Improves forward jet-tagging

HCAL:Scintillator/Brass: Barrel stays but endcap replaced

- Has sufficient time resolution to provide energy in correct 12.5 ns BX with 40 MHz sampling. Readout may be able to produce 80 MHz already.

ECAL: $PBWO_4$ Crystal: Stays

- Also has sufficient time resolution to provide energy in correct 12.5 ns BX with 40 MHz sampling, may be able to produce 80 MHz output already.
- Exclude on-detector electronics modifications for now -- difficult:
  - Regroup crystals to reduce $\Delta\eta$ tower size -- minor improvement
  - Additional fine-grain analysis of individual crystal data -- minor improvement

Conclusions:

- Front end logic same except where detector changes
- Need new TPG logic to produce 80 MHz information
- Need higher speed links for inputs to Cal Regional Trigger

Trig. Prim.: CMS Endcap Muon

- D. Acosta

4 stations of CSCs: Bunch Crossing ID at 12.5 ns:

- Use second arriving segment to define track BX
  - Use a 3 BX window
- Improve BX ID efficiency to 95% with centered peak, taking 2nd Local Charged Track, requiring 3 or more stations
  - Requires 4 stations so can require 3 stations at L1
- Investigate improving CSC performance: HV, Gas, …
  - If 5 ns resolution ⇒ 4 ns, BX ID efficiency might climb to 98%

Occupancy at 80 MHz: Local Charged Tracks found in each station

- Entire system: 4.5 LCTs /BX
- Worst case: inner station: 0.125/BX (others 3X smaller)
- P(≥ 2) = 0.7% (spoils di-μ measurement in single station)
- Conclude: not huge, but neglected neutrons and ghosts may be under-estimated⇒ need to upgrade trigger front end to transmit LCT @ 80 MHz

Occupancy in Track-Finder at 80 MHz:

- Using 4 BX window, find 0.5/50 ns in inner station (every other BX at 25 ns!)
  - ME2-4 3X smaller, possibly only need 3 BX
- Need studies to see if these tracks generate triggers

Trig Primitives: CMS DT & RPC

DT:

- Operates at 40 MHz in barrel
- Could produce results for 80 MHz with loss of efficiency…or…
- Could produce large rate of lower quality hits for 80 MHz for combination with a tracking trigger with no loss of efficiency

RPC:

- Operates at 40 MHz
- Could produce results with 12.5 ns window with some minor external changes.
- Uncertain if RPC can operate at SLHC rates, particularly in the endcap

CMS SLHC L-1 Tracking Trigger
Ideas & Implications for L-1

THE UNIVERSITY of WISCONSIN MADISON

Additional Component at Level-1

- **Actually, CMS could have a rudimentary L-1 Tracking Trigger**
  - Pixel z-vertex in $\Delta\eta \times \Delta\phi$ bins can reject jets from pile-up
    - **Cable not hooked up in final version**
- **SLHC Track Trigger could provide outer stub and inner track**
  - Combine with cal at L-1 to reject π^0 electron candidates
  - Reject jets from other crossings by z-vertex
  - Reduce accidentals and wrong crossings in muon system
  - Provide sharp P_T threshold in muon trigger at high P_T
- **Cal & Muon L-1 output needs granularity & info. to combine w/ tracking trig. Also need to produce hardware to make combinations**

Move some HLT algorithms into L-1 or design new algorithms reflecting tracking trigger capabilities

- Local track clusters from jets used for 1st level trigger signal → jet trigger with σ_z = 6mm!
- Program in Readout Chip track cluster multiplicity for trigger output signal
- Combine in Module Trigger Chip (MTC) 16 trig. signals & decide on module trigger output

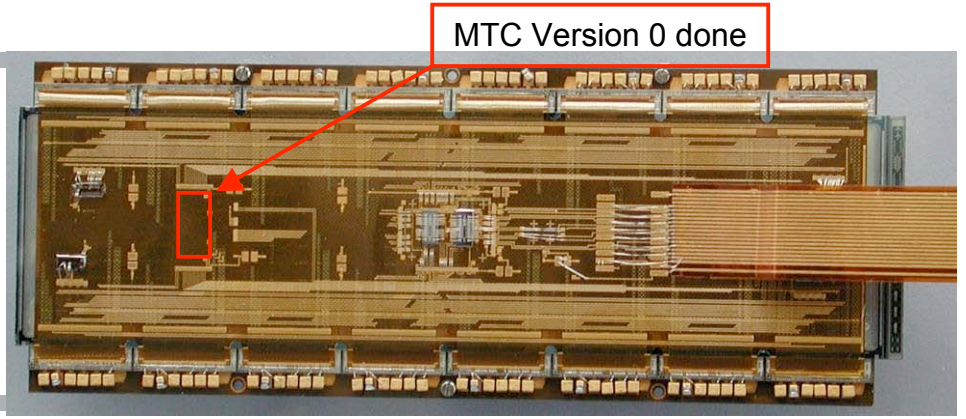

CMS ideas for trigger-capable tracker modules -- very preliminary

- Use close spaced stacked pixel layers
- Geometrical p_T cut on data (e.g. ~ GeV):
- Angle (γ) of track bisecting sensor layers defines p_T ($\Rightarrow$ window)
- For a stacked system (sepn. ~1mm), this is ~1 pixel
- Use simple coincidence in stacked sensor pair to find tracklets
- More details & implementation next slides

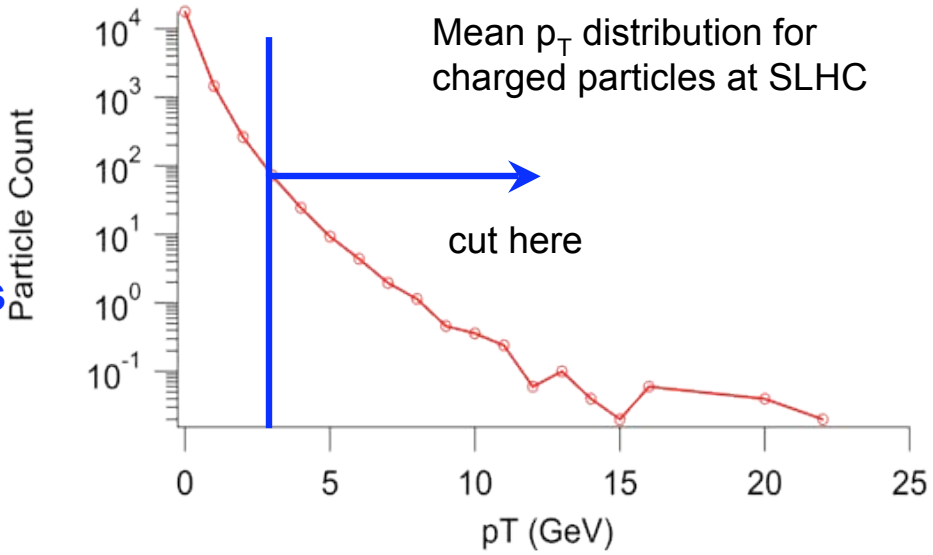

-- C. Foudas & J. Jones

A track like this wouldn't trigger:

