


Tracking and Trigger Upgrades to CMS

19 November 2014

University of Bristol
Brunel University
Imperial College London
Rutherford Appleton Laboratory

Geoff Hall



CMS: Compact Muon Solenoid



CMS Tracker and its sub-systems

- Two main sub-systems: Silicon Strip Tracker and Pixels
 - pixels quickly removable for beam-pipe bake-out or replacement
 - SST not replaceable in reasonable time


Microstrip tracker	Pixels
~210 m ² of silicon, 9.3M channels	~1 m ² of silicon, 66M channels
73k APV25s, 38k optical links, 440 FEDs	16k ROCs, 2k olinks, 40 FEDs
27 module types	8 module types
~34kW	~3.6kW (post-rad)





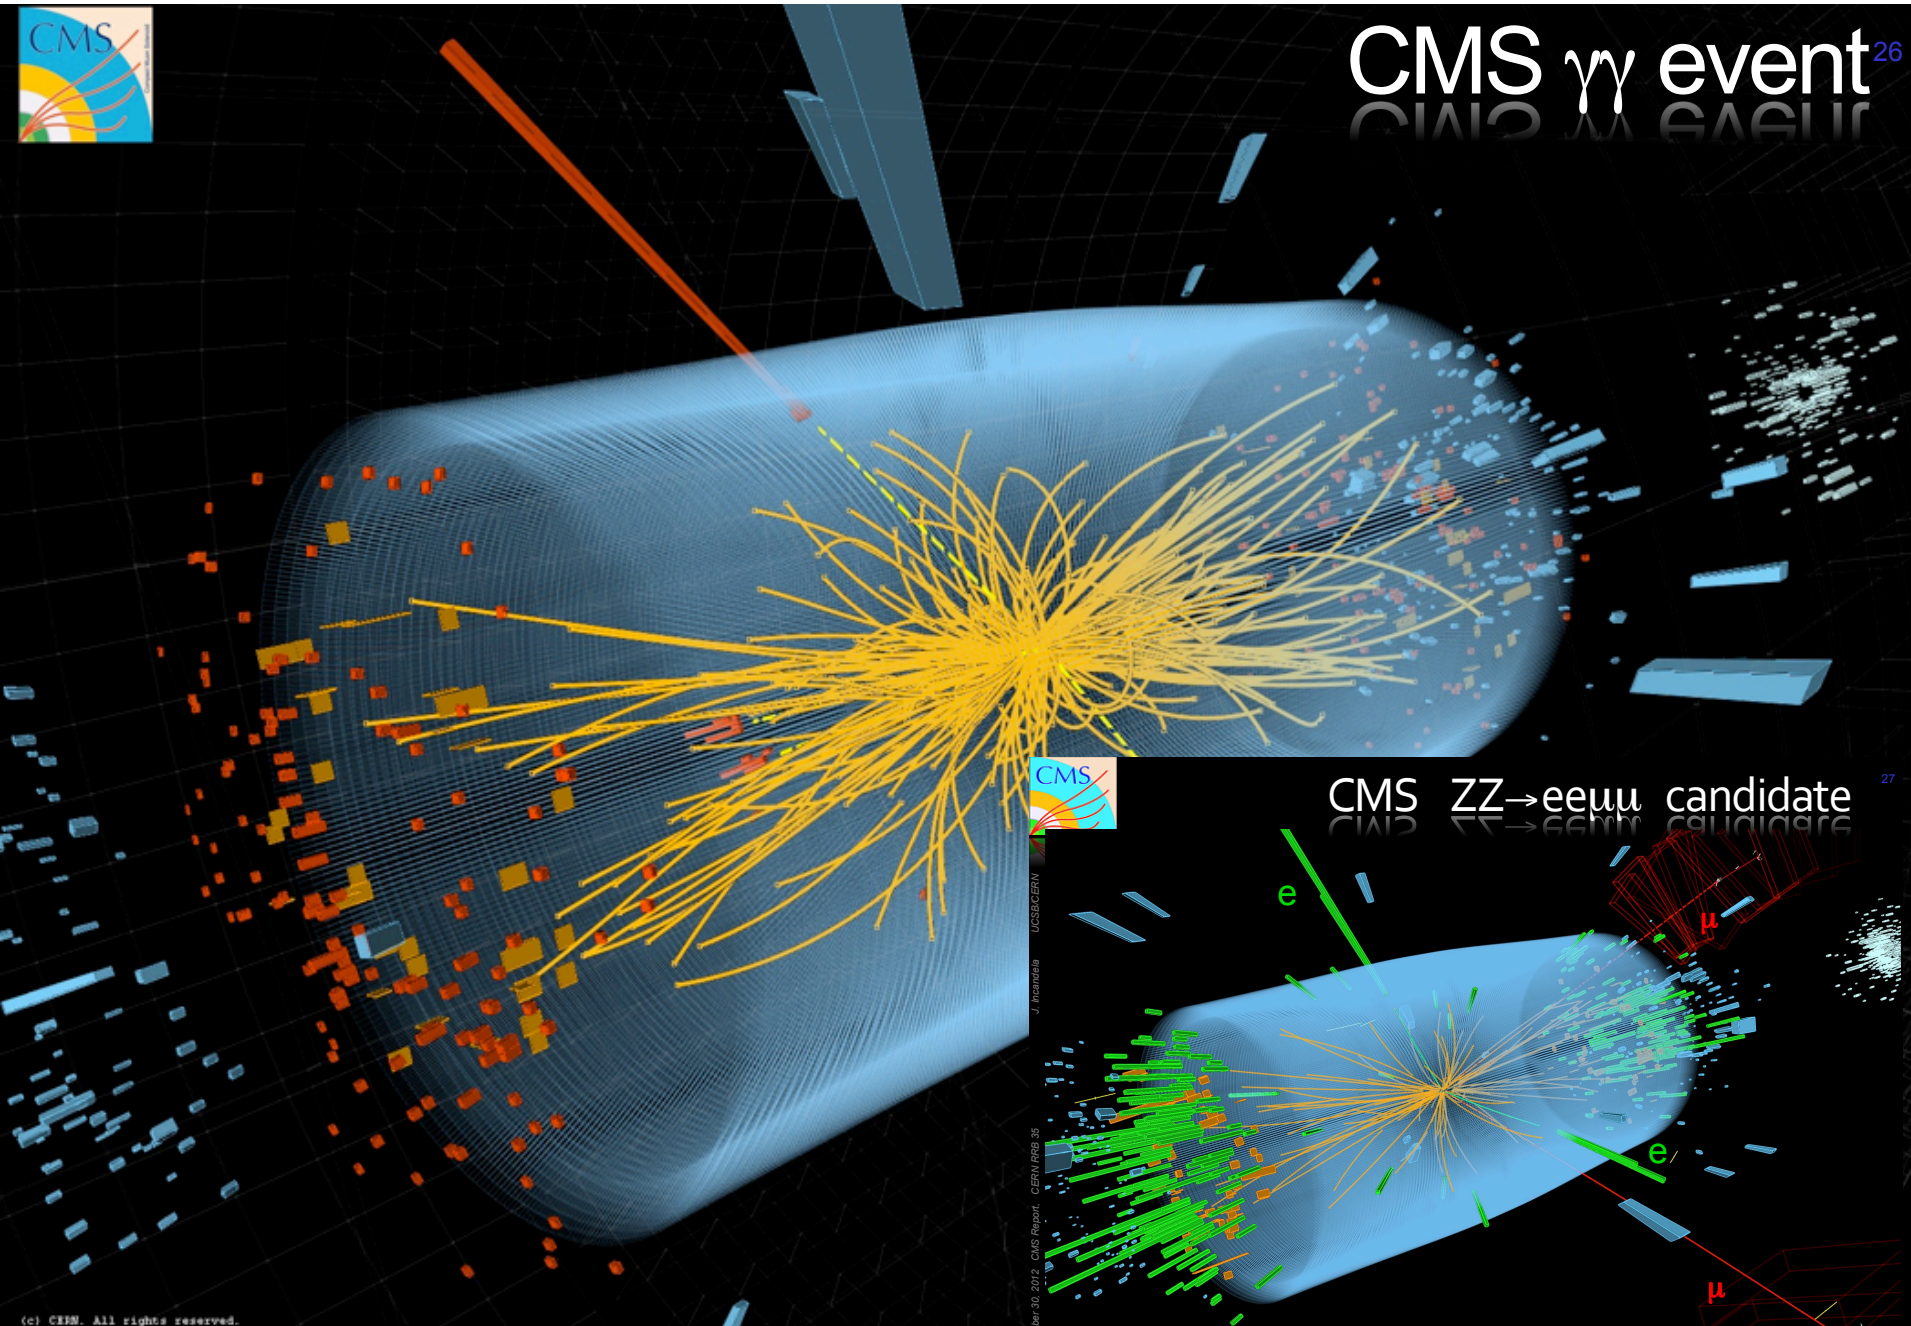
Current CMS L1 Trigger

- Pipelined trigger at 40 MHz
- Mix of FPGAs and ASICs
- Many copper parallel links
- Internal bandwidth constraints e.g. jet finding
- Detectors designed to accommodate:
 - 100 kHz maximum L1 rate
 - Max latency 4 μ s







