

CMS R&D for HL-LHC L1 Tracking



Kristian Hahn – Northwestern University
on behalf of the CMS Track Trigger Working Group



ACES 2016 – Fifth Common ATLAS CMS Electronics Workshop for LHC Upgrades
CERN – 03.08.16


The HL-LHC

High luminosity is crucial for future CMS physics goals

- Instantaneous: $5\text{-}8 \times 10^{34} \text{ cm}^2\text{s}^{-1}$
- Integrated: 3000 fb^{-1} (10 yr)


Technical Proposal:

<http://cds.cern.ch/record/2020886/files/LHCC-P-008.pdf?version=3>



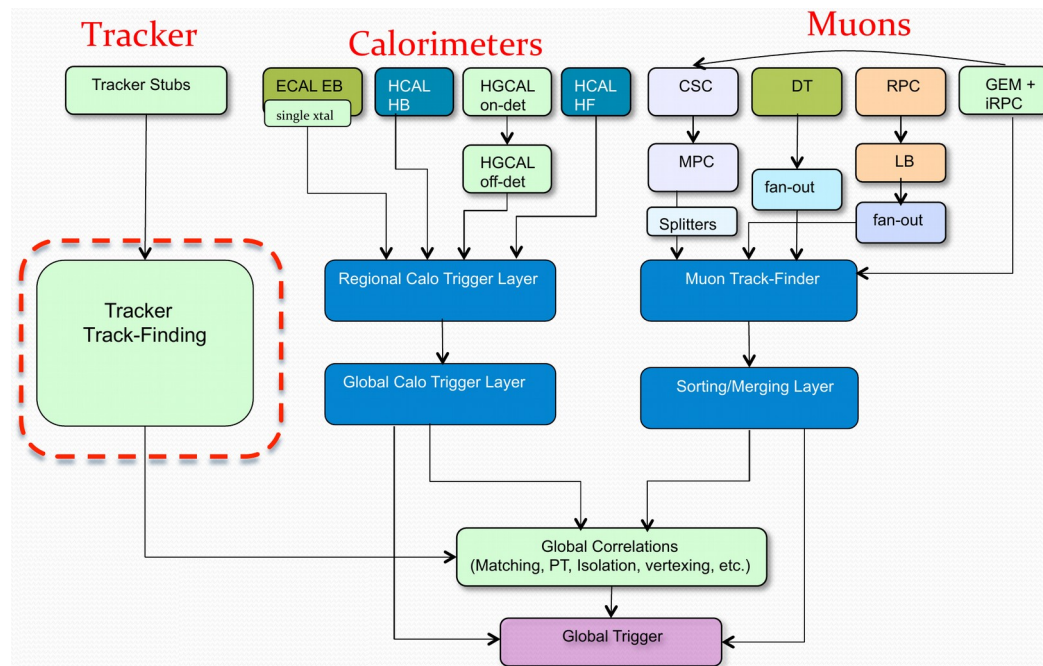
And is incredibly challenging for the CMS trigger!

- Up to 200 pile-up interactions
 - In and out of time ...
- $\sim 6\text{k}$ tracks / BX
- While maintaining sensitivity to “low-mass” (eg: Higgs) physics





HL-LHC trigger must utilize *tracking information @ L1, otherwise ...*

- EG rate @ 25 GeV > 100 kHz
- Muon rates plateau
- Total rate > 1.5 MHz
- ...




Greatly improves L1 object reconstruction

- Lowering rates
- Lowering thresholds
- Improving / maintaining physics sensitivity





L1 tracking enabled by Phase-2 Outer Tracker design

- 6 Barrel + 5 Endcap layers
- Trigger sent to BE via GBTs
- Pixel information not part of L1