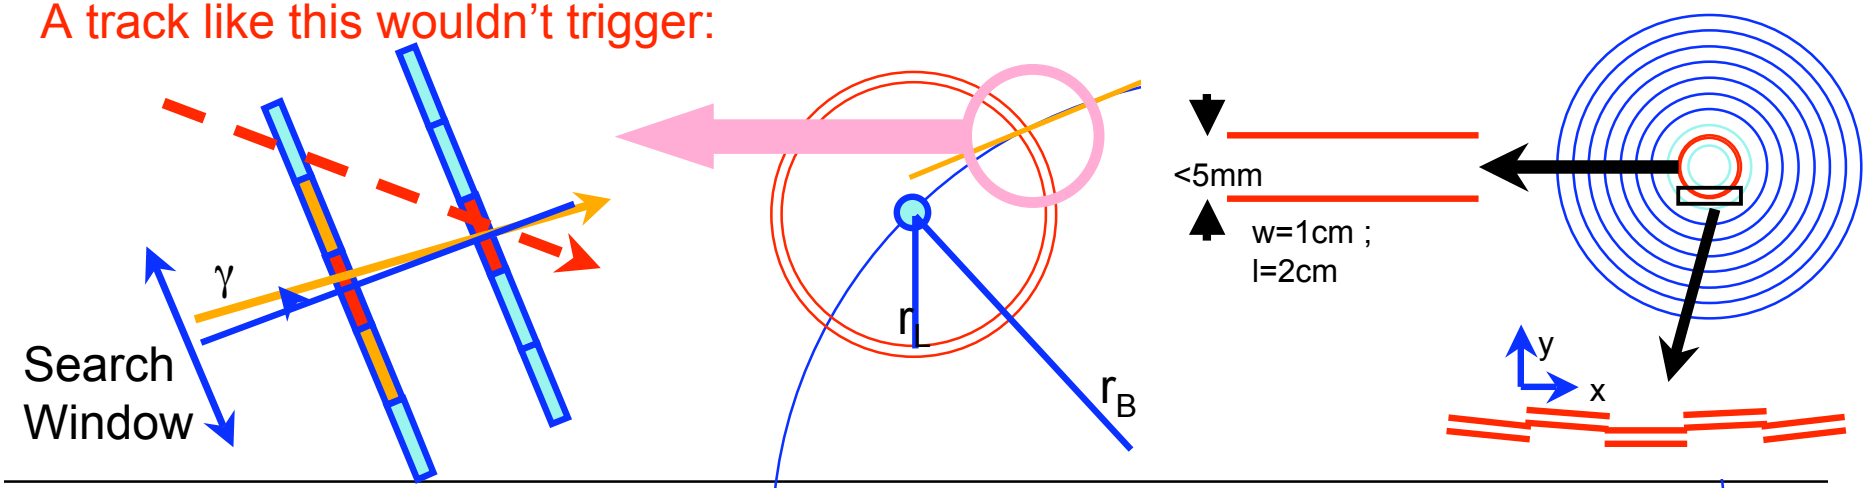

SLHC Tracker Layout

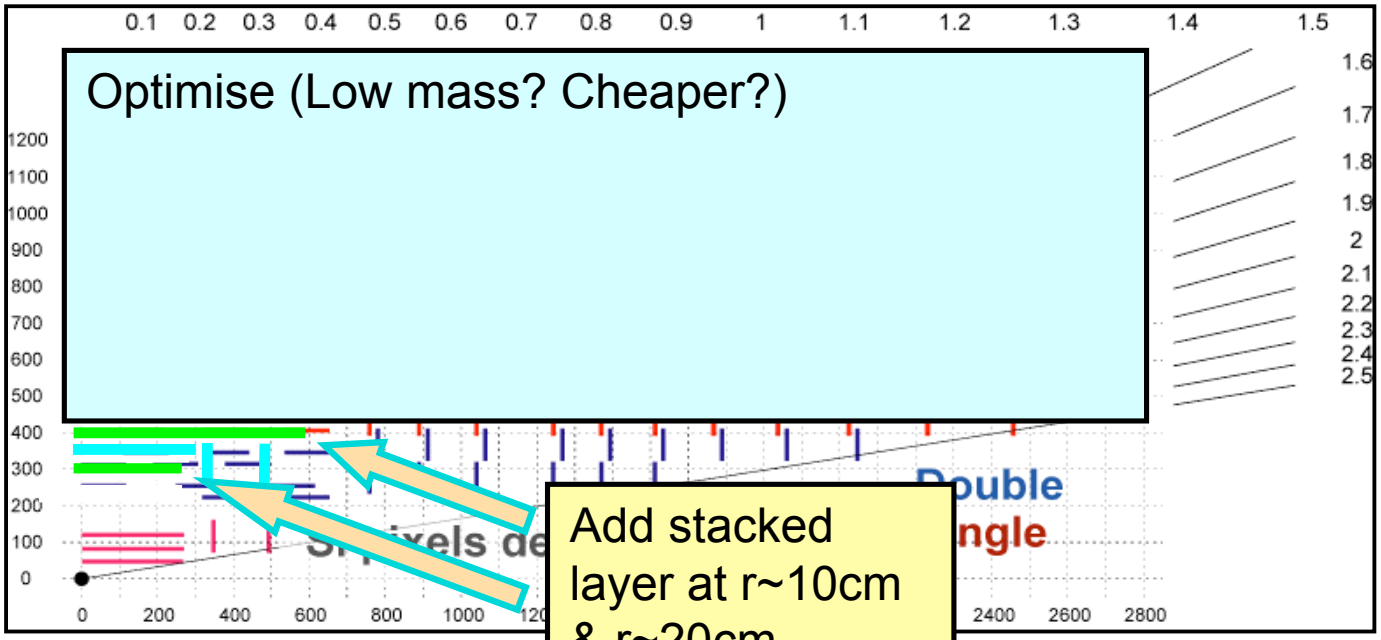

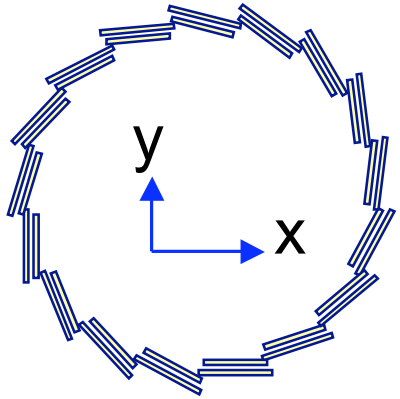

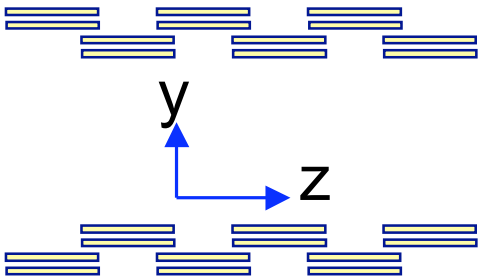

Detector Dimensions:
120cm(z)x20cm(r) & 60cm(z)x10cm(r)

Pixel Pitch: 10μm(r)x200μm(z)x20μm(ϕ)

Data Rate for SLHC pixels @ r = 10 cm

Pixel occupancy in SLHC ~ 4 hits / $(1.28\text{cm})^2$ @ 80 MHz BX (or 8 @ 40 MHz)

Assume 20-bit pixel coding scheme (1024x1024 array)

Base data rate is 80x106 x 4 x 20 / (1.28)2 = 3.9 Gbit/cm^2/s

BUT have ignored:

- Charge sharing → x2
- Error correction on optical links (Hamming coding / 8b10b) → x1.25

Should add margin (~ 20%) for e.g. data coding overheads

3.125 x 2 x 1.25 x 1.2 = ~12 Gbit/cm^2/s

Difficult to implement

- Power, cabling, etc…..

Charged Particles vs. p_T

- J. Jones

Mean p_t distribution for charged particles at SLHC

- Pythia 6.2772; 10,000 min. bias events
- CMKIN 4.2, standard datacard

Cut at ~ GeV ⇒removes much background

- But… minbias doesn't include high-p_T leptons & LO QCD?

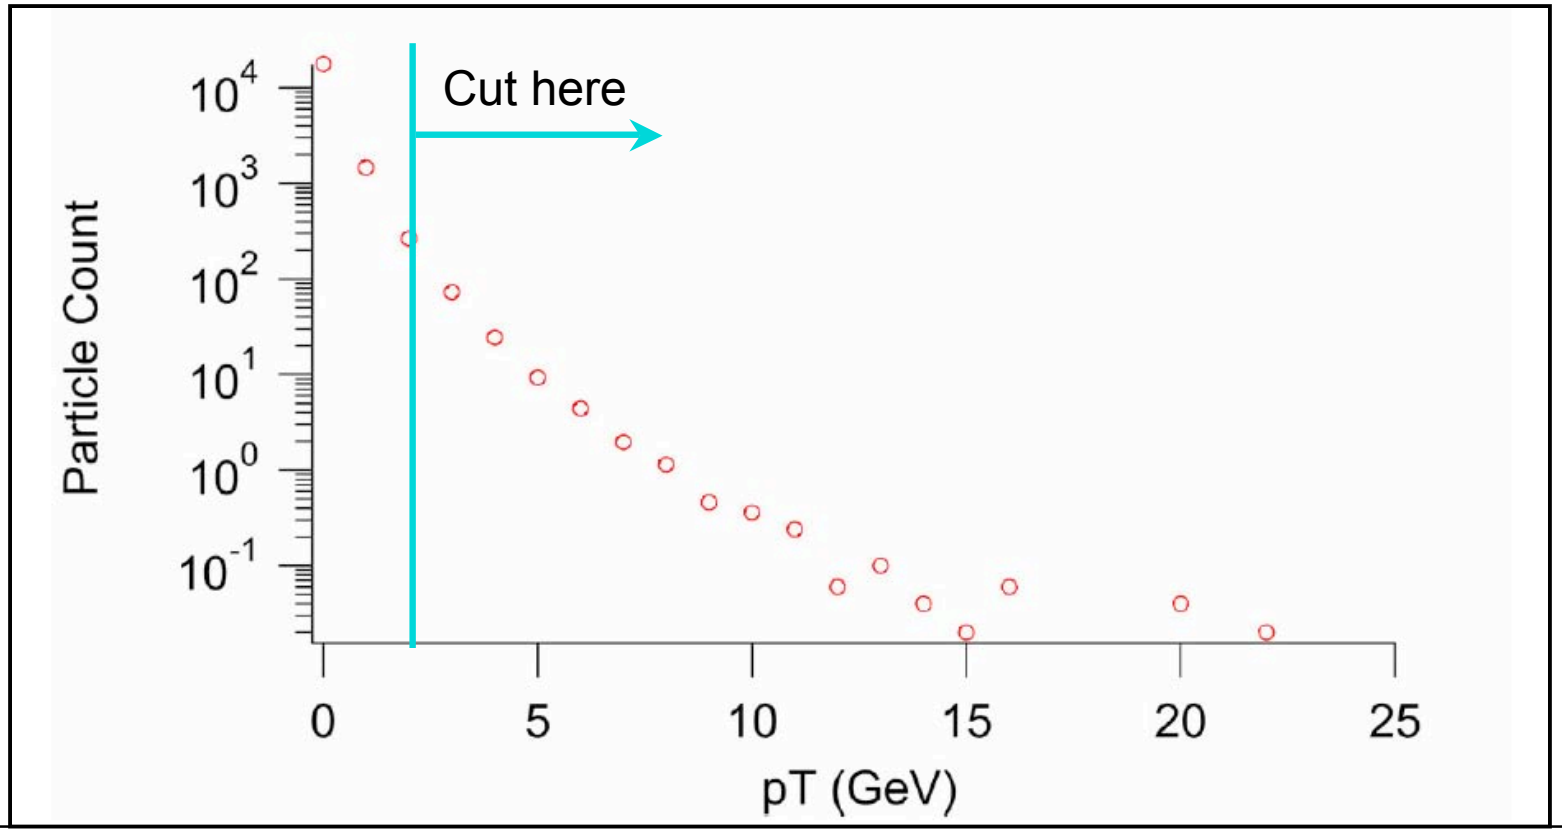

Tangent-Point Reconstruction

- Assume IP r=0
- Angle α determines p_T of track

Smaller α = greater p_T

- Can find high-p_T tracks by looking for small angular separation of hits in the two layers
- Correlation is fairly 'pure' provided separation is small and pixel pitch is small