CMS $\gamma\gamma$ event²⁶



CMS $ZZ \rightarrow eeu\mu$ candidate²⁷



October 30, 2012 CMS Report CERN/RBB/35 J. P. Prentice UCSB/CERN




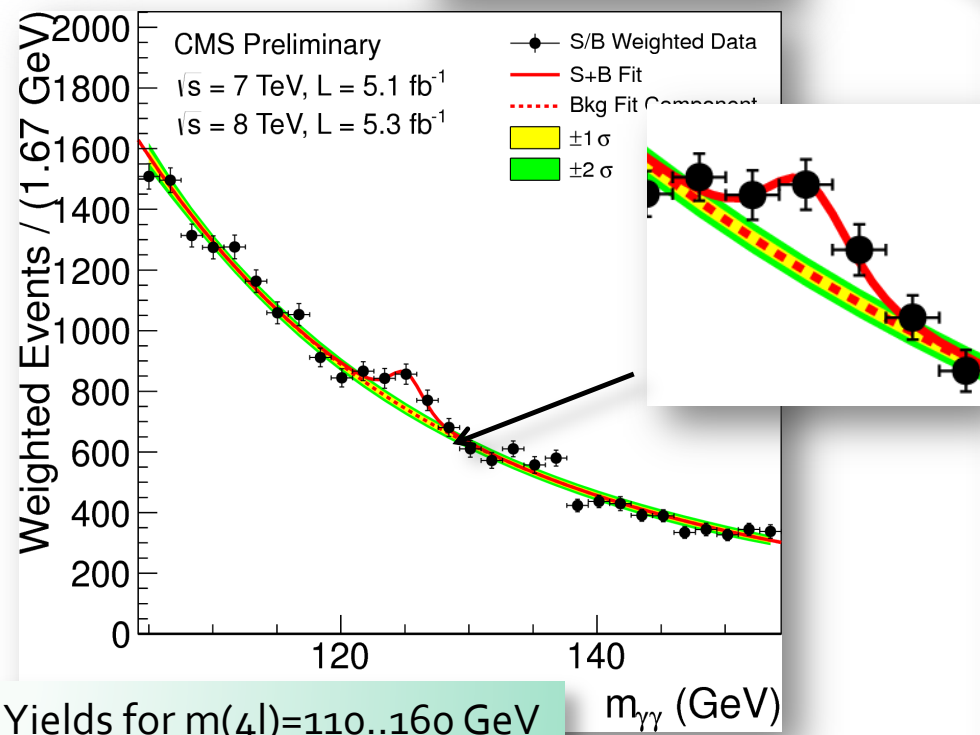
Discovery!

UCSB/CERN

J. Incandela

CERN RRB 35

October 30, 2012 CMS Report,



Yields for $m_{4l}=110..160$ GeV

Channel	4e	4 μ	2e2 μ	4 ℓ
ZZ background	2.65 ± 0.31	5.65 ± 0.59	7.17 ± 0.76	15.48 ± 1.01
Z+X	$1.20^{+1.08}_{-0.78}$	$0.92^{+0.65}_{-0.55}$	$2.29^{+1.81}_{-1.36}$	$4.41^{+2.21}_{-1.66}$
All backgrounds	$3.85^{+1.12}_{-0.84}$	$6.58^{+0.88}_{-0.81}$	$9.46^{+1.96}_{-1.56}$	$19.88^{+2.43}_{-1.95}$
$m_H = 126$ GeV	1.51 ± 0.48	2.99 ± 0.60	3.81 ± 0.89	8.31 ± 1.18

164 events expected in [100, 800 GeV]
 172 events observed in [100, 800 GeV]


Event-by-event errors

What next?


- Detector upgrades planned around essential shutdowns
 - LS1 – currently underway – collisions in 2015
 - upgrade energy to 13-14 TeV
 - LS2 – 18 months 2018-2019
 - collimation, injector and cryogenic upgrades
 - LS3 - 30 months 2023-2025
 - prepare for levelled high-luminosity running
- Some extra activities possible in Year End Technical Stops
 - e.g. probable extended 2016 YETS for CMS pixel installation
- Major upgrade essential for Run 3 (post-2025)

HL-LHC - goals

- Prepare machine for operation beyond **2025 and up to 2035**
- Devise beam parameters and operation scenarios for:
 - total integrated luminosity of **3000 fb⁻¹** in around 10-12 years
 - an integrated luminosity of **250 fb⁻¹ per year**
 - $\mu \leq 140$ (peak luminosity of $5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$)



2010 - 2035



Motives for upgrades

- Many big physics questions outstanding
 - perhaps improving the Higgs precision will be the only way?
 - experimentally challenging – both for detectors and analyses
 - if SUSY is discovered in Run 2, may expect some revision of plans
 - if discoveries are absent should expect further theoretical ideas
- probably unwise to assume today's experimental and theoretical landscape will be static
 - LHC machine and experiments have been running only three years
 - gains from software and analyses have been impressive
- Eventually important parts of detectors will be under great stress
 - radiation damage
 - data volumes and rates
 - performance improvement from technology evolution

GPDs: scope of Phase II detector upgrades

- Most sub-detectors are foreseen to survive to 3000 fb^{-1}
 - with on-going maintenance and refurbishment where possible
- Trackers must be completely replaced
 - radiation damage limits their lifetimes to $<500 \text{ fb}^{-1}$
- New tracker readout systems are therefore also essential
 - based on more modern technologies, which improve performance
 - all sub-system readout systems must remain compatible
- Triggers must also be substantially upgraded
 - designed for $10^{34} \text{ cm}^{-2}\text{s}^{-1}$, $\langle N_{\text{ev}} \rangle \sim 25$
 - with safety factors – but exploited to maximise acceptance


BASIC REQUIREMENTS FOR ATLAS AND CMS

ttbar event with 140 pile-up events
(ATLAS simulation)



- **Radiation hardness**
 - Ultimate integrated luminosity considered $\sim 3000 \text{ fb}^{-1}$ (original $\sim 400 \text{ fb}^{-1}$)
 - Radiation hard sensor material
 - New readout electronics required
- **Granularity**
 - Resolve 140-200 collisions per bunch crossing
 - Maintain occupancy below % level
 - Requires much higher granularity

- **Improve tracking performance**
 - Reduce material in the tracking volume
 - Improve performance at low p_t
 - Reduce rates of nuclear interaction, photon conversions, Bremsstrahlung...
 - Reduce average pitch
 - Improve performance at high p_t



CMS upgrade summary

New Tracker

- Radiation tolerant - high granularity - less material
- Tracks in hardware trigger (L1)
- Coverage up to $\eta \sim 4$

Muons

- Replace DT FE electronics
- Complete RPC coverage in forward region (new GEM/RPC technology)
- Investigate Muon-tagging up to $\eta \sim 3$

Barrel ECAL

- Replace FE electronics
- Cool detector/APDs

Endcap Calorimeters


- Radiation tolerant
- High granularity

Trigger/DAQ

- L1 (hardware) with tracks and rate up ~ 750 kHz
- L1 Latency $< 12.5 \mu\text{s}$
- HLT output rate 7.5 kHz