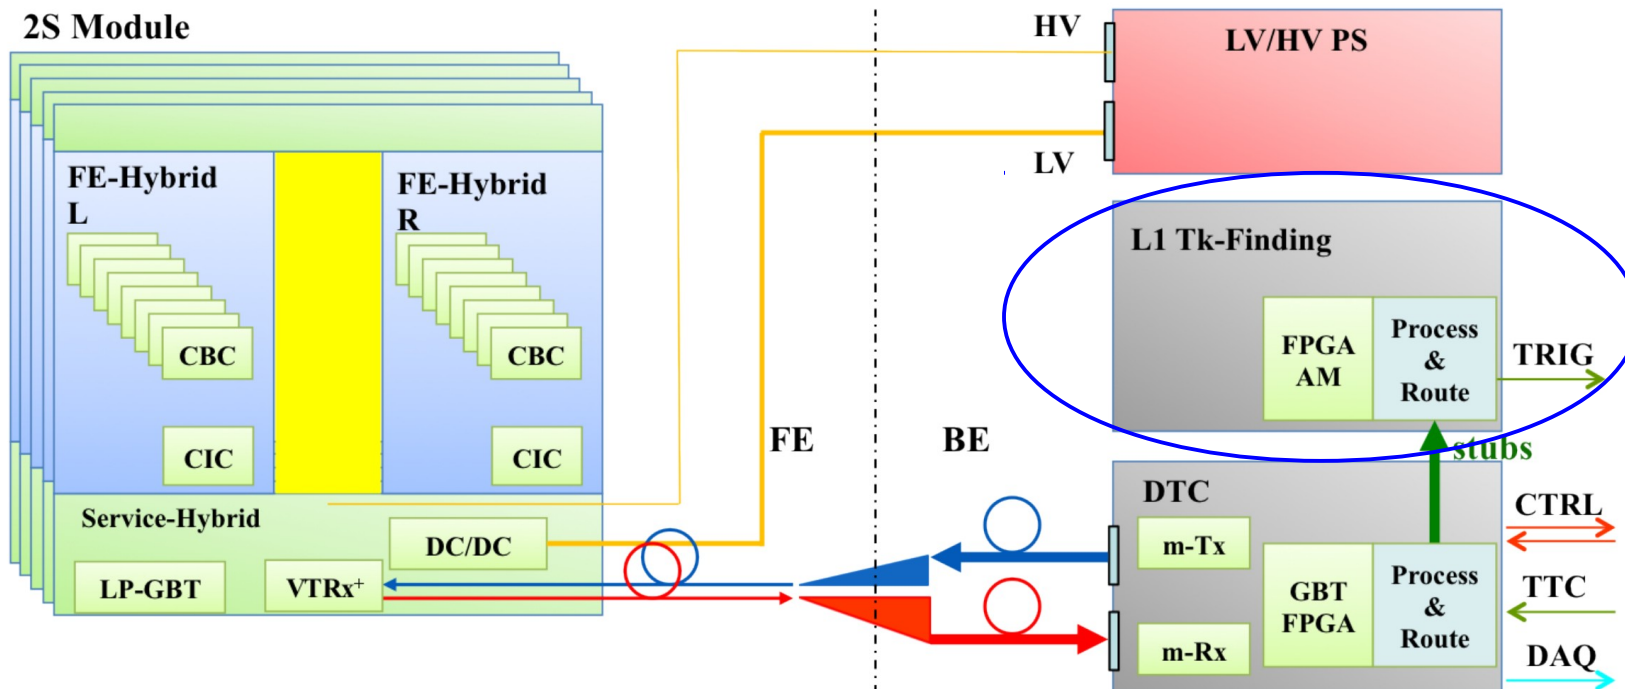


pT (bend) measurements at the detector front-end

- Double-sided modules, correlate pairs of hits into “stubs”
- Send only data that corresponds to high momentum tracks
- Immediate rate reduction by a factor of >10!



Front End ↔ Back End



Data from front-ends received at Data, Trigger Control Board (DTC)

- Likely an ATCA form factor, 50-70 modules per DTC
- FPGA emulation of GBT links to the front-ends
- Approximately 4x100Gpbs outputs to L1Tk system

L1Tk system builds and delivers track primitive to the L1 trigger

- On-going R&D for this system is the focus of this talk ...

(see talk from F. Vasey on Wednesday for more details on the Tk FE / BE)

L1 tracking finding poses formidable technical challenges!

- Data rates > 50-100 Tbps
- Occupancy & combinatorics: $O(10^4 \text{ hits/BX}) @ 200 <PU>$
- Latency: 4+1 μs for tracking

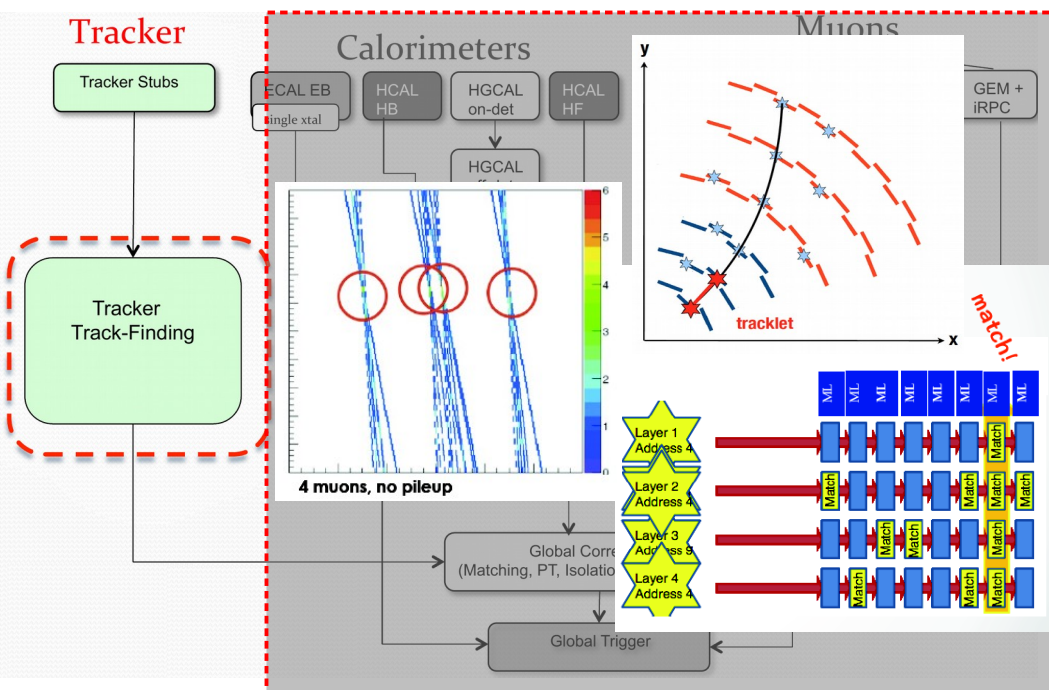


Three CMS R&D efforts tackling the problem

- 2 FPGA-only
- 1 ASIC-assisted

No precedent for a tracking trigger at this scale ...

- All options must be explored!
- Each working toward technical demonstration by Q4'16





CMS Track Finding R&D



The algorithmic challenge : combinatorics & pattern recognition

- L1Tk In general: stub/track association → track fitting → track cleaning
- Approaches primarily differ in their means of stub-track association
- Each seeks to significantly reduce number of stub combinations to process


The technical challenge: rates & latency

- All approaches will leverage high-speed links (optical, backplane, FMC)
- Likely all will target ATCA or similar
- All will seek to exploit modern, dense FPGAs
 - Additional challenges with pattern density (AM), though trade offs wrt FPGA performance
- Technical demonstration will indicate how far technology must evolve ...