Matching hits tend to be from the same track

- If sensors are precisely aligned, column number for hit pixels in each layer can be compared
- Finding high-p_T tracks becomes a relatively simple difference analysis

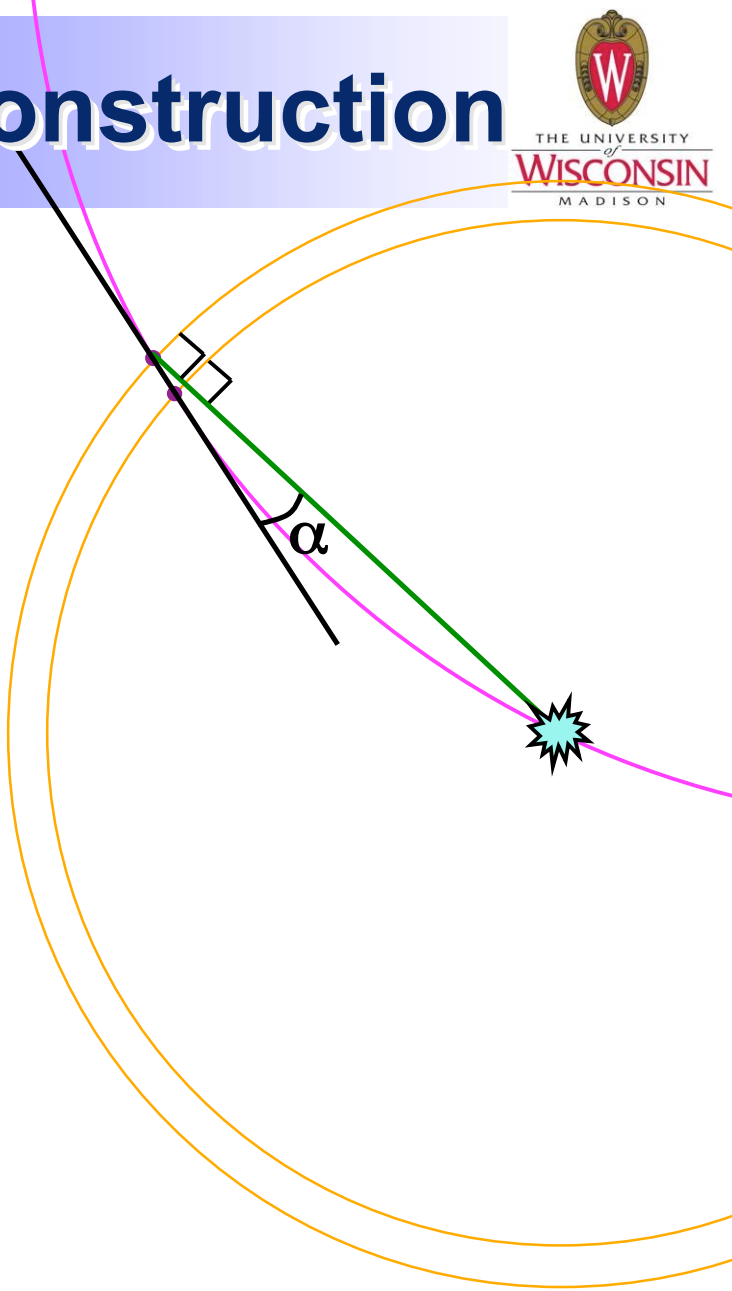

Difference Analysis in Practice

- Nearest-neighbor example

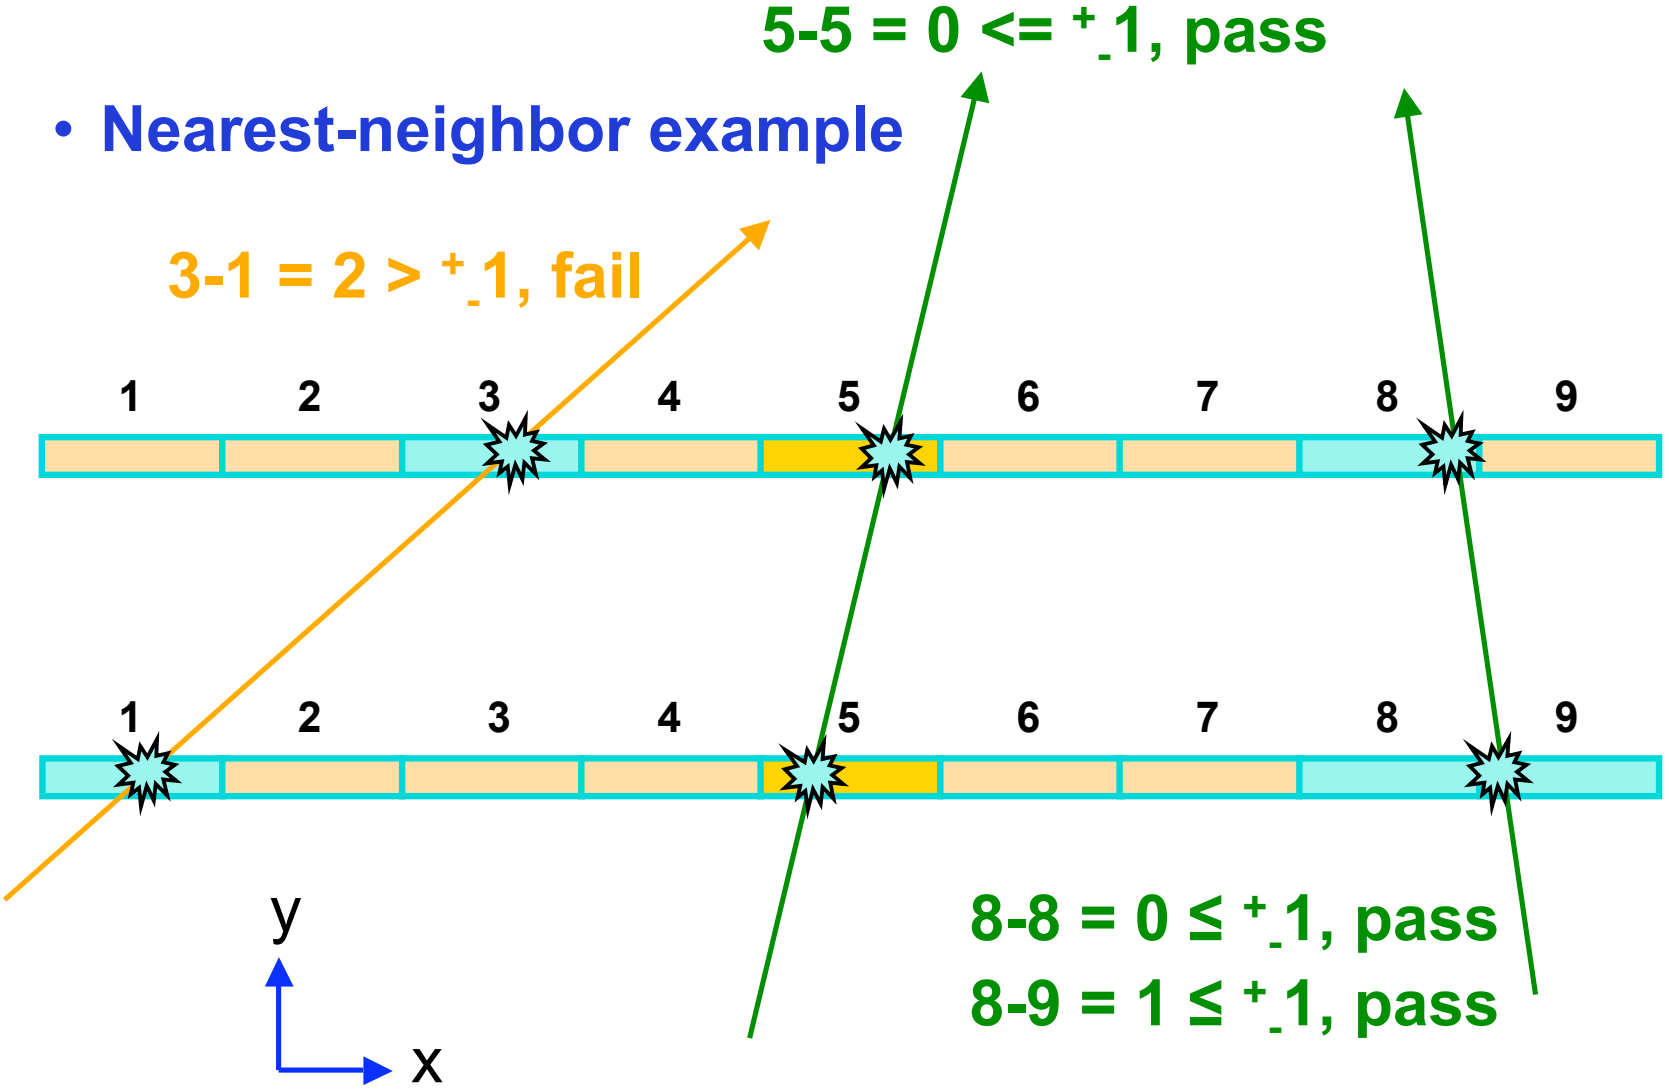

Correlator Architecture

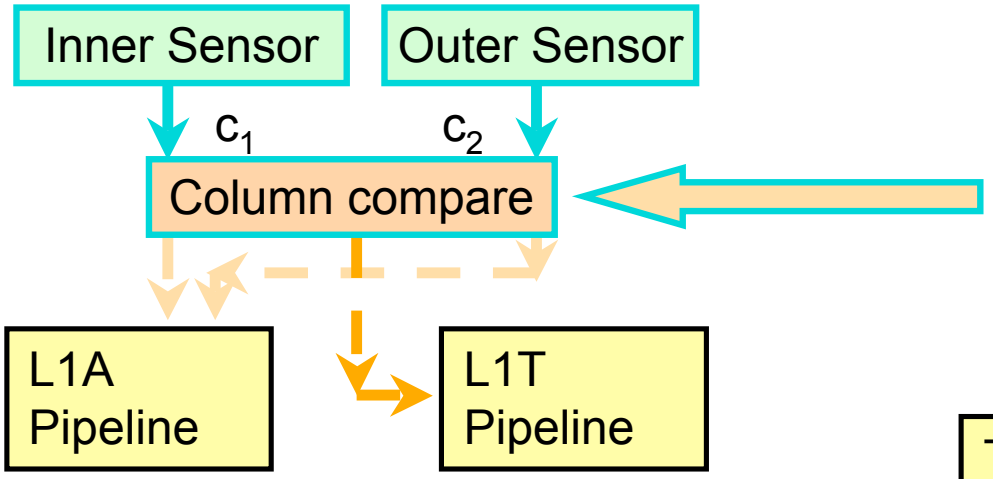

- If $c_2 > c_1 + 1$, discard c_1
- If $c_2 < c_1 - 1$, discard c_2
- Else copy c2 & c1 into L1 pipeline, next c1

This determines your search window
In this case, nearest-neighbour

At low luminosity, all hits could be read out
Put a 'bypass' switch in correlator

L1A pipeline: 512 BX @ 80MHz x 4 hits / $(1.28\text{cm})^2$
→ ~10kByte event buffer