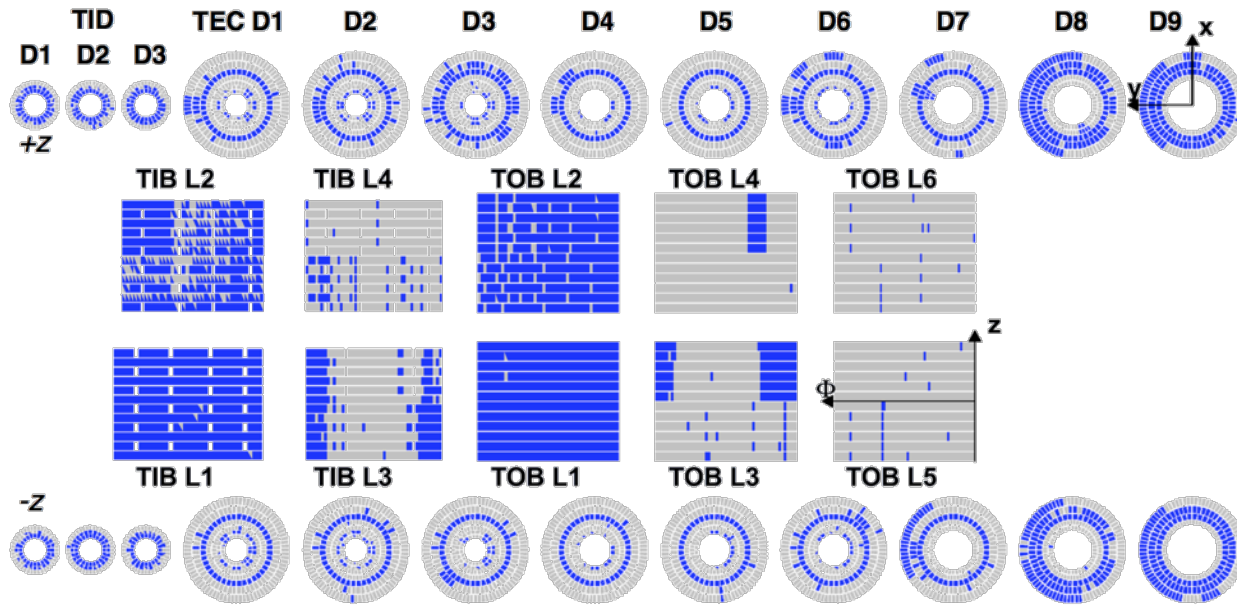
Other R&D

- Fast-timing for in-time pileup suppression
- Pixel trigger





Need for Tracker replacement

J Mans
ECFA 2014



Blue tracker modules are inactive after 1000 fb^{-1} due to very high leakage currents induced by hadron fluence.




- Extensive R&D campaigns happened in all experiments. defined with options to follow up.
 - For ATLAS and CMS - Outer Tracker well defined
 - **Common ATLAS & CMS Market Survey** for Outer Tracker for AC-coupled sensors
 - More studies necessary for inner pixel
 - Some common ATLAS/CMS wafer submissions planned

	Strips/strixel baseline	Pixel <u>outer</u> layers baseline / options	Pixel <u>inner</u> layers baseline / options	Special
ALICE	MAPS (Monolithic Active Pixels)			
ATLAS	<ul style="list-style-type: none"> • n-in-p planar FZ 300µm thick AC-coupled • and/or HV-CMOS 	<ul style="list-style-type: none"> • n-in-p (n) planar • and/or HR/HV-CMOS 	<ul style="list-style-type: none"> • n-in-n planar 100-200µm active thickness • and/or HR/HV-CMOS • and/or 3D • and/or diamonds 	<div style="border: 1px solid purple; padding: 5px; transform: rotate(-15deg); display: inline-block;"> With ~700m² needs to be factored in the global orders </div>
CMS	<ul style="list-style-type: none"> • n-in-p planar FZ 200µm <i>active</i> thickness AC- and DC-coupled • and/or MCz (pref) • and/or 300 µm 	<ul style="list-style-type: none"> • n-in-p planar 100-200µm active thickness 	<ul style="list-style-type: none"> • n-in-p planar 100-200µm active thickness • and/or 3D sensors 	
LHCb Geoff Hall	UT planar n-in-p • or p-in-n	VELO planar n-in-p • or n-in-n		

New issues for trigger

- $L \sim 5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ (levelled) $\Rightarrow N_{ev}/BX \sim 140 - 200$
- Calorimeters
 - isolation of $e/\gamma/\tau$ degraded by pile-up from π^0 γ s and hadrons
 - many more jets, which overlap
- Muon systems
 - increased combinatorial fakes, enhanced by multiple scattering
- Outcome: much higher rate of L1 triggers
 - usual response is to increase thresholds, which risks physics
 - even worse - raising thresholds does not look effective



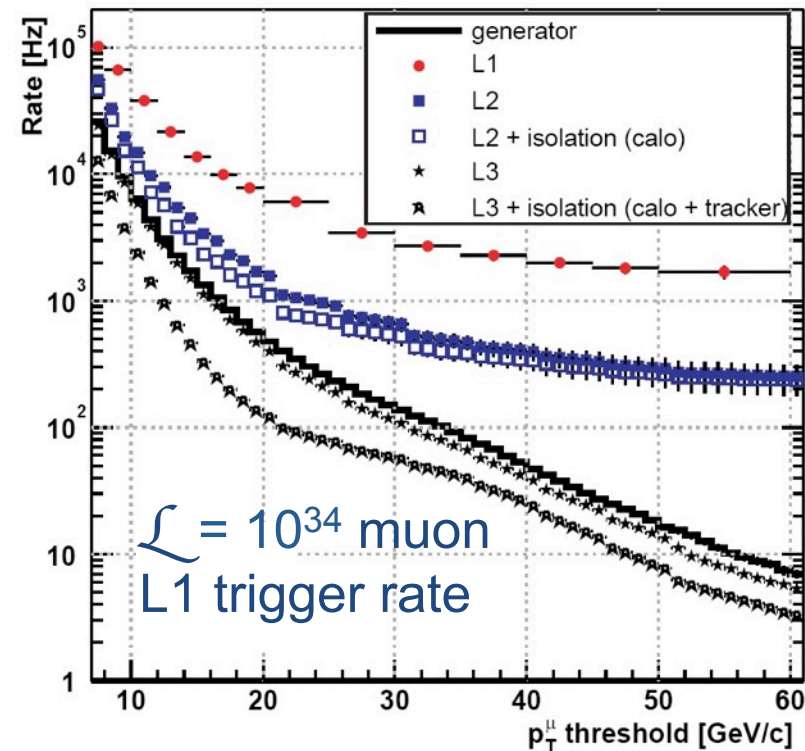
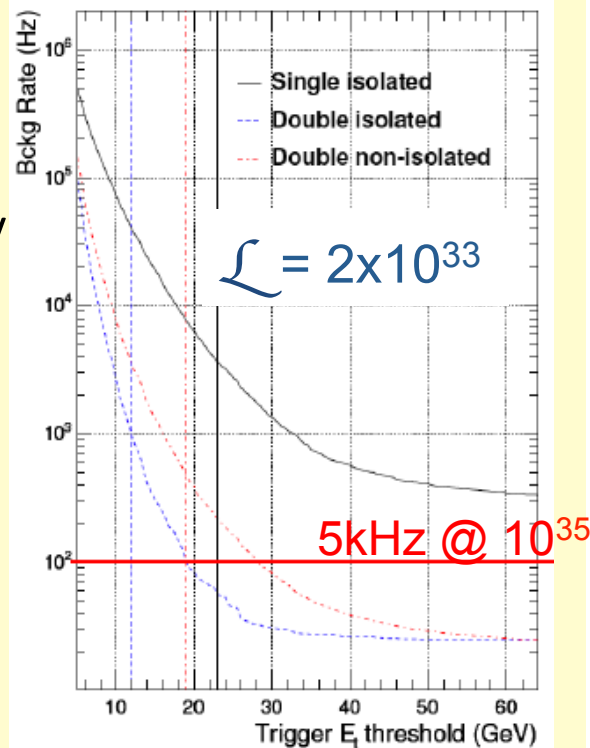
Why tracker input to L1 trigger?