Rest of the talk : detailed review of the 3 CMS approaches

- Overviews & some recent highlights
- Emphasis on demonstration setups, rather than on final system visions

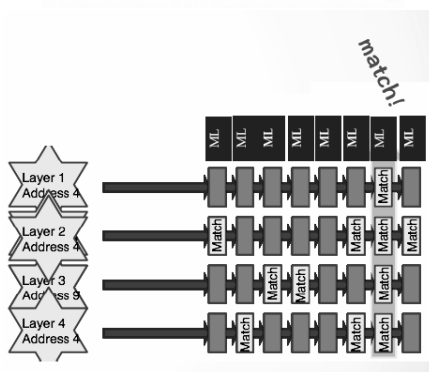
CMS Track Finding R&D



TMT Hough




Tracklet



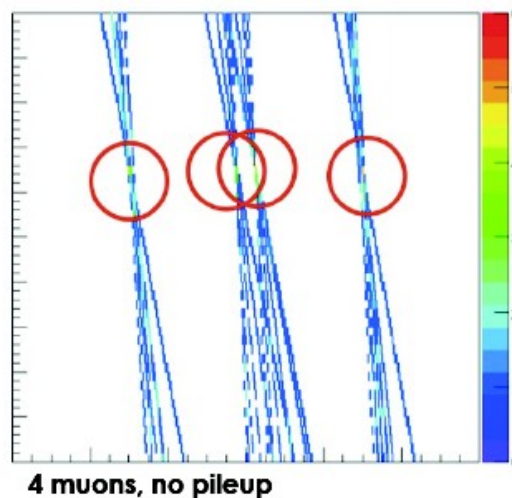
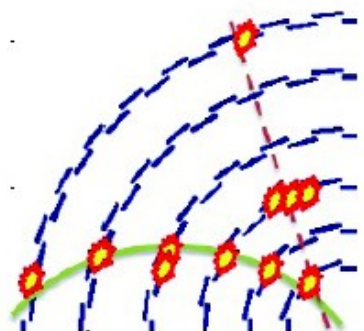
AM+FPGA

TMT / Hough Transform

- **Fully time multiplexed** → all module data for given BX to a single L1 track finding board (mux factor: presently 24)




- **FPGA-based pattern recognition & track fitting**
 - r - Φ Hough transform
 - Linearized χ^2 track fit



Circular trajectories in the r - ϕ plane

- Described with two parameters: q/pT , ϕ_{65}
- Create 2D histogram of q/pT , ϕ_{65}
- For each stub, fill those bins for which a track with corresponding q/pT , ϕ_{65} passes through the stub
 - Individual stubs mapped to lines
 - Where lines cross: a track!



3D pattern recognition via equivalent procedure in the r - z plane

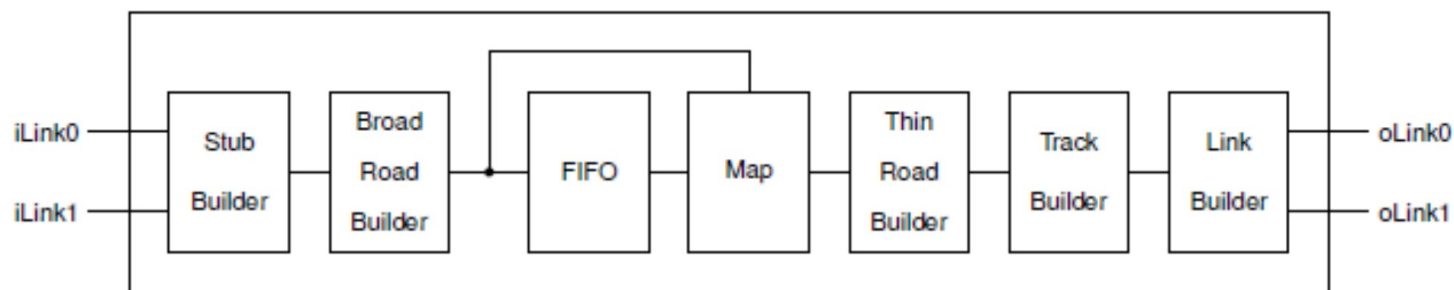

- Simpler methods in r - z also explored

Resolution and segmentation

- Tracker divided into $5\eta \times (36 \text{ or } 64)\phi$ sectors with independent HTs

HT implemented as a pipelined array

- In a 36ϕ sector scheme
- A single track finding board receives data from all 36 sectors
- Each sector processed with an independent 32×32 HT (below)



32 q/p_T columns processed in parallel, stubs from ≥ 5 stub tracks then sent on for track fitting

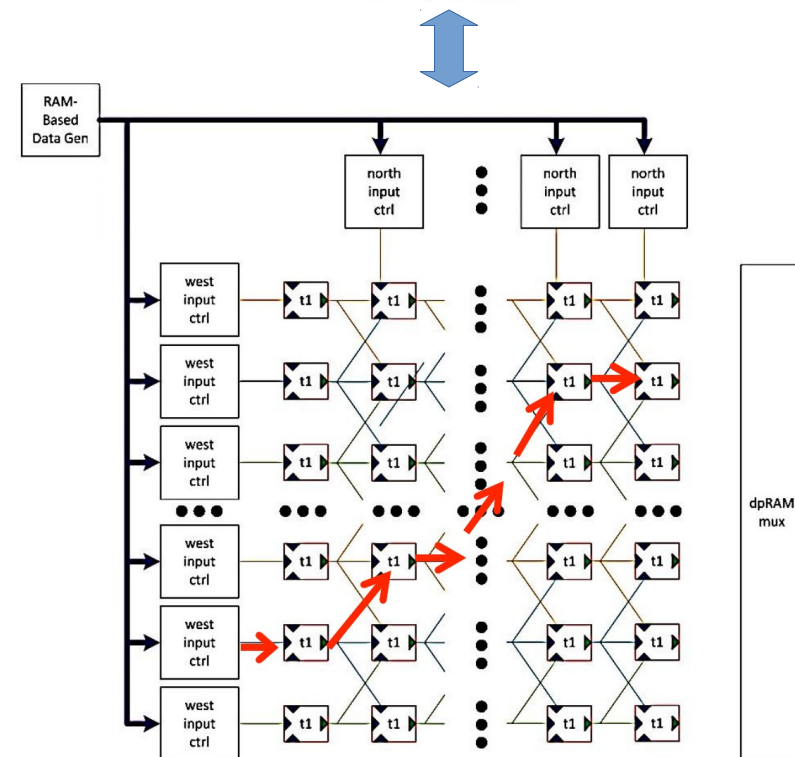

- Virtex 7 currently fits 8/36 sectors
- Firmware redesign in progress to increase #sectors