p_T Cuts in a Stacked Tracker – p_T Cut Probabilities

- J. Jones

- **Depends on:**

Layer Sepn. & Radius

Pixel Size

Search Window

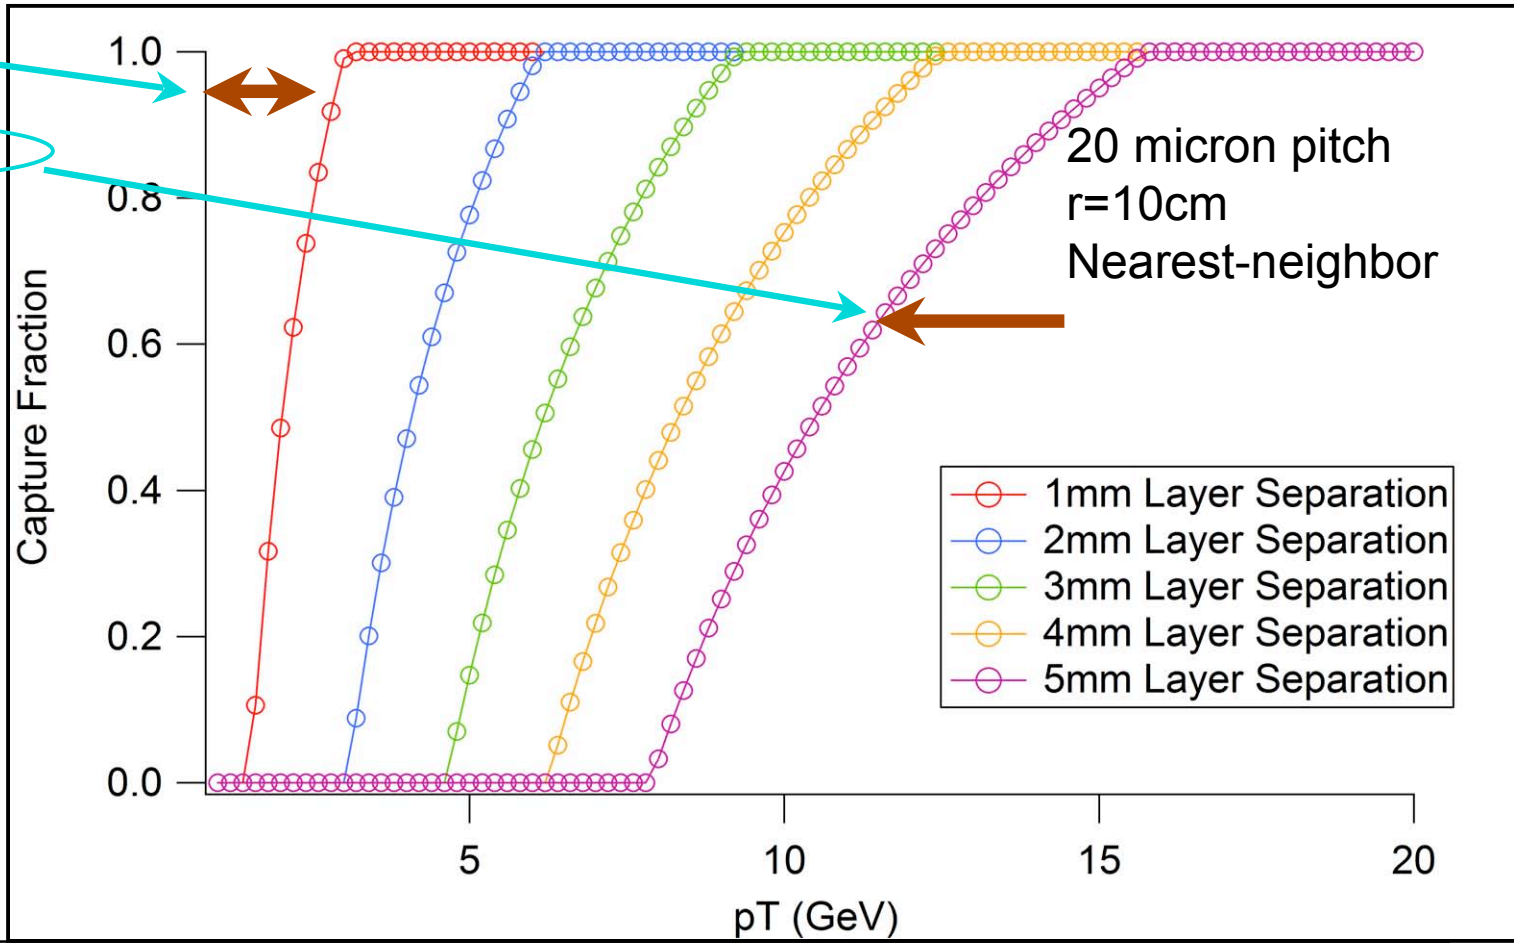

There is an additional 'blurring' caused by charge sharing…

Use of CMS L1 Tracking Trigger

- D. Acosta

Combine with L1 μ trigger as is now done at HLT:

- **Attach tracker hits to improve P_T assignment precision from 15% standalone muon measurement to 1.5% with the tracker**
  - Improves sign determination & provides vertex constraints
- **Find pixel tracks within cone around muon track and compute sum P_T as an isolation criterion**
  - Less sensitive to pile-up than calorimetric information *if* primary vertex of hard-scattering can be determined (~100 vertices total at SLHC!)