- Single μ and e L1 trigger rates will greatly exceed 100kHz
 - similar behaviour for jets

Single electron trigger rate

$\langle p_T \rangle \approx$ few GeV/bx/
trigger tower

Isolation criteria alone are insufficient to reduce rate at $\mathcal{L} = 10^{35} \text{ cm}^{-2} \cdot \text{s}^{-1}$







Trigger levels


- Not feasible to achieve sufficient data reduction in L1 single step
 - 100 kHz -> ~500 kHz in Run 3
- When decisions were made on L2, two points of view
 - (custom) hardware processors needed to reduce data volume in ~50ms
 - sufficient computing power would evolve to avoid intermediate level
- ATLAS and CMS therefore diverged, with future implications
 - CMS must always store data on-detector until L1 decision
 - hardware **trigger latency limited by shortest buffer length**
 - **transfer large data volume** quickly to HLT = large BW
 - ATLAS can transfer selected data to L2 buffers
 - potentially much **longer trigger latency** possible
 - much smaller fraction of data, but **more complexity**
- **Latency is an important constraint: target < 12 μ s**

Silicon tracker with trigger-stub capability

Strip/Strip Modules
90 μm pitch/5 cm length






Inner Pixel
Covers up to $\eta=4.0$



Strip/Pixel Modules
100 μm pitch/2.5 cm length
100 μm x 1.5 mm “macropixels”


Stacked-tracker principle

- Compare pattern of hits in contiguous sensor elements in closely spaced layers
 - p_T cut set by angle of track in layer
 - primarily depends on layer separation
 - but increasing separation worsens fake combinations
 - details depend on
 - pitch
 - thickness
 - charge sharing
 - track impact point
 - ...




CMS Tracker ASIC evolution


- 1999: APV25 0.25 μ m
 - 7 mm x 8mm (128 chan)






- 2011: CBC 0.13 μ m
 - 7mm x 4mm (128 chan)



- 2013: CBC2 0.13 μ m
 - 11mm x 5mm (254 chan)





beam profile



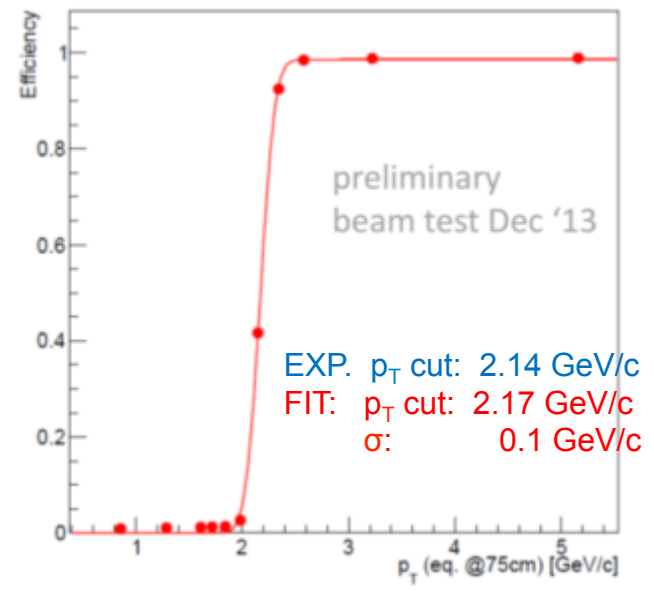
bump-bondable,
cluster & correlation logic

CMS Outer Tracker trigger

- ~15000 modules transmitting
 - p_T -stubs to L1 trigger @ 40 MHz
 - full hit data to HLT @ 0.5-1 MHz



~8400 2S-modules

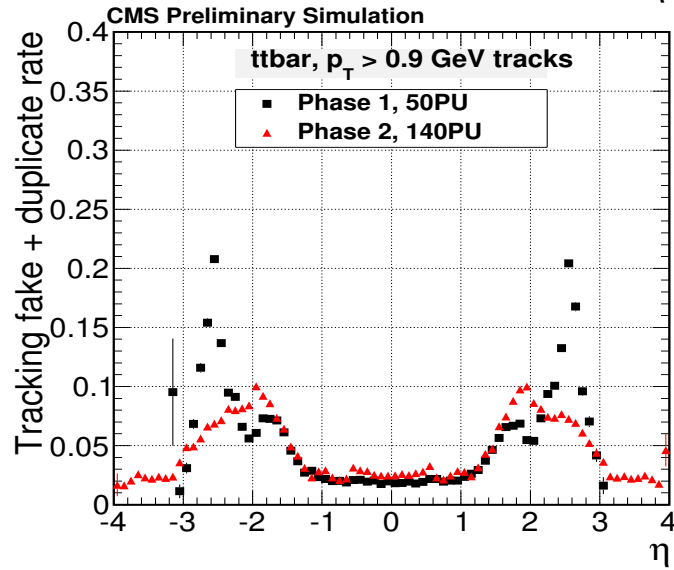





reconstructed p_T cut of $r=75\text{cm}$ layer

Efficiency, resolution and fake rate

J Mans
ECFA 2014

Outer Tracker
substantially lighter!



Pixel used in simulation results to date is identical to the Phase 1 Pixel detector with additional forward disks.
Further optimization of pixel parameters for b-tagging and forward track parameter resolution is planned

Improvements To Current Trigger

Run 2 will already require an improved trigger because of energy and luminosity increases

Have not yet exploited fully potential improvements from present detector

e.g. 10 years ago technology limited data transfer rate which enforced some data suppression

now 10 Gbps optical links are standard,

along with fast, flexible, more powerful processing



Motivation and requirements


● Motivation

- Trigger rates are driven by the increase in luminosity, the centre-of-mass energy, and by the higher PU (especially hadronic objects)
- CMS detector electronics are limited to a L1 trigger rate of 100 kHz. It will be a major challenge to maintain physics acceptance within this rate
 - Example: factor of 3 in inst. lumi x factor of 2 from CoM energy → 6 times the rate compared to 2012 → single lepton trigger thresholds ~50 GeV

● Requirements

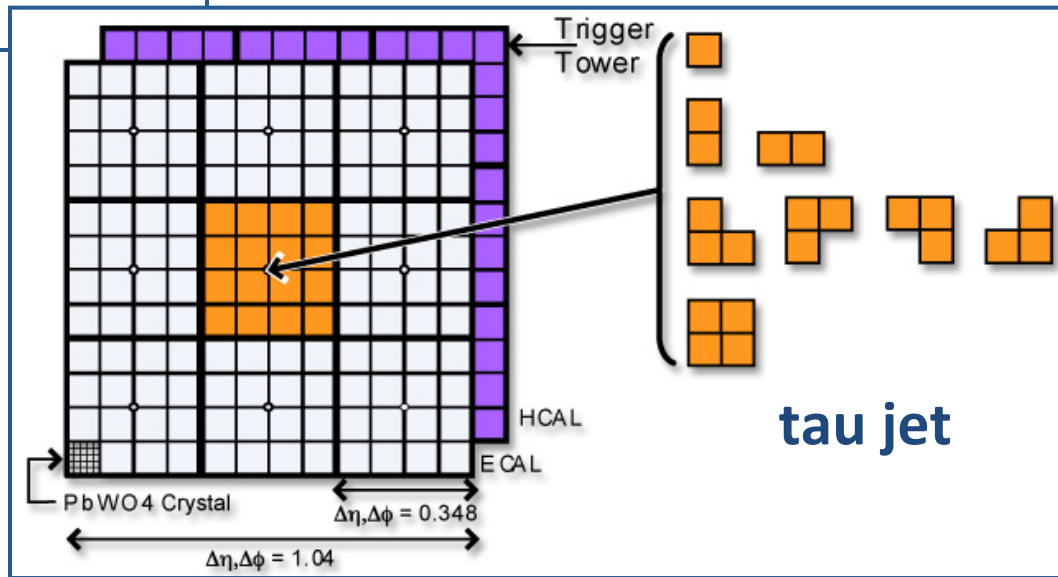
- Maintain sensitivity for electroweak scale physics and for TeV scale searches similar to what was achieved before LS1
- Making significant changes to the trigger system could put CMS data-taking at risk. The plan must allow parallel commissioning of the new trigger to mitigate this risk

Calorimeter Algorithms



- Electron/photon
 - Large deposition of energy in small region, well separated from neighbour
 - pileup worsens the separation for lower p_T objects

- jets
 - hadrons – large, likely overlapping objects
 - τ - isolated irregular, narrow energy deposits
 - simulations identify likely patterns to accept or veto



Trigger rates

Table 2.1: Extrapolated trigger rates for several current Level-1 Triggers and thresholds, for a range of instantaneous luminosity (given in $\times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$) and pile-up ($\langle PU \rangle$). The rate in the third column used a typical run in 2012, and the next two columns are determined using the high pile-up fills of 2012 at 8 TeV. There is no scaling for higher beam energy. The final column uses a Monte Carlo sample to account for the increased beam energy and out-of-time pile-up. The second column gives the offline E_T (or p_T) values at which the trigger reaches 95% efficiency, to enable a fair comparison with menus presented later.