HT implemented as a systolic array

- Each cell in the array corresponds to an HT bin
- A single HT array processes all sectors
 - Presently 14x12 cells in Virtex 7
 - Requires larger number of sectors (64) to achieve desired granularity

- Stubs injected on periphery of array
- Stored in each compatible cell
- Then passed on to neighboring cells

Timing met at 240 MHz, 61% utilization (~20% from infrastructure logic, eg: links, control)



Bend Filtering

Utilize pT consistency to reduce fake rate

- Stub bend value checked against pT coordinate of HT array
- If stub is not consistent, don't store

BEND FILTER

❖ Only add stub to HT cell if stub bend is consistent with predicted bend based on q/Pt of cell & (r,z) coords. of stub.

$$| \text{bend}(\text{stub}) - \text{bend}(\text{predicted}) | < 2.5 * \text{bend_resolution}$$

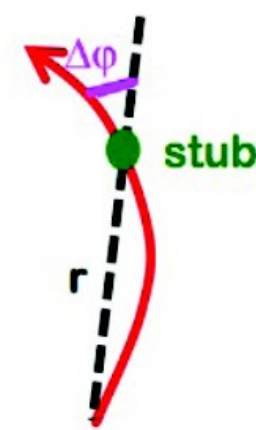
with: $\text{bend}(\text{predicted}) = \Delta\phi / \rho$,

where:

$$\Delta\phi = \text{const} * r(\text{stub}) * q/Pt(\text{cell})$$

and:

In Barrel: $\rho = \text{sensor_pitch} / \text{sensor_separation}$
 In Endcap: $\rho = (\text{sensor_pitch} / \text{sensor_separation}) * (z/r)_{\text{stub}}$



Bend information also employed in the other L1Tk approaches

Goal: demo 1 geometric/time slice of the full system


- Target resolution: 6mrad φ , 0.0012 q/pT
- Mock-up future, larger FPGAs w/ >1 daisy-chained boards

Demo uses Imperial MP7 (μ TCA)

- Virtex 7 690T, 0.94 Tbps FP I/O
- 6x12 channel miniPOD, 11.3 Gbps / link
- Expertise from CMS Phase-1 Calo trigger

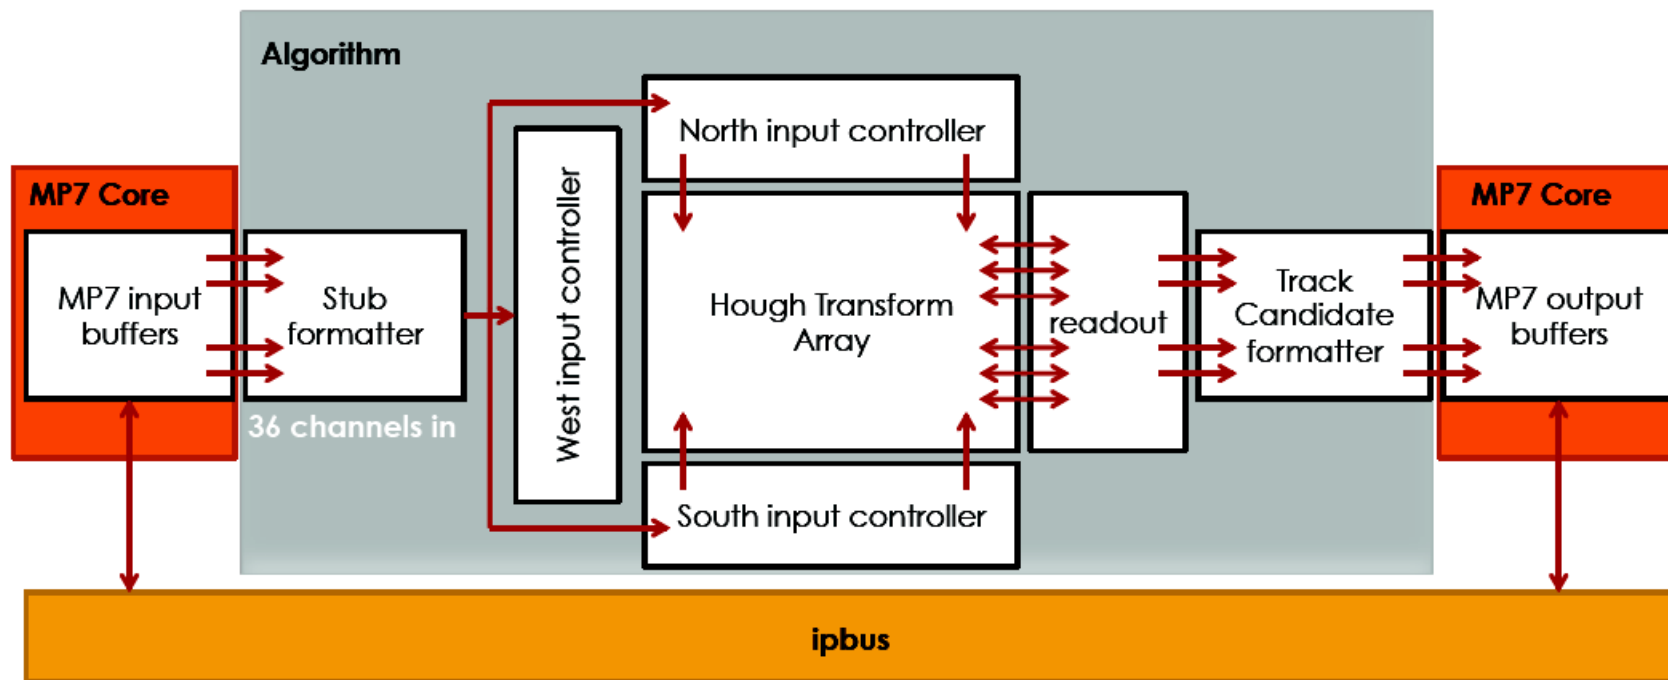
System presently a single track finding board and a data source