To do this requires η–φ information on muons finer than the current 0.05–2.5°

- **No problem, since both are already available at 0.0125 and 0.015°**

CMS Muon Rate at $L = 10^{34}$

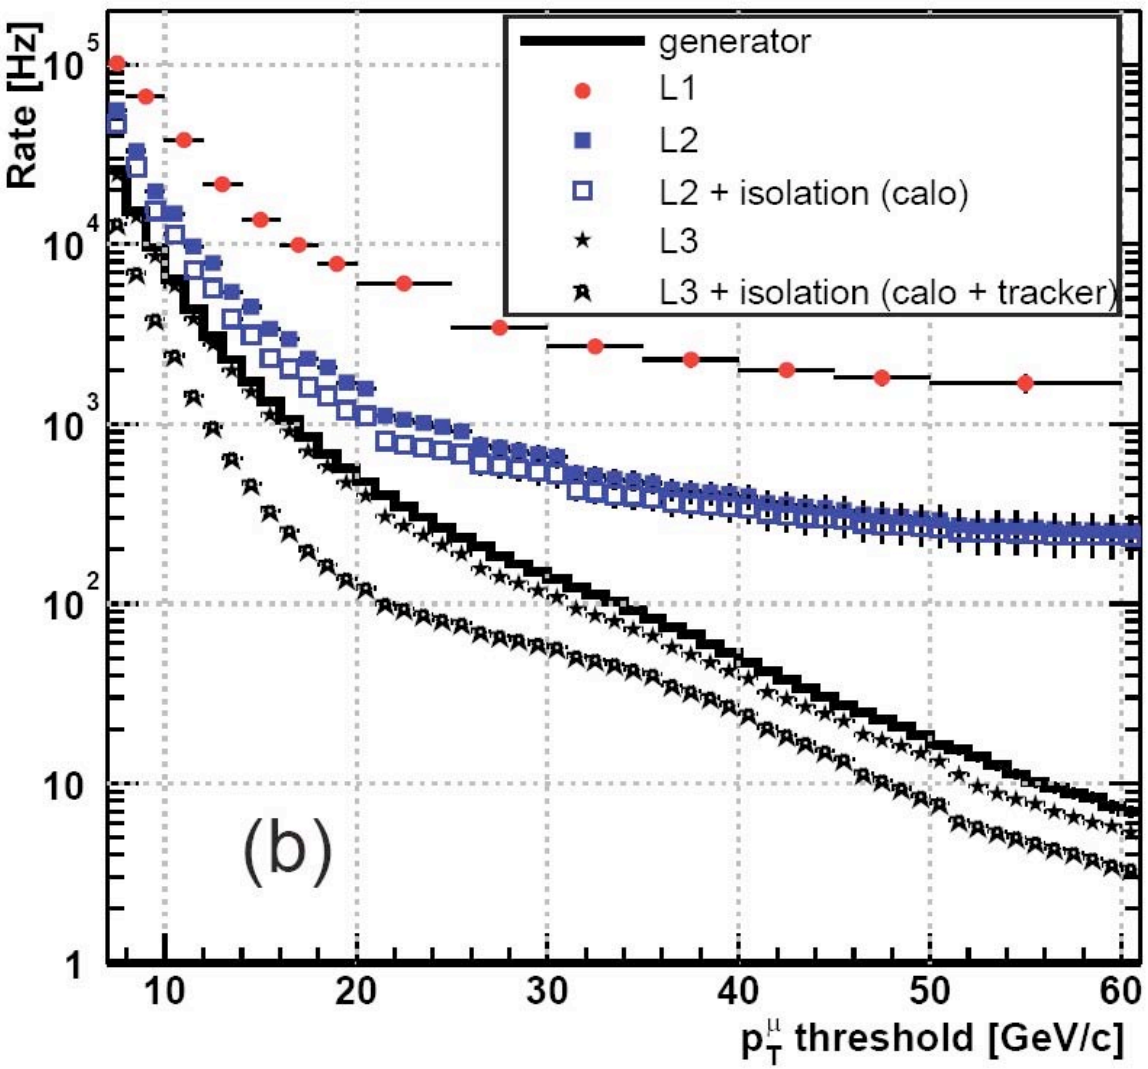

From CMS DAQ TDR

Note limited rejection power (slope) without tracker information

CMS SLHC e/γ/τ object clustering

e/γ/τ objects cluster within a tower or two

- **Crystal size is approximately Moliere radius**
  - Trigger towers in ECAL Barrel contain 5x5 crystals
- **2 and 3 prong τ objects don't leak much beyond a TT**
  - But, they deposit in HCAL also

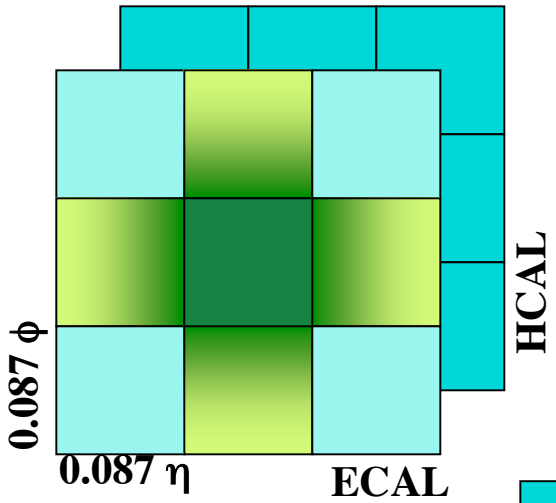

E_T scale: 8-bits

e/γ E_T = 1 x 2 or 2 x 1 sum
e/γ H/E cut for all 9 towers
e/γ isolation patterns:

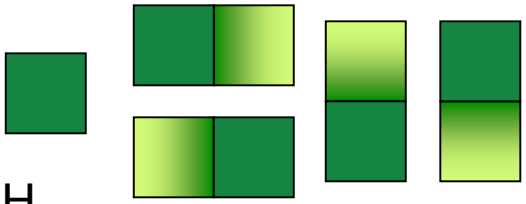

τ E_T = 3 x 3 sum of E + H
τ isolation patterns include E & H:

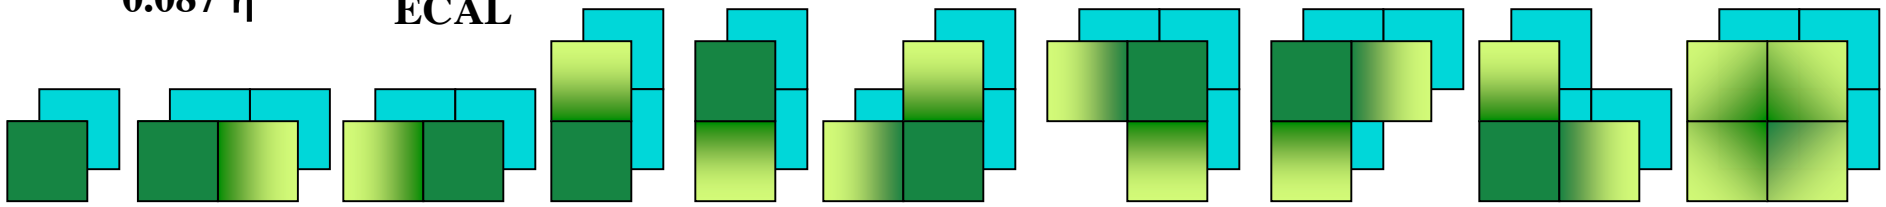

CMS SLHC e / γ / τ object ⇔track correlation

Use e / γ / τ objects to seed tracker readout

- Track seed granularity 0.087ϕ x 0.087η ⇔ 1 x 1
- Track seed count limited by presorting candidates
  - e.g., Maximum of 32 objects?

Tracker correlation

- Single track match in 3x3 with crude P_T (8-bit ~ 1 GeV)
  - Electron (same for muons)
- Veto of high momentum tracks in 3x3
  - Photon
- Single or triple track match
  - Tau

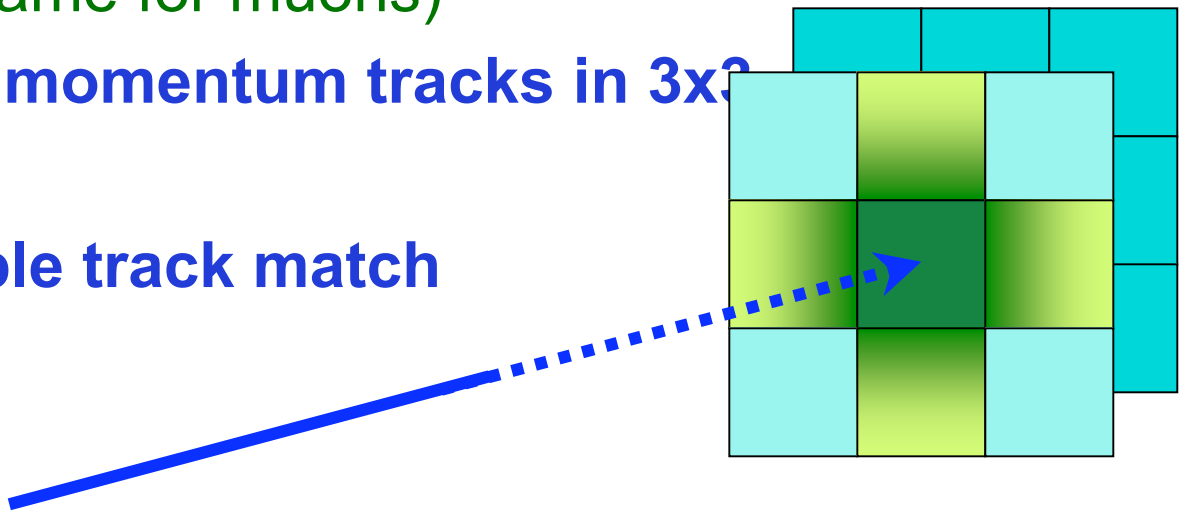

CMS SLHC Jet Clustering

Cluster jets using 2x2 primitives: 6x6, 8x8, 10x10

- Start from seeds of 2x2 E+H (position known to 1x1)
- Slide window at using 2x2 jet primitives
- E_T scale 10-bits, ~1 GeV

Jet Primitive is sum of E_T in E/HCAL

Provide choice of clustering?