Level-1 Trigger, 2012 Threshold	95% Threshold [GeV]	Rate [kHz] $L = 0.4$ $\langle PU \rangle = 15$	Rate [kHz] $L = 1.1$ $\langle PU \rangle = 45$	Rate [kHz] $L = 1.6$ $\langle PU \rangle = 66$	Rate [kHz] $L = 1.1$ (14 TeV) $\langle PU \rangle = 50$
Single iso e/γ , 18 GeV, $ \eta < 2.1$	26	6.3	19	27	40
Single Mu, 14 GeV, $ \eta < 2.1$	19	4.1	11	15	27
Single Jet, 128 GeV	150	1.0	3.6	6.6	14
H_T , 150 GeV	280	0.9	55	$\sim 10^3$	110
Double e/γ , 13, 7 GeV	20,13	5.4	15.4	24	47



Requirements



- Identified the following areas for upgrade:
 - Improved electromagnetic object isolation using calorimeter energy distributions with pile-up subtraction
 - Improved jet finding with pile-up subtraction
 - Improved hadronic tau identification with a much narrower cone
 - Improved muon p_T resolution
 - Isolation of muons using calorimeter energy distributions with pile-up subtraction
 - Improved global Level-1 trigger menu with a greater number of triggers and with more sophisticated relations involving the input objects

Improvements to $e\text{-}\gamma\text{-}\tau$



Calorimeter trigger: MP7


- IN
- 72 x 12.5 Gbps
= 0.9 Tbps
- OUT
- 72 x 12.5 Gbps
= 0.9 Tbps
- ~80W




G Hall

New Trigger Architecture

Conventional Trigger



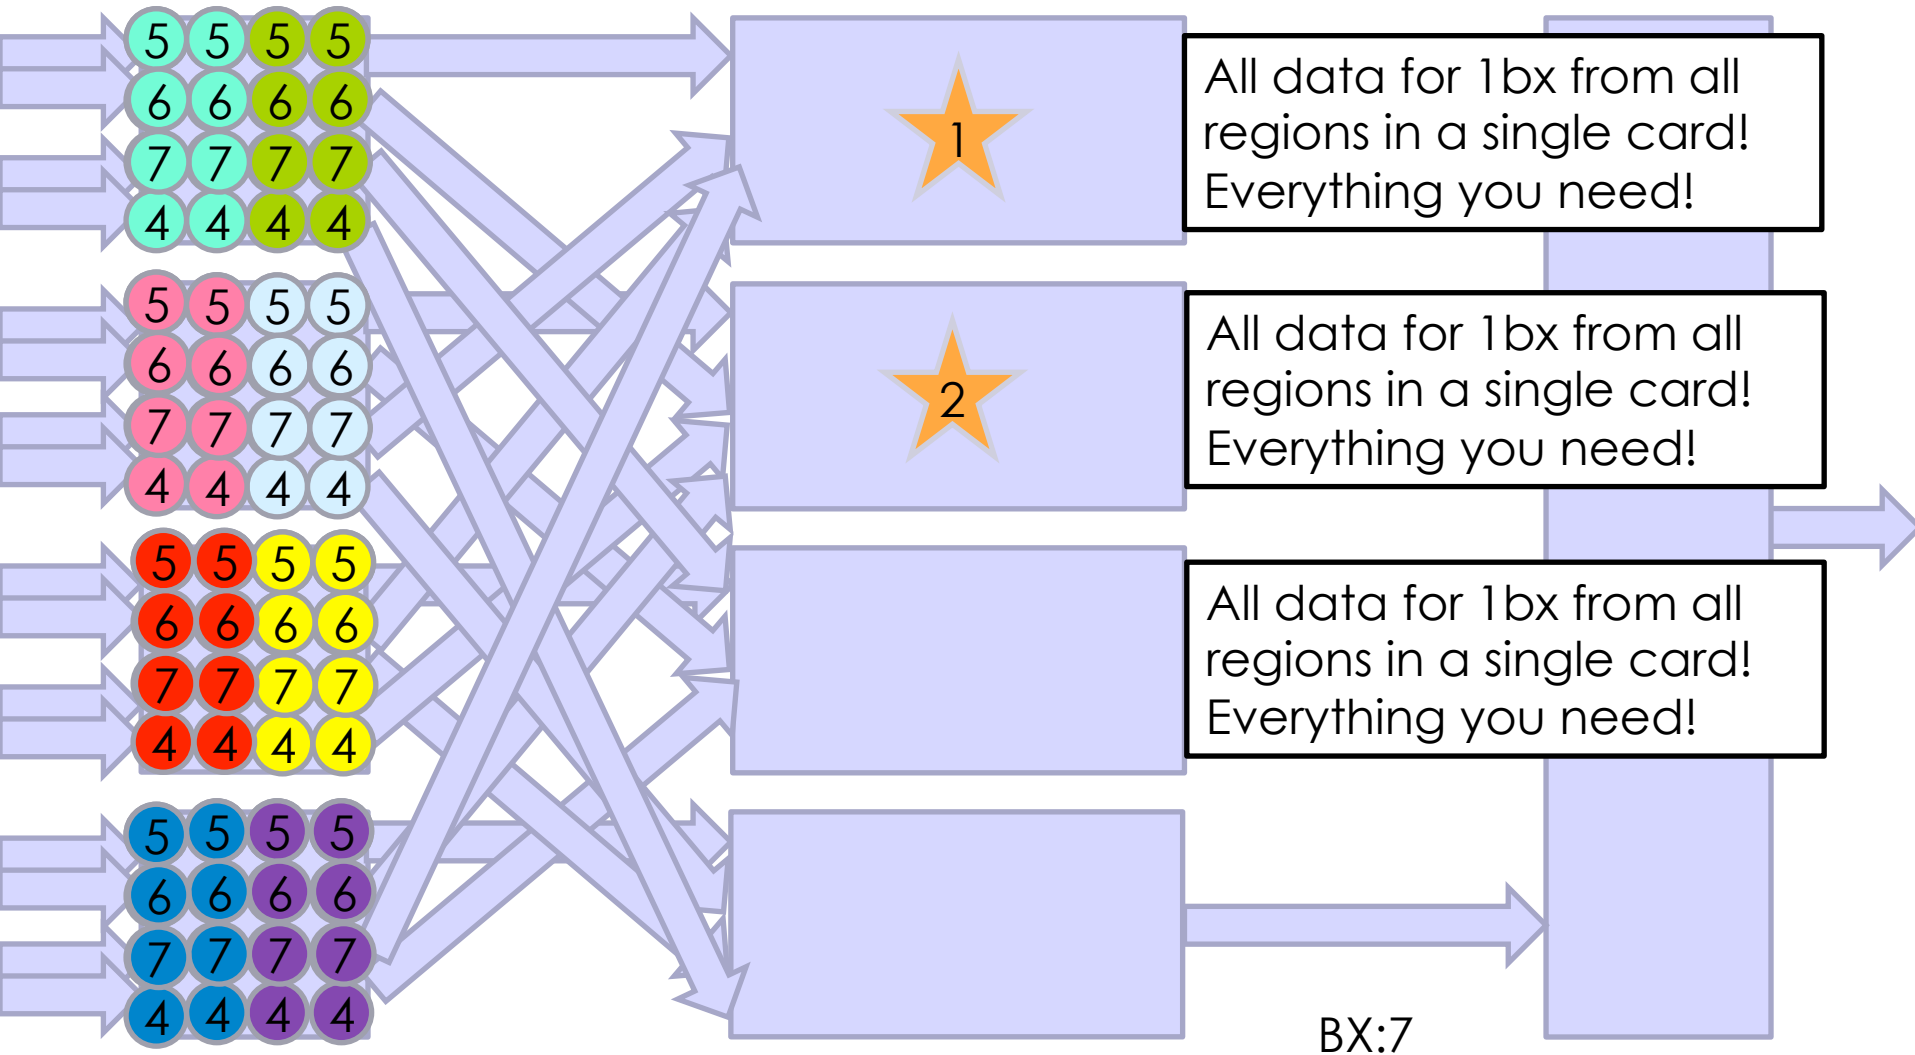
Time Multiplexed Trigger



What is a Time Multiplexed Trigger?