- Running both pipelined and systolic HT firmware implementations



<http://www.hep.ph.ic.ac.uk/mp7/>

Block diagram of preliminary demonstration firmware (systolic)



- MP7 core firmware and IPBus interfaces well established
- 14x12 HT → 61% logic utilization: 25% from array, rest glue logic

Focused now on evaluating performance of pipeline vs systolic

- Collapse to a single implementation, freeing cycles for downstream development (eg: r-z stub consistency)



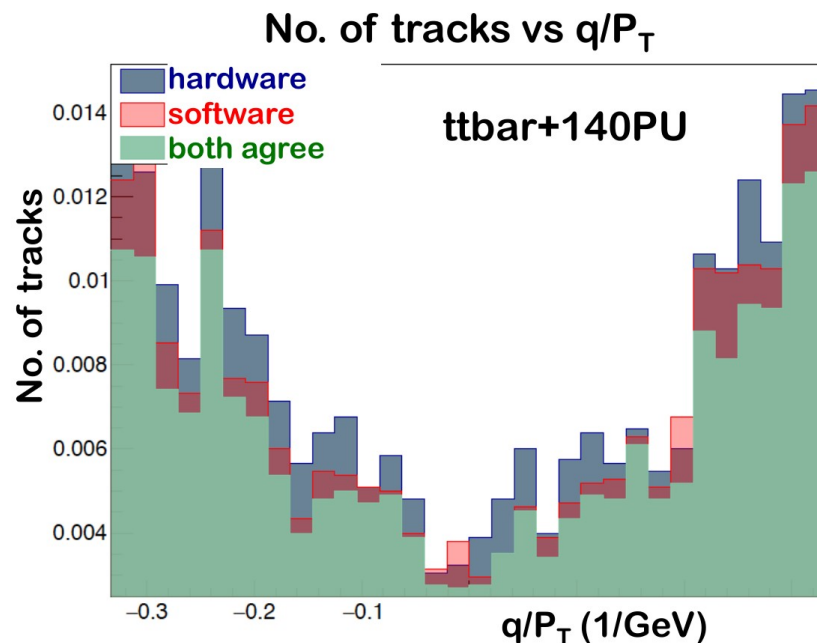
TMT Preliminary Results

Input event data from simulation, compare kinematics of found tracks in HW with SW


- Fairly good agreement ... ~10% discrepancy under investigation

HT latency proportional to #stubs

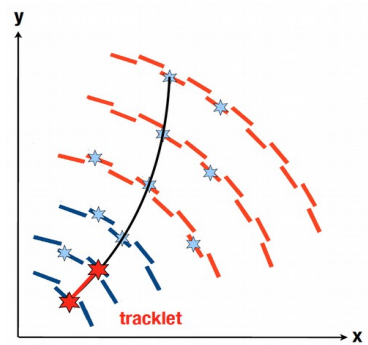
- ~2 μs in busiest $t\bar{t}$ +140 PU events
- Includes TMUX period, results depend on ultimate data rates/segmentation
- Corresponding efficiency, fake rate to be estimated



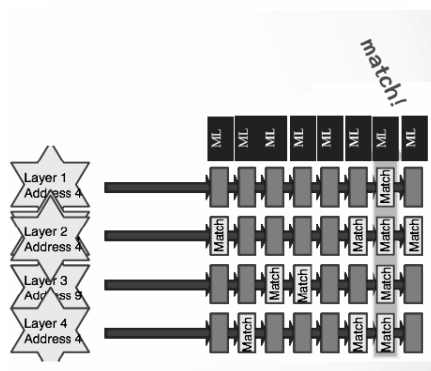
CMS Track Finding R&D



TMT Hough



Tracklet



AM+FPGA

Stubs in pair of layers form a seed (tracklet)


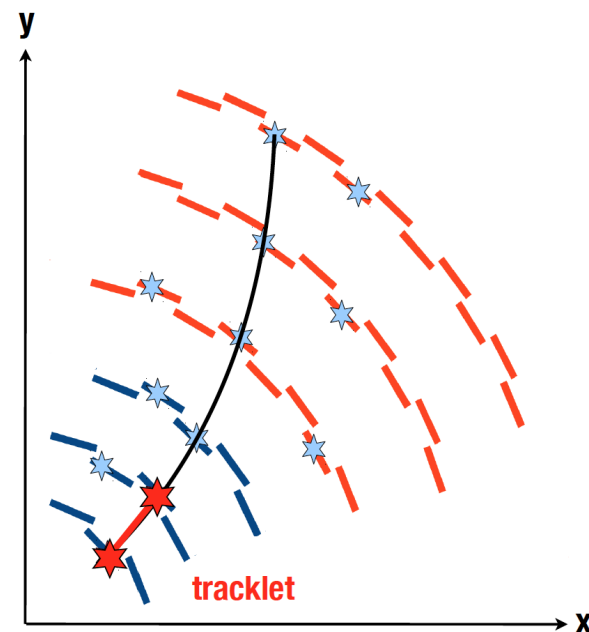
- Seeds in layer pairs 1+2, 3+4, 5+6
- Use tracklet + IP to project into other layers
- Look for stubs near the projected track
- Refit final track parameters using all stubs

To handle combinatorics, split detector into 28 phi sectors (spanning η)