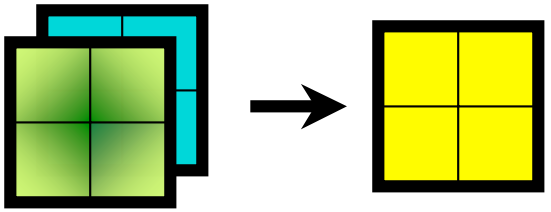

6x6 Jet

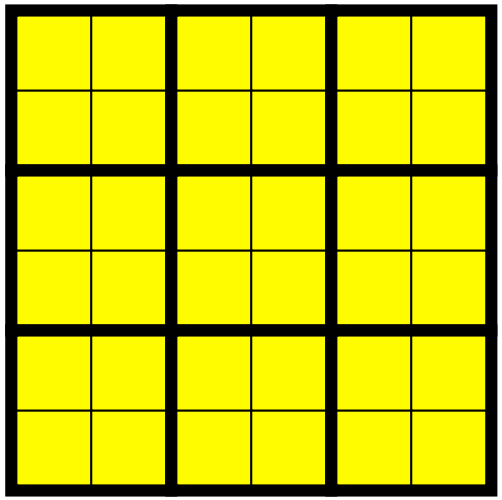

8x8 Jet

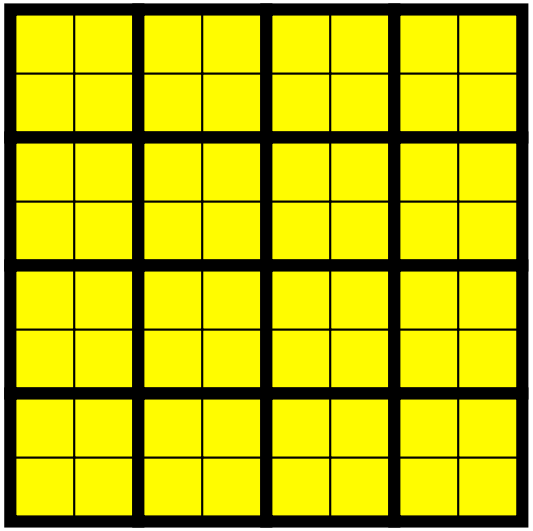

10x10 Jet

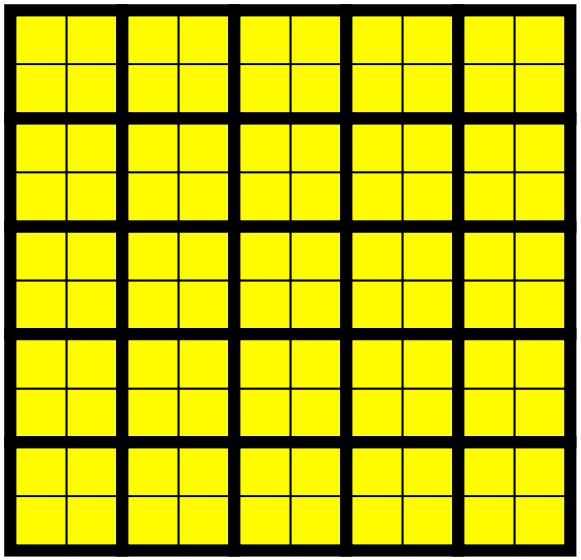

CMS tracking for electron trigger

Present CMS electron HLT

- C. Foudas & C. Seez

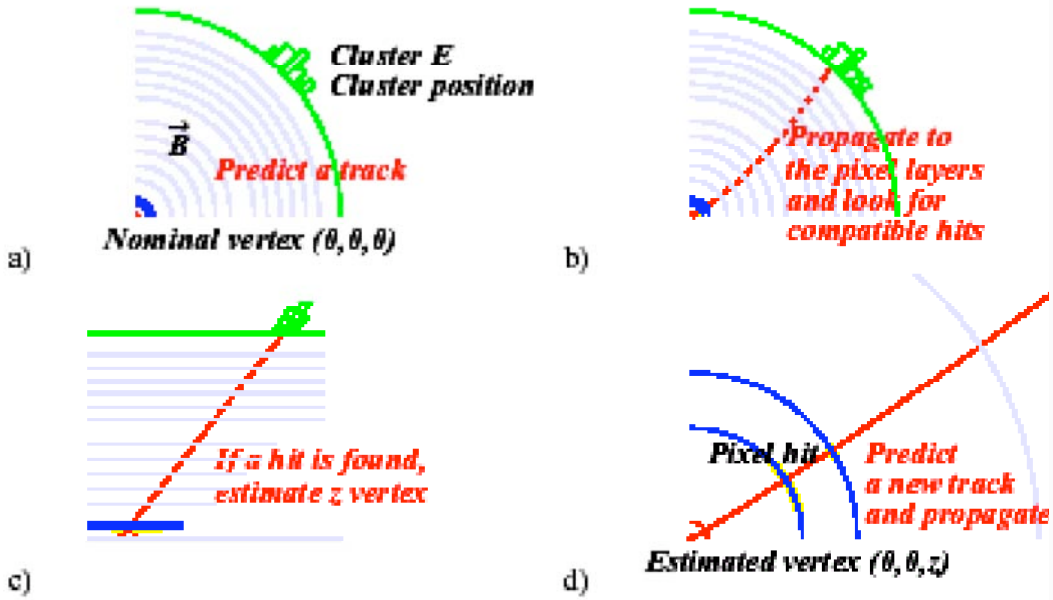

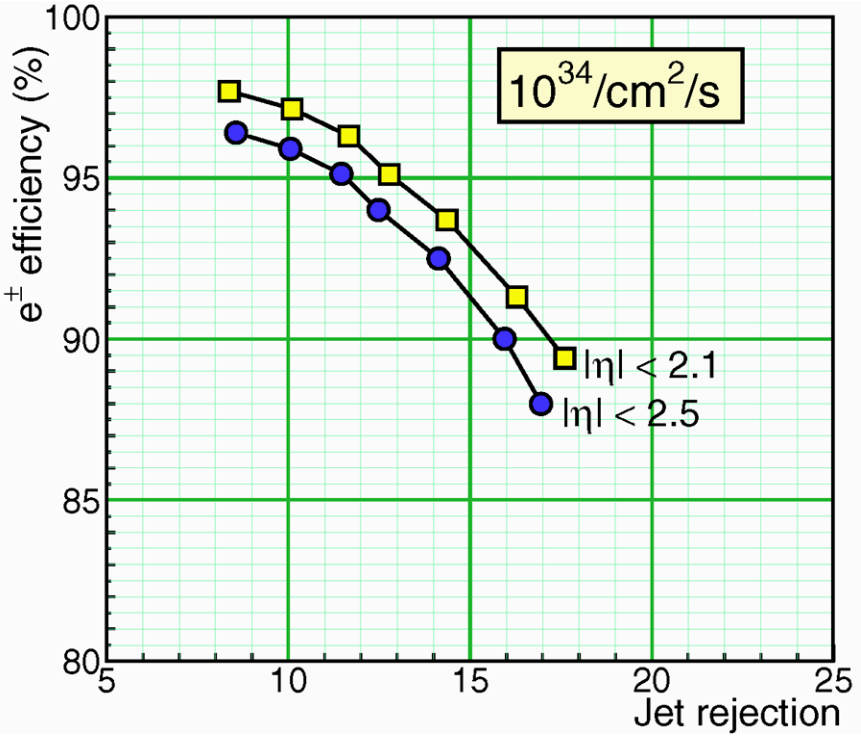

Factor of 10 rate reduction

γ: only tracker handle: isolation

- Need knowledge of vertex location to avoid loss of efficiency

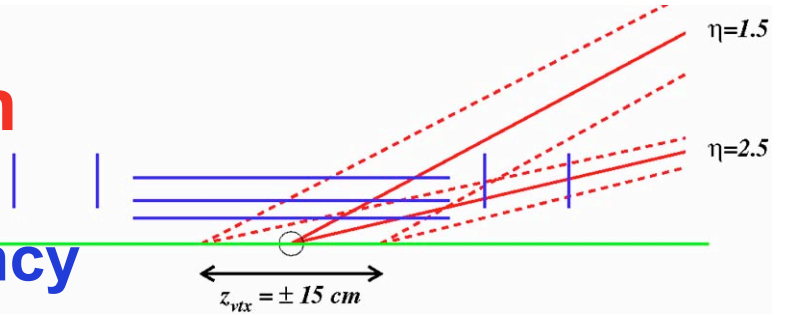

CMS tracking for τ-jet isolation

τ-lepton trigger: isolation from pixel tracks outside signal cone & inside isolation cone