- Multiple sources send to single destination for complete event processing
 - as used, e.g., in CMS High Level Trigger
- Requires two layers with passive switching network between them
 - can be “simple” optical fibre network
 - could involve data processing at both layers
 - could also be data organisation and formatting at Layer 1, followed by data transmission to Layer 2, with event processing at Layer 2
 - illustration on next slide

Time-multiplexing



What are advantages of TMT?


- “All” the data arrive at a single place for processing
 - in ideal case avoids boundaries and sharing between processors
 - however, does not preclude sub-division of detector into regions
- Architecture is naturally matched to FPGA processing
 - parallel streams with pipelined steps at data link speed
- Single type of processor, possibly for both layers
 - L1= PP: Pre-Processor L2 = MP: Main Processor
- One or two nodes can validate an entire trigger
 - spare nodes can be used for redundancy, or algorithm development
- Many conventional algorithms explode in a large FPGA
 - timing constraints or routing congestion for 2D algorithms
- Synchronisation is required only in a single node
 - not across entire trigger

Calorimeter trigger current status

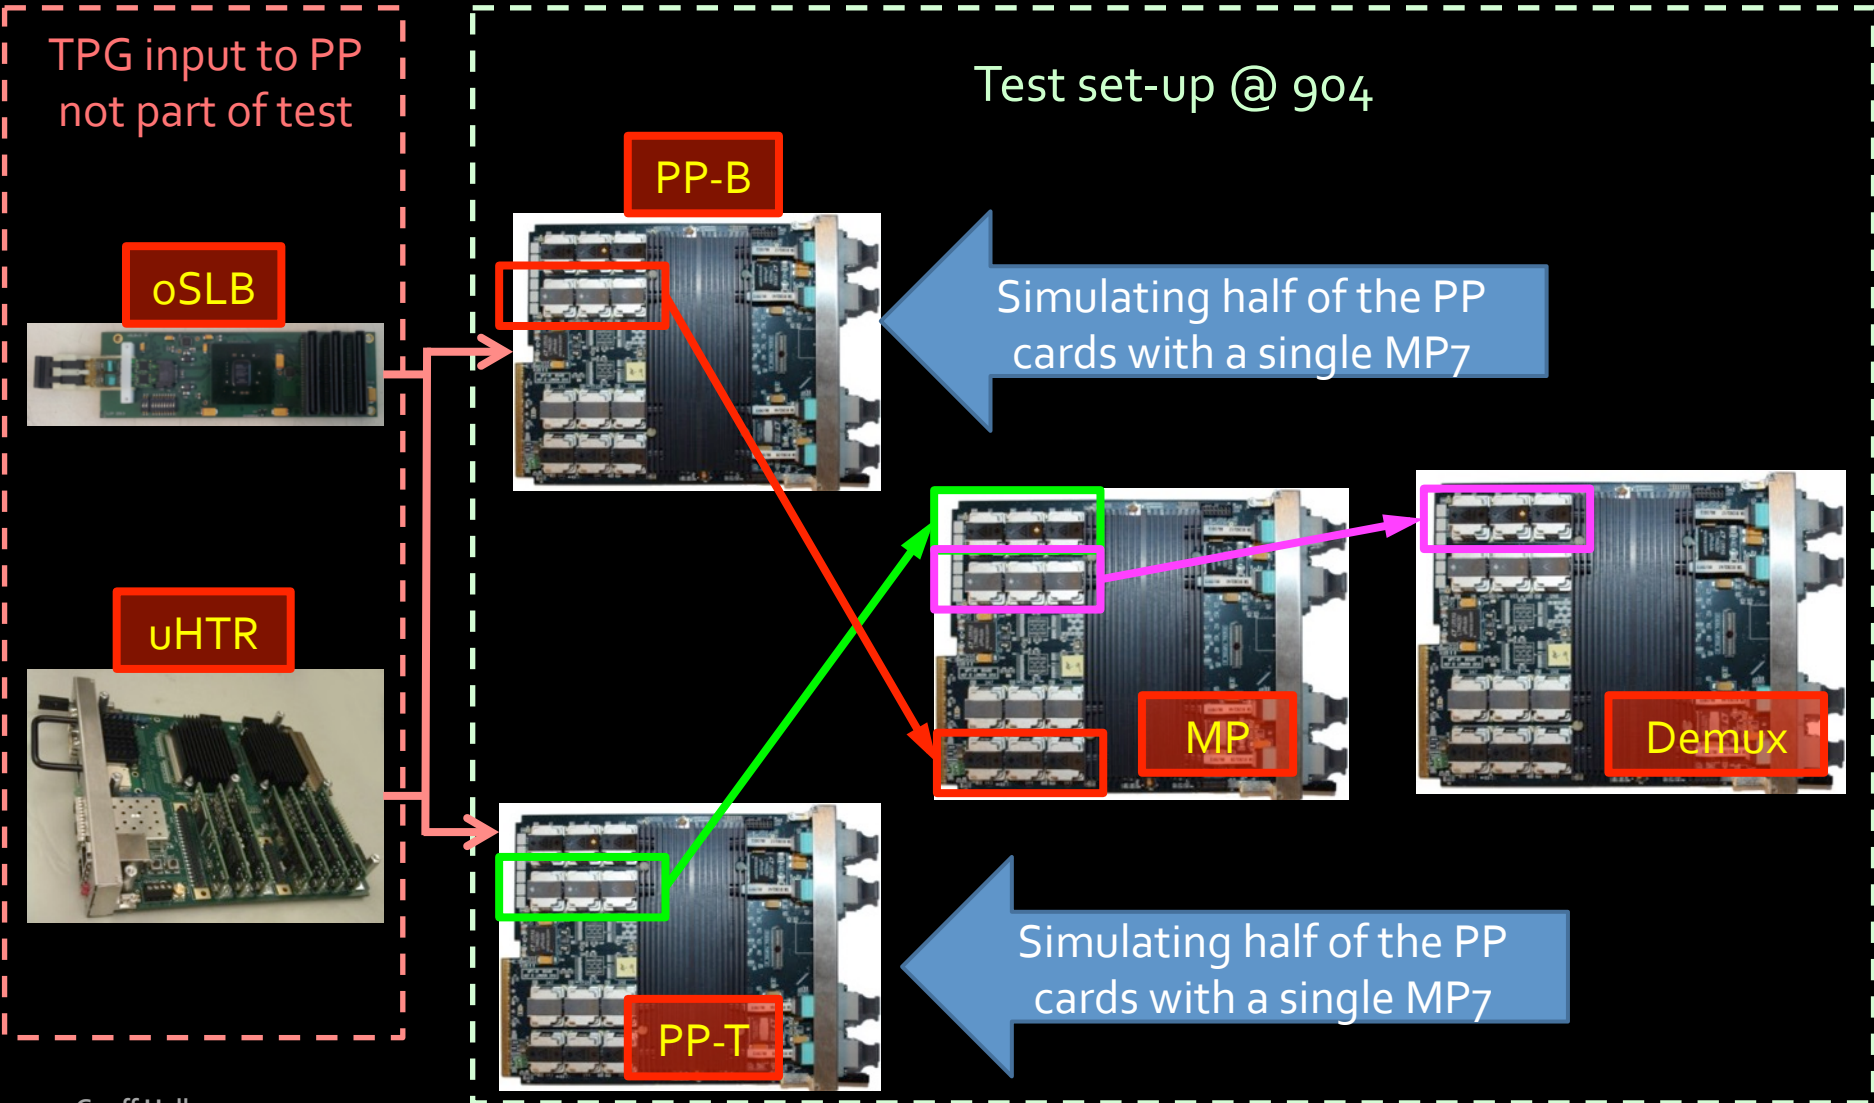
- Currently being installed at LHC Point 5 in CMS for parallel operation with the existing trigger during 2015
 - following a series of demonstrators
 - time-multiplexed data transfer from layer-1 to layer-2
 - pseudo-random data and emulated physics data with pileup
 - implementation of realistic algorithms in firmware
 - bit-level confirmation of algorithm operation – software vs firmware
 - verification of latency estimates
- Huge FPGA + huge IO = all problems solved?
 - not quite so simple...