- 1 sector primarily processed by 1 board
- 2 GeV tracks will be in ≤ 2 φ -sectors
- **Neighboring board I/O required**


Timing model

- Factor of 4 TMUX envisioned
- **Fixed-time processing model**, truncate if a time allocated for a given step is exceeded



L1 Tracking

(DTC emulation separately)




Total of 4 boards (one shared for input & output)

One board acts as DTC emulator that sends input stubs, track sink that receives output tracks

3 sector boards that have a central sector and its positive and negative phi neighbors

Used to test and validate the algorithm including board-to-board communication



Tracklet demonstrator @ CERN

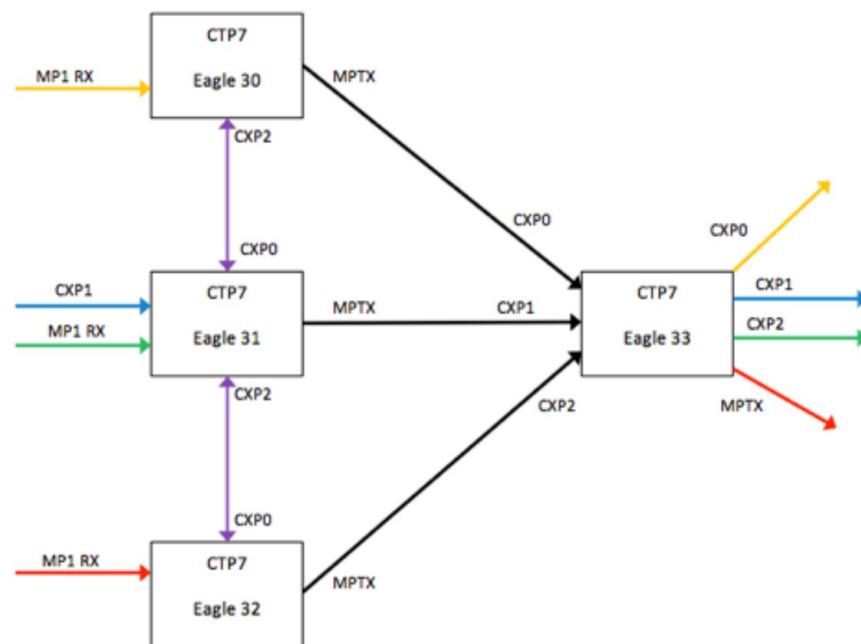
Using Wisconsin CTP7 (μ TCA) boards for the demonstration

- Used in the CMS Phase-1 Calo trigger
- Virtex-7 690 FPGA
- Zynq SoC with dual Cortex-A9 ARM
- GTH: 80 RX & 61 TX
- 3 CTP7 boards are used as sector boards, 1 CTP7 handles sending input stubs and receiving output tracks

AMC13 card is used for central clock distribution

Mixed-mode I/O @ 10 Gbps

- 12Tx/31Rx on miniPOD
 - For communication with data source
- 3x12 Tx/Rx on CXP + links @ 32 bits
 - For projections & matches
- I/O sufficient for demonstrator needs