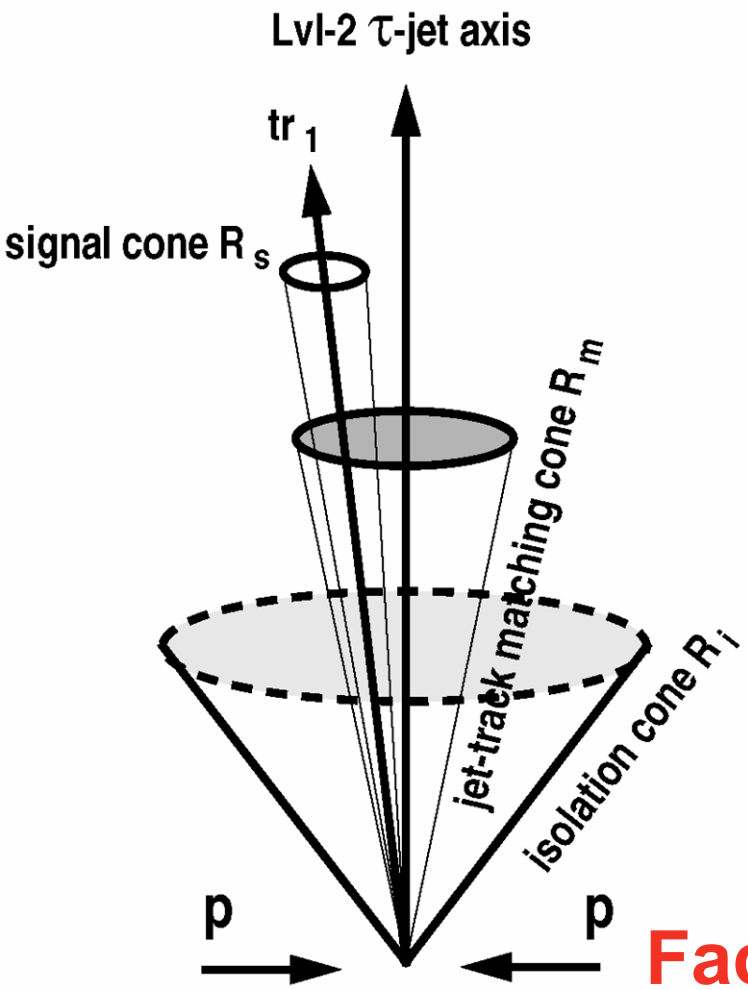

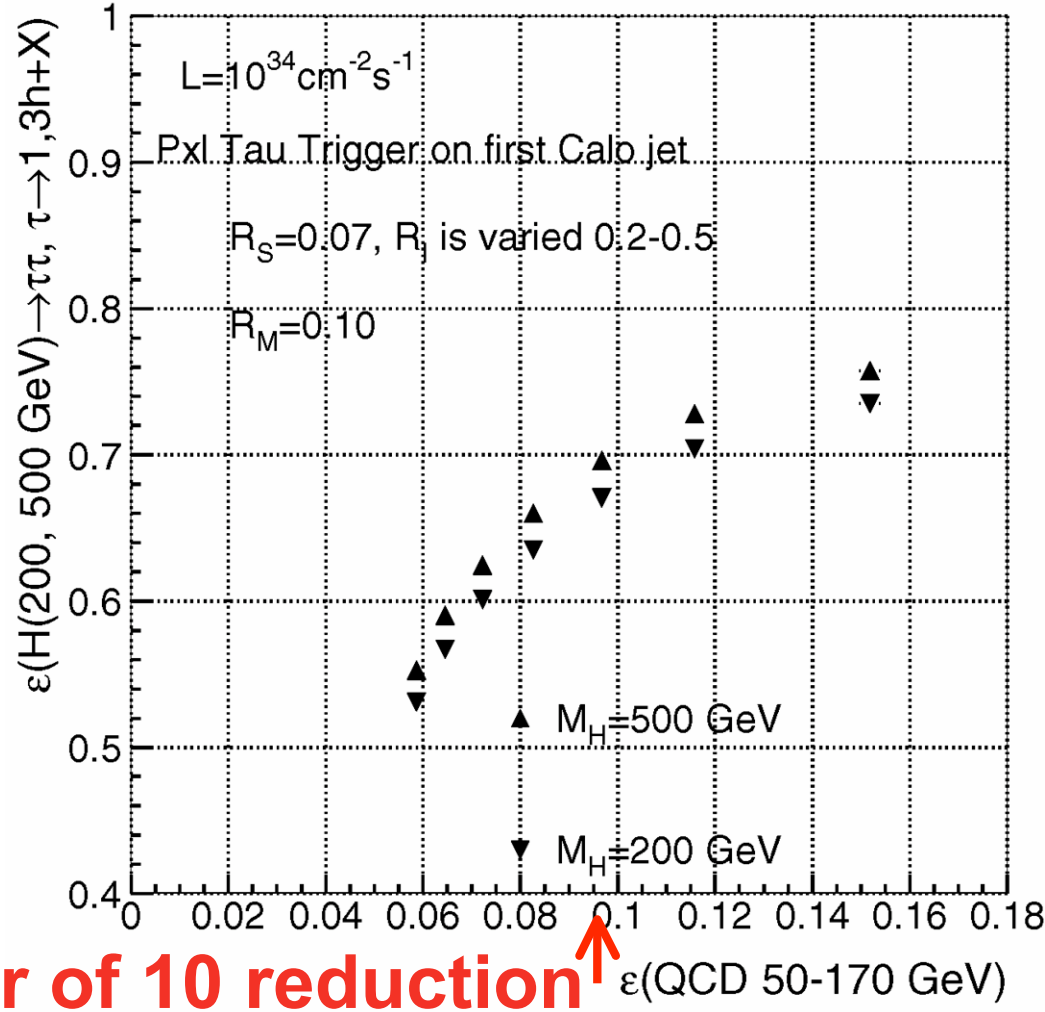

Factor of 10 reduction

CMS L1 Algorithm Stages

Current for LHC:

TPG ⇒ RCT ⇒ GCT ⇒ GT

Proposed for SLHC (with tracking added):

TPG ⇒ Clustering ⇒ Correlator ⇒ Selector

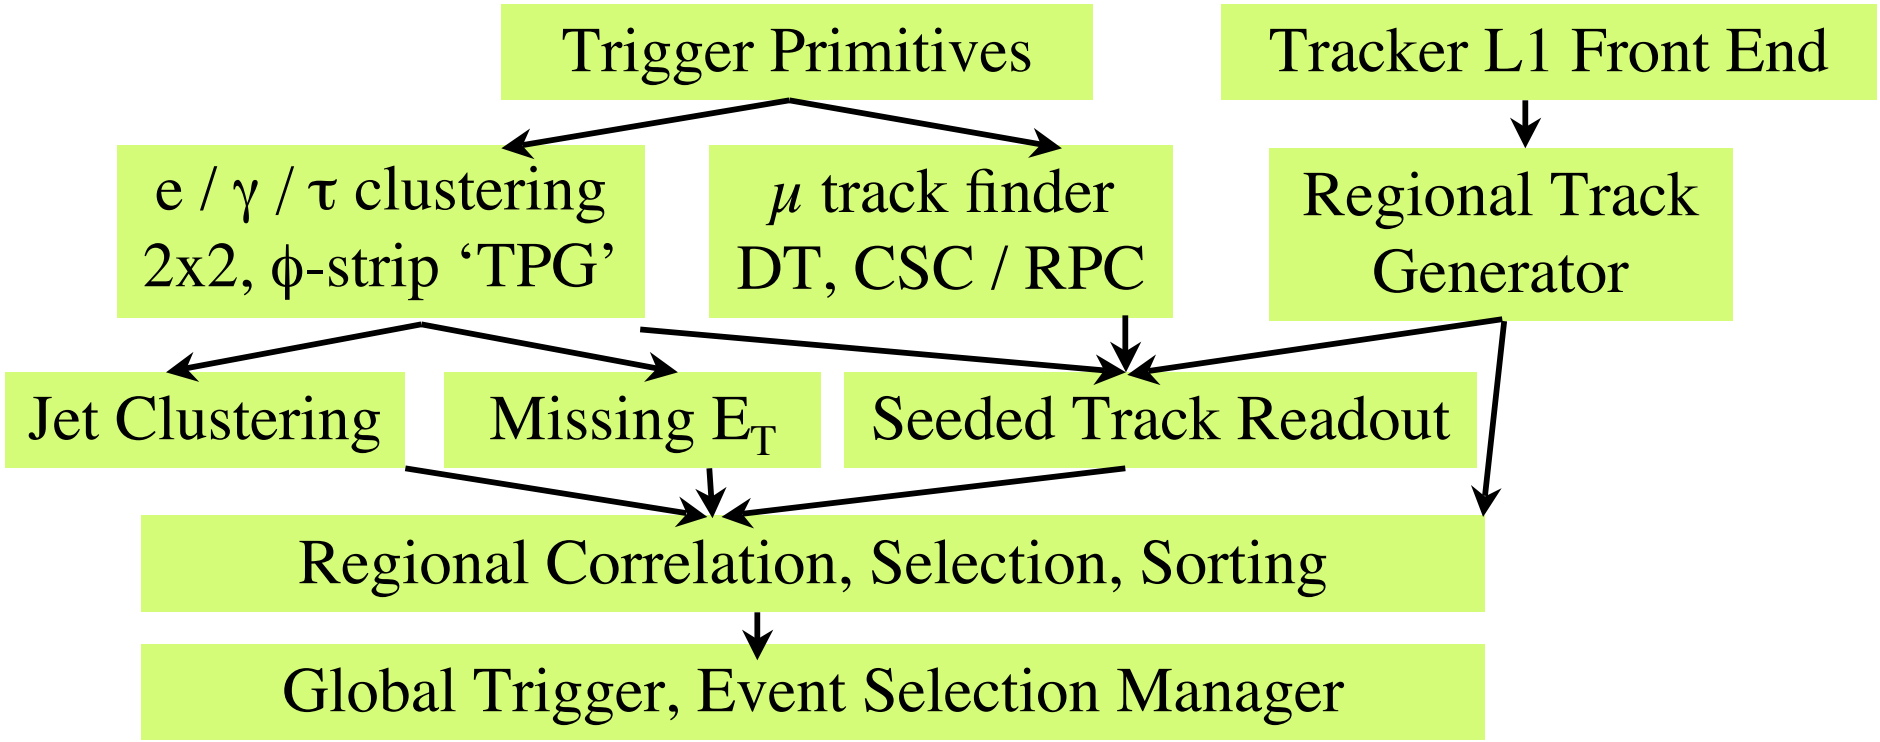

CMS SLHC Trigger Architecture

LHC:

- Level 1: Regional to Global Component to Global

SLHC Proposal:

- Combine Level-1 Trigger data between tracking, calorimeter & muon at Regional Level at finer granularity
- Transmit physics objects made from tracking, calorimeter & muon regional trigger data to global trigger
- Implication: perform some of tracking, isolation & other regional trigger functions in combinations between regional triggers
  - New “Regional” cross-detector trigger crates
- Leave present L1+ HLT structure intact (except latency)
  - No added levels --minimize impact on CMS readout

CMS Level-1 Latency

CMS Latency of 3.2 μsec becomes 256 crossings @ 80 MHz

- Assuming rebuild of tracking & preshower electronics will store this many samples
- Calorimeters keep 40 MHz sampling: 128 crossings at 3.2 μsec

Do we need more?

- Yield of crossings for processing only increases from ~70 to ~140
  - It's the cables!
- Parts of trigger already using higher frequency

How much more? Justification?

- Combination with tracking logic
- Increased algorithm complexity
- Asynchronous links or FPGA-integrated deserialization require more latency
- Finer result granularity may require more processing time
- ECAL digital pipeline memory is 256 40 MHz samples = 6.4 μsec
  - Propose this as CMS SLHC Level-1 Latency baseline

CMS SLHC L-1 Trigger Summary

Attempt to restrict upgrade to post-TPG electronics as much as possible where detectors are retained

- **Only change where required -- evolutionary -- some possible pre-SLHC?**
  - Inner pixel layer replacement is just one opportunity.