A note on Xilinx 7-series FPGAs

- Large 7-series FPGAs are beasts!
- Scale of routing problem has outstripped performance of PCs
- Xilinx has invested heavily in improving tools
- Still have to work a lot harder at optimising designs than we used to



CMS Calo TMT demonstrator (Sep 2013)




Status of TMT jet algorithms

- Jets

- 9×9 sum of trigger towers at every site
- Fully asymmetric jet veto calculation
- Local (“Donut”) or Global pile-up estimation
- Full overlap filtering
- Pile-up subtraction
- Pipelined sort of candidates in ϕ
- Accumulating pipelined sort of candidates in η

- Ring sums

- Scalar and Vector (“Missing”) ET
- Scalar and Vector (“Missing”) HT




9×9 jet at tower-level resolution

50% LUT utilization INCLUDING links ,
buffers, control, DAQ, etc.

Runs at 240 MHz


Results (from September test)

- Implementation of an algorithm and successful transmission of data through it
- Random data passed through an emulator was used in the testing of the algorithms



Compared emulated results (solid line) with those from the MP7 (markers)

C++ emulator and hardware match precisely



Results – Latency Measurement

- Verification of latency and how it compares to TDR value -in particular the SerDes link

Source of Latency	BX (TDR)	BX - measured in Sept 2013
L1 processing + TM	10	7
L1/L2 SerDes (Tx+Rx) @ 10Gbps	5	5
L1/L2 SerDes Align Data	1	1
L1/L2 cable (20m)	4	4
L2 Processing	8	5.5 (clustering, jets, ring sums)
L2/GT SerDes (Tx+Rx)	5.5	5 (identical link to L1/L2 above)
L2/GT SerDes Align Data	1	1 (identical link to L1/L2 above)
L2/GT cable	0.5	0.5
De-multiplex	6	7
TOTAL	41	36


A simple example of Routing Congestion: 1


- (Giles) Created simple design to find routing limit
 - 30x36 2x2 tower clusters (“electrons”) with 10bit energy
 - 432 Gb/s (without 8B/10B)
 - Approximately $\frac{3}{4}$ of CMS
 - Sum 16 clusters to create “pseudojets”
 - No other firmware (e.g. no sort, no transceivers, no DAQ, etc)
 - XC7VX485T – Place & Route fails even though LUT usage only at 29%

Bare minimum
“physics” algorithm

but number of LUTs is not the whole story...

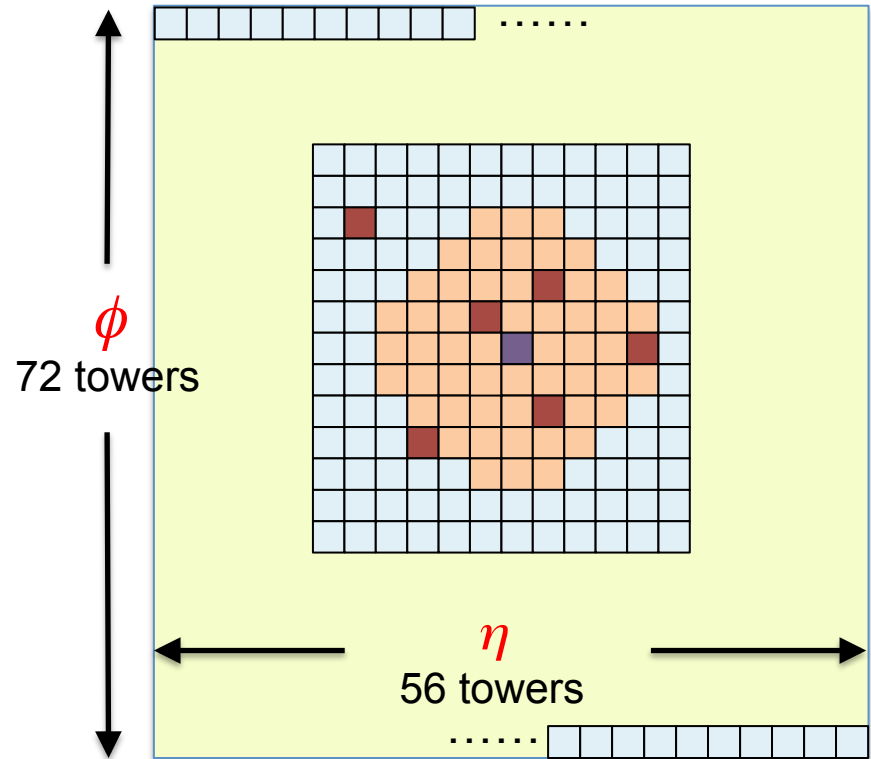
A bigger FPGA may not solve all the problems...

 Route:463 - The router has detected a very dense, congested design. It is extremely unlikely the router will be able to complete the design in an excessive run time the router will exit with a partially routed design. This behavior will allow you to identify difficult design constraints, putting too much logic into this device, or an issue with the implementation. If you would prefer a design that uses less logic, remove some logic from the design, or change the placement before running router again. If you are willing to accept a long run time behavior.

 Route:543 - This design is experiencing routing congestion. Please review the Xilinx Routing Optimization White Paper, [in resolving this issue.](#)

Routing Congestion 2 : A nasty example..

- (Giles) Implemented a proposed circular isolation algorithm
 - using pipelined design
- Searches every tower location in 56 x 72 region
 - 4032 sites
- Counts the number of objects above threshold within a circular ring of diameter 9 towers or clusters
 - Result passed into LUT with the energy to determine object status



Operates up to 400MHz
Compact: < 1% of the FPGA
Low latency - 9 clks (no overlap)
1.5 BX @ 240 MHz


* Only synthesised 36 towers in eta, rather than 56, but in the small FPGA

Why was Time Multiplexed Trigger not already used?

- Mainly technology limitations
 - It is reliant on high performance hardware
 - large & powerful FPGAs
 - many high speed (optical) links
- More recent objections to latency penalty in L1-L2 transmission
 - but this is mostly a myth!
 - If properly organised, data processing does not need to wait for entire event data.
 - It can begin as soon as first cycle's worth of data arrive

The track-trigger challenge

- Impossible to transfer all data off-detector for decision logic so on-detector data reduction (or selective readout) essential
 - The hit density means high combinatorial background
 - 99% tracks < 2 GeV/c
 - after ~ 2 GeV/c cut,
 - ~ 60 - 70% hits associated with real track
- What are track-trigger requirements?
 - original assumption: a few points would help
 - but simulations did not show big rate reductions
 - hence quasi-full track reconstruction for $p_T > \sim 3$ GeV
 - separation of primary vertices is required
- application in trigger based on plausible simulations
 - i.e. 2025 trigger not well defined
 - isolation, pileup subtraction, vertex association, object pointing,....





Possible layout of CMS TM Track-Trigger

split tracker into **phi** regions

constrained problem by looking at **minimum** number of trigger regions (TRs) required, and imposed constraint that **one module cannot be shared across more than two TRs**


1 GeV/c




- 5 TRs in phi only

- 1 GeV/c boundary region assumed

- e.g. could allow for better reco @ 2GeV/c in case of e^+e^- , brem, low p_T , multiple scattering etc.



the time multiplex period is not a completely free parameter

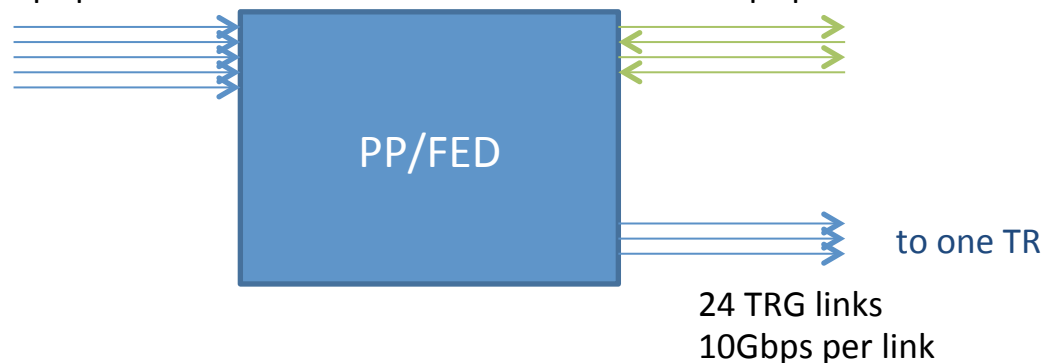
small TM period	large TM period
full event must be quickly assembled into one MP	could allow more efficient processing of pipelined data into MP
reduces data volume per event from PP to MP (or requires increased number of links)	increases data volume per event from PP to MP (or reduces number of links)
reduces latency	increases latency
reduces number of MPs	increases number of MPs
min ~15bx (PP output bandwidth without more Trigger Regions)	max ~34bx (68 links/2 Trigger Regions)
	 preferred direction

TM period of 24BX chosen for case study (could be optimised in future)

from non-shared modules

68 FE links
3.2Gbps per link

4 bidirectional DAQ links
10Gbps per link




24 TRG links
10Gbps per link


from shared (boundary) modules

68 FE links
3.2Gbps per link

4 bidirectional DAQ links
10Gbps per link




48 TRG links
10Gbps per link




4 DAQ links per PP/FED allows a maximum bandwidth of 40Gbps (~588Mbps available per tracker module)

Organisation for 5 TRs in Phi & 24BX TM Period



(numbers from tkLayout...)

full tracker L1 and trigger data can be read out with a total of **353 MP7s**



after 24BX, **all trigger data** from tracker from first event are assembled into **5 MPs**
each MP corresponds to a trigger region in phi

But processing should have started earlier than BX 25

subsequent events are available in next set of 5 MPs after one extra BX

no boundary sharing is required after duplication in PP;


no post-TM removal of duplicates necessary

What processing is possible in MPs?

track-finding? track fitting? data processing before an AM stage?

2 MP7s emulate event data from 1 out of 5 regions, one out of every 24BX

emulated PP/FEDs
PP->MP links
PP links/MP
PP->MP links total
TM nodes



- This demonstrator already exists
 - just need to program source data to be ready to try algorithms

phi1

phi2

phi3


phi4

phi5

Firmware design

- Still to establish the best way of finding tracks at HL-LHC
 - latency and efficiency, as well as (firmware) programming challenges
 - we know from MP7 that large FPGAs present **serious** challenges, e.g.:
 - exceeding RAM resources
 - logic fails to synthesize within timing constraints after many hours
- Possible approach based on Hough transform
 - locate series of hits on a trajectory, for which $y = mx + c$
 - For fixed (m,c) : every “y” corresponds to single “x”
 - For fixed (x,y) : every “c” corresponds to single “m”
 - point (m,c) -> line (x,y) point (x,y) -> line (m,c)
 - All hits from a real track have same (m,c)
 - For each data point (x,y) , hypothesize “m” and calculate “c”
 - When multiple hits have same (m,c) , send for fitting
 - identify by histogramming entries into array

pipelined systolic array as proposed by Andy Rose



stub enters array, close to entry row and finds its way

Summary

- The TMT is now a proven architecture in CMS
 - which will operate in the CMS calorimeter trigger from 2016
- The hardware is very flexible and can be deployed for a TMTT
 - only a fraction of the system is required to validate the concept
 - installing and building should only require replicating identical nodes
 - a track-trigger could be built using present technology
 - safe to assume further technological progress in next decade
- The next challenge is to prove algorithms can be implemented
 - under way

Backup



Trigger Algorithm	Current Level-1 $L = 1.1 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$			Current Level-1 $L = 2.2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$		
	Rate [kHz]	95% Threshold [GeV]	Plateau Efficiency	Rate [kHz]	95% Threshold [GeV]	Plateau Efficiency
Single e/γ	12	46	1.0	10	67	1.0
Single iso e/γ	10	38	0.9	9.4	52	0.9
Single Mu	12	23	0.95	11	42	0.95
Single iso Tau	10	65	0.3	9.2	72	0.3
iso $e/\gamma + e/\gamma$	10	24 15	0.9	16	26 16	0.9
Mu + Mu	6.3	18 10	0.9	7.4	20 12	0.9
Tau + Tau	7.5	36 36	0.1	8.2	36 36	0.1
iso $e/\gamma + \text{Mu}$	9.6	21 11	0.85	6.2	24 12	0.85
Mu + e/γ	3.3	18 14	0.95	5.0	20 15	0.95
Single Jet	6.4	170	1.0	5.4	205	1.0
Double Jet	4.6	140 140	1.0	5.8	170 170	1.0
Quad Jet	9.4	4@71	1.0	4.8	4@96	1.0
Single iso $e/\gamma + \text{Jet}$	7.5	32 68	0.9	8.5	38 82	0.9
Single Mu + Jet	8.6	22 43	0.95	7.5	27 54	0.95
Single iso $e/\gamma + H_{\tau}^{\text{miss}}$	10	29 110	0.9	8.2	38 120	0.9
Single Mu + H_{τ}^{miss}	4.6	18 89	0.95	9.8	20 93	0.95
H_{τ}	3.9	500	1.0	5.4	580	1.0
Total Rate	94			92		

Trigger Algorithm	Current Level-1 $L = 2.2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$			Upgraded Level-1 $L = 2.2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$		
	Rate [kHz]	95% Threshold [GeV]	Plateau Efficiency	Rate [kHz]	95% Threshold [GeV]	Plateau Efficiency
Single e/γ	10	67	1.0	11	57	1.0
Single iso e/γ	9.4	52	0.9	15	31	0.90
Single Mu	11	42	0.95	14	22	0.90
Single iso Mu	NA	NA	NA	15	19	0.82
Single Tau	NA	NA	NA	12	100	0.95
Single iso Tau	9.2	72	0.3	13	83	0.7
iso $e/\gamma + e/\gamma$	16	26 16	0.9	12	23 16	0.9
(iso)Mu + Mu	7.4	20 12	0.9	9.4	15 10	0.8
(iso)Tau + Tau	8.2	36 36	0.1	7.2	64 62	0.67
iso $e/\gamma + \text{Mu}$	6.2	24 12	0.85	11	21 10	0.85
(iso)Mu + e/γ	5.0	20 15	0.95	8.3	18 15	0.83
iso $e/\gamma + \text{Tau}$	NA	NA	NA	8.3	21 57	0.86
isoMu + Tau	NA	NA	NA	5.8	14 47	0.8
Single Jet	5.4	205	1.0	5.9	205	1.0
Double Jet	5.8	170 170	1.0	4.2	130 130	1.0
Quad Jet	4.8	4@96	1.0	5.0	4@55	1.0
Single iso $e/\gamma + \text{Jet}$	8.5	38 82	0.9	11	27 78	0.90
Single Mu + Jet	7.5	27 54	0.95	9.7	18 52	0.93
Single iso $e/\gamma + H_T^{\text{miss}}$	8.2	38 120	0.9	12	27 110	0.90
Single Mu + H_T^{miss}	9.8	20 93	0.95	11	18 86	0.93
H_T	5.4	580	1.0	3.0	380	1.0
Total Rate	92			95		

Synchronisation & Structure

- Synchronization is only required per-time-node, not across the whole trigger
 - De-synchronization of a node only affects that node.
 - A timing glitch in a CT disrupts the whole trigger.
- The efficient pipelined logic should lead to lower latency, and the eventual clock speed can be as fast as the FPGA allows.
- Firmware build times should be significantly shorter due to the pipelined, as opposed to combinatorial, nature of the architecture.

first wafer probed manually



	Total	Genuine		Combin.		Unknown	
Layer B1	2500	1573	(62.9%)	788	(31.5%)	139	(5.6%)
Layer B2	1778	1350	(75.9%)	243	(13.7%)	185	(10.4%)
Layer B3	1263	947	(75.0%)	110	(8.7%)	206	(16.3%)
Layer B4	818	557	(68.1%)	68	(8.3%)	193	(23.6%)
Layer B5	642	430	(70.0%)	27	(4.2%)	185	(28.8%)
Layer B6	520	327	(62.9%)	14	(2.7%)	179	(34.4%)
Layer E1	2326	1795	(77.2%)	468	(20.1%)	63	(2.7%)
Layer E2	899	669	(74.4%)	178	(19.8%)	52	(5.8%)
Layer E3	493	351	(71.2%)	97	(19.7%)	45	(9.1%)
Layer E4	469	297	(63.4%)	109	(23.2%)	63	(13.4%)
Layer E5	342	239	(69.9%)	37	(10.8%)	66	(19.3%)
Layer E6	265	179	(67.6%)	17	(6.4%)	69	(26.0%)

Table 3.7: Average total number of stubs produced per event in each barrel and endcap layer, as well as the number and proportion of the total of each class of stub, averaged over 100 events using the single seed stub producer.