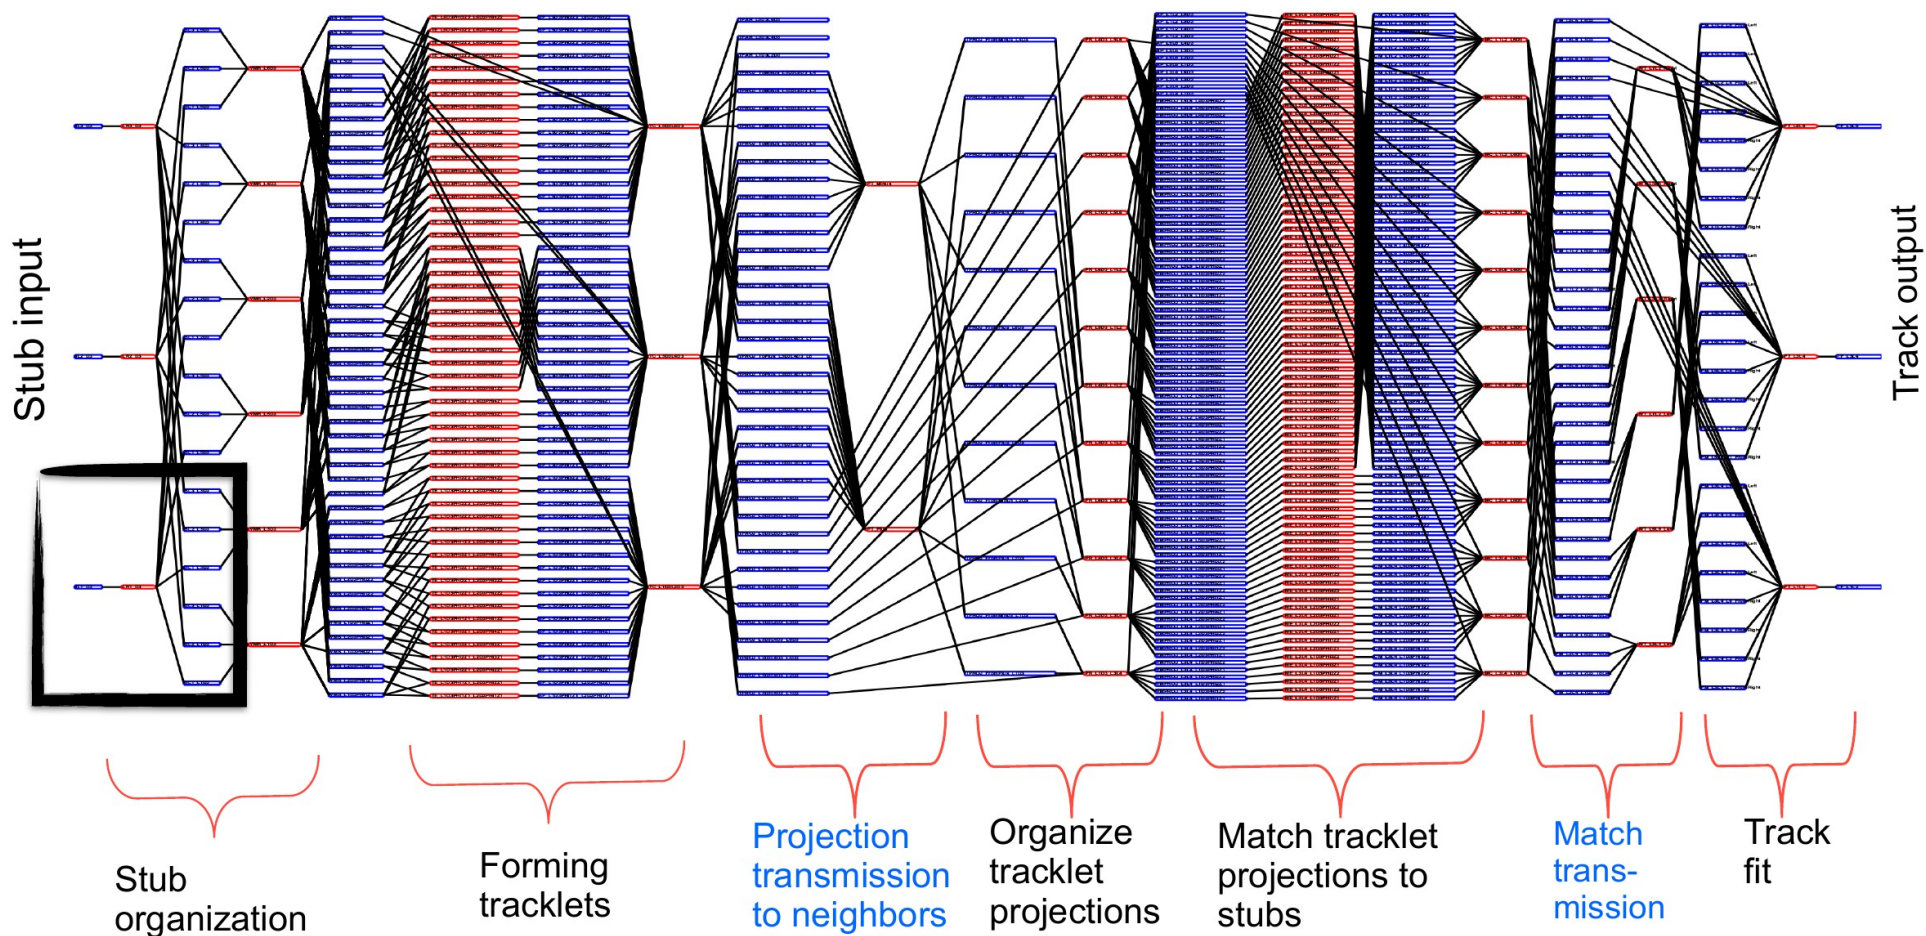
Inter-board communication based on Wisconsin's protocol

- 8b/10b encoding, 10 Gbps link
- Protocol produces data packets (24 words)
- Sent every 100ns (4x TMUX)
- Channel bonding via sync FIFOs

Board synchronization via AMC13


Tracklet Firmware

Tracklet diagram for 1/4 barrel in a ϕ sector.



Firmware and emulation generated with SW configuration tool!

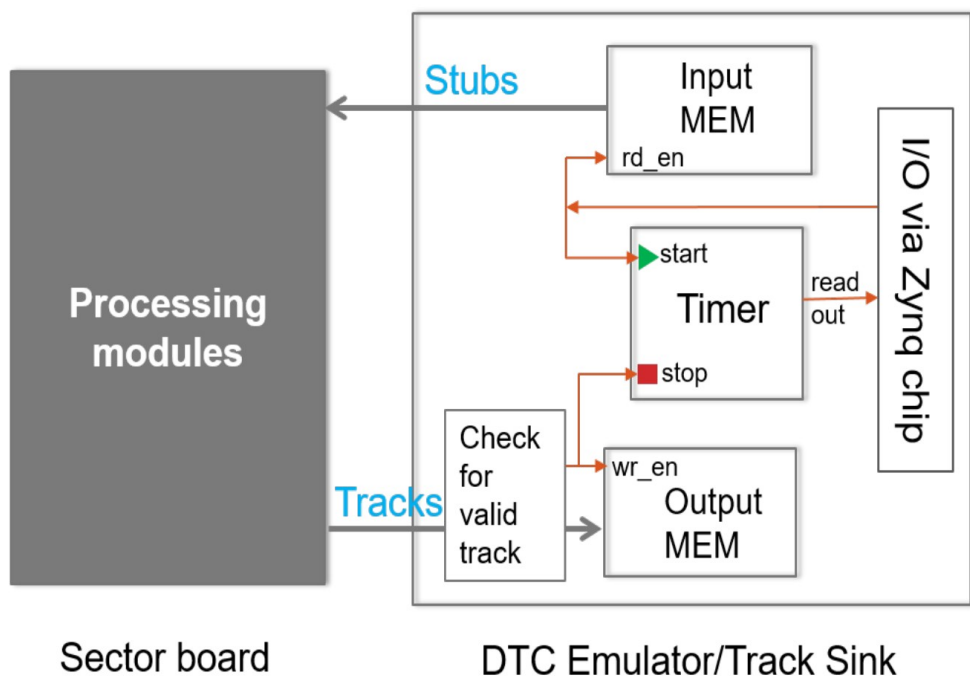
- Memories
- Processing modules



Present implementation is $\frac{1}{4}$ Barrel for a full ϕ sector

- Work underway on expanding to $\frac{1}{2}$ Barrel and including Disks
- Expected resource usage should permit $\frac{1}{2}$ Barrel for the demonstration

Tracklet Latency



- For **single track**, TMUX factor of 4,
 - 3.2 μ s measured for TMUX = 6
- **Measurements of efficiency & fake rate in realistic events TBD**

- **Measurement with a clock counter**
 - 240 MHz clock
 - same as processing clock
 - implemented on the DTC emulator
 - start: read enable of input memory
 - stop: write enable of output memory
- **Measured latency includes**
 - stub input links from DTC board to sector board
 - processing of each step, including “inter-board” communication
 - track output links back to DTC board
- **Result:**
 - 638 clock cycles
 - or 2658.3 ns


Preliminary Results

Great agreement between HW, emulation and simulation for low occupancy events

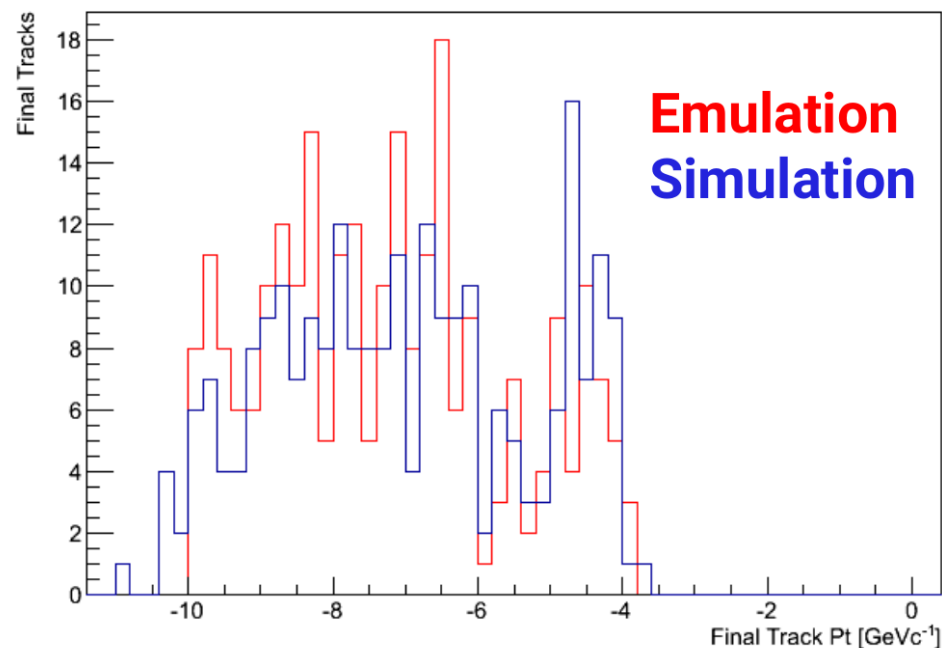
- Eg: single or double tracks

Differences observed for higher occupancy events

- Some of these expected due to a few missing features in the firmware




Single Track Validation




High Occupancy events

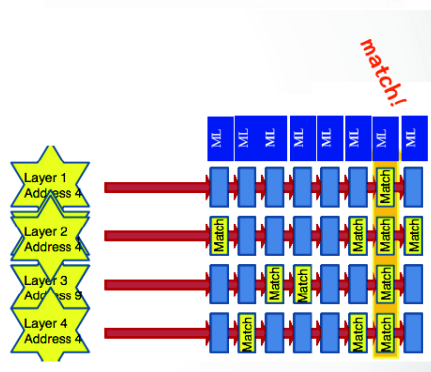
CMS Track Finding R&D



TMT Hough



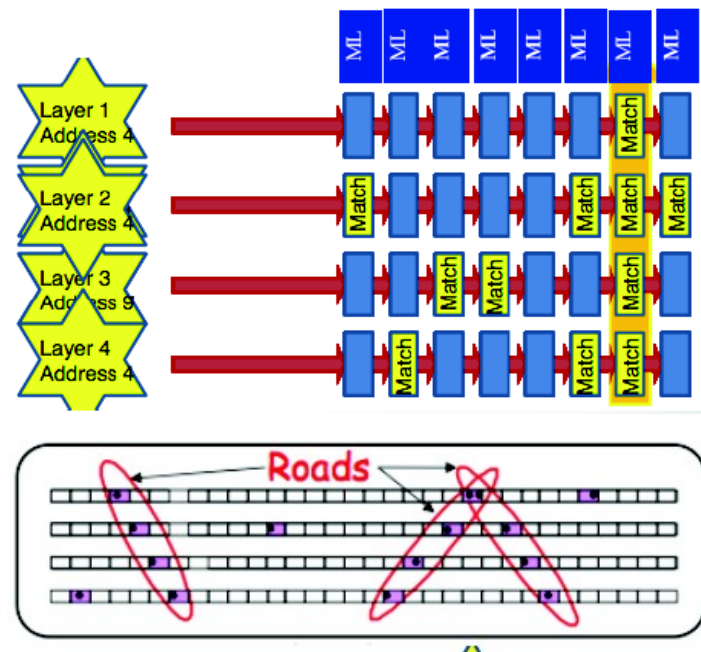
Tracklet



AM+FPGA

Parallelize computationally expensive PR tasks with AMs *match!*

- CAM cells + majority logic
 - CAMs match hits in layers to those stored
 - Hits="SuperStrips", coarse groupings of stubs
 - ML associates hits to across layers, matching patterns
 - "Roads" = coarse, multi-layer hit patterns
- Completes as soon as all hits arrive



Simplifies downstream track fitting!


- Avoids power-law dependence of execution time on occupancy
- Already employed to this end in L2 applications (SVT, FTK)

Finer grained PR and param estimation from FPGA track fitting

The AM + FPGA Approach



Split detector into trigger “towers”, nominally 48 ($6\eta \times 8\phi$)

- 3 types of towers (barrel, endcap, hybrid)



Utilize modern, high-speed ATCA backplane

- Nominally 1 shelf per tower, 1 mezzanine per BX




Pulsar