New Features:

- **80 MHz I/O Operation**
- **Level-1 Tracking Trigger**
  - Inner pixel track & outer tracker stub
  - Reports "crude" P_T & multiplicity in ~ 0.1x 0.1 $\Delta\eta \times \Delta\phi$
- **Regional Muon & Cal Triggers report in ~ 0.1 x 0.1 $\Delta\eta \times \Delta\phi$**
- **Regional Level-1 Tracking correlator**
  - Separate systems for Muon & Cal Triggers
  - Separate crates covering $\Delta\eta \times \Delta\phi$ regions
  - Sits between regional triggers & global trigger
- **Latency of 6.4 μsec**

CMS DAQ: Possible upgrade

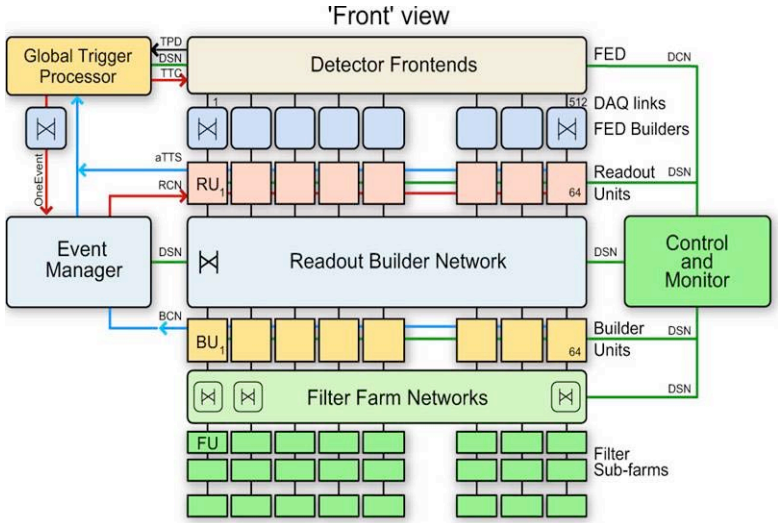

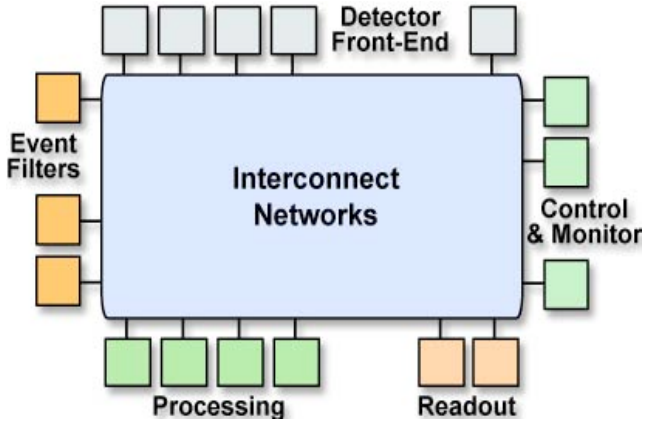

- S. Cittolin

LHC DAQ design:

A network with Terabit/s aggregate bandwidth is achieved by two stages of switches and a layer of intermediate data concentrators used to optimize the EVB traffic load.
RU-BU Event buffers ~100GByte memory cover a **real-time interval of seconds**

SLHC DAQ design:

A **multi-Terabit/s network** congestion free and scalable (as expected from communication industry).
In addition to the Level-1 Accept, the Trigger has to transmit to the FEDs additional information such as the event type and the event destination address that is the processing system (CPU, Cluster, TIER..) where the event has to be built and analyzed.
The event fragment delivery and therefore the **event building will be warranted by the network protocols** and (commercial) network internal resources (buffers, multi-path, network processors, etc.)
Real time buffers of Pbytes temporary storage disks will cover a **real-time interval of days**, allowing to the event selection tasks a better exploitation of the available distributed processing power.

New SLHC Fast Controls, Clocking & Timing System (TTC)

80 MHz:

- **Provide this capability "just in case" SLHC can operate at 80 MHz**
  - Present system operates at 40 MHz
- **Provide output frequencies close to that of logic**

Drive High-Speed Links

- **Design to drive next generation of links**
  - Build in very good peak-to-peak jitter performance

Fast Controls (trigger/readout signal loop):

- **Provides Clock, L1A, Reset, BC0 in real time for each crossing**
- **Transmits and receives fast control information**
- **Provides interface with Event Manager (EVM), Trigger Throttle System**
  - For each L1A (@ 100 kHz), each front end buffer gets IP address of node to transmit event fragment to
  - EVM sends event building information in real time at crossing frequency using TTC system
    - **EVM updates 'list' of avail. event filter services (CPU-IP, etc.) where to send data**
    - **This info.is embedded in data sent into DAQ net which builds events at destination**
  - Event Manager & Global Trigger must have a tight interface
- **This control logic must process new events at 100 kHz → R&D**

SLHC DAQ: Readout

Front End: more processing, channels, zero suppression

- Expect VLSI improvements to provide this
- But many R&D issues: power reduction, system complexity, full exploitation of commercial data-communications developments.

Data Links: Higher speeds needed

- Rx/Tx available for 40G, electronics for 10G now, 40G soon, accepted protocols emerging: G-ethernet, Fibre Channel, SDH/Sonet
- Tighter integration of link & FE --R&D on both should take place together

Radiation tolerance: major part of R&D

- All components will need testing
- SEU rate high: more error detection & correction

SLHC Front End Electronics

- A. Marchioro
- K. Einsweiler

Power

- **Key problem -- for everyone!**
  - Major difficulties: power density & device leakage
- **Power impact on services (cooling)**

Radiation -- Example: 130 nm Deep Sub-Micron CMOS

- **Total Integrated Dose**
  - Enclosed transistor circuits do well, linear layout some problems
- **Single Event Upset**
  - Higher sensitivity but enclosed transistors give rate ~ 250 nm DSM
- **Single Event Lockup**
  - Not observed and not expected for careful designs
- **Tentative Conclusion: better than 250 nm DSM**

Complexity

- **Modes involve more neighbors due to capacitive cross-couplings**

Cost

- **Per IC cost is lower but cost of mask set over 0.5 M$!**
- **Wafer cost much higher but more IC's per wafer**
- **Engineering run: 0.25 μm: 150 k$, 0.13 μ: 600 k$**

SLHC Electronic Circuits

ADC's -- benefit from technology development

- Today: CMS ECAL in 0.25 μm: 11.1 bit @ 40 Ms/s @ 125 mW
- SLHC: Design in 65 nm, apply scaling: 6 bit @ 80 MS/s @ 2.5 mW

Technology Choice

- Tradeoff between power and cost (SiGe BiCMOS vs. CMOS DSM)
- Evaluate 90 nm, 65 nm: long & expensive process (need access to design rules)

Power regulators

- Distribute regulation over many small regulators to save power
  - Local DC-DC converters & "Serial powering": build regulators into chips

Need new designs to save power in digital circuits

- Reduce voltage where possible
- Design architecture to reduce power
  - # FF's, Inverters/FF, Capacitance/Inverter
- Turn off digital blocks when results not needed
  - Gate input or clocks to blocks
  - Turn off entire chips when not needed (temp monitor)
- Use data compression wherever possible
  - If occupancy remains low, transmit hit channels
  - For calorimeter data, try Huffman encoding on differences?

SLHC Link Electronics

- F. Vasey

Faster link electronics available

- Si-Ge & Deep Sub-Micron

Link electronics has become intrinsically rad-tolerant

More functionality incorporated

- Controls, diagnostics, identification, error correction, equalization

Link electronics now available up to 10G

Industrial development mostly digital

- Easier to store, buffer, multiplex, compress

For now all links use LHC crossing clock as timing ref.

- Possible to run with other clocks with buffers (latency)

Optical Links:

- Transmitters & Receivers available up to 10G & 40G
- Variety of fibers available
- Variety of packages are available
- Possibility to use frequency mult. to better use bandwidth

FPGA Technology

Available Now:

- 8M Usable Gates
- 1500 Fine Pitch Ball Grid Array Pacakges
- 1200 (Altera) or 1100 (Xilinx) I/O pins
- Core Voltage 1.5 V
- Flexible internal clock management
- Built in Multi-Gigabit Transceivers: 0.6 - 11 Gbps
- Built-in I/O serializer/deserializer (latency)

Upgrade:

- Logic Speed, Usable Gates, Logic Volume plenty
- Use of these devices becomes difficult, limiting factor
  - Packaging, routing, mounting, voltages all difficult
  - Need to explore new I/O techniques - built in serdes?

Data Link Technology

Integration:

- **Discrete deserializers vs. integration in FPGAs**
  - Issue: deserializer latency (improving)

Connections:

- **CAT6,7,8 cables for 1G and 10G Ethernet**
- **Parallel Optical Links**
- **Parallel LVDS at 160 MHz**

Backplanes:

- **Use cable deserializer technology**
- **Exploit new industry standard full-mesh and dual-star serial backplane technology & PCI Serial Express:**
  - Each serial link operates at 2.5 GHz bit rate (5 GHz in development) 8B/10B encoding → 2.0 (4.0) Gbps data rate.
- **Issues: latency for deserialization & circuitry for synchronization**

Power:

- **Providing power & cooling infrastructure a challenge**

SLHC Trigger & DAQ Summary

Significant Challenges:

- Occupancy: degraded algorithms, large event size
- High trigger rates: bandwidth demands on DAQ
- Radiation damage: front end electronics
- Increased channel counts, data volume: electronics power

Promising directions for development:

- Use of tracking & finer granularity in Level-1 Trigger
- More sophisticated calculations using new FPGAs
- Higher speed data links & backplanes
  - FPGA link/serializer integration
- New DAQ architecture to exploit commercial developments
- Smaller feature size Deep Sub-Micron CMOS for front ends
  - Good radiation tolerance with appropriate design rules
- Designs for lower power electronics
  - Lower voltages, architecture, shut-off when not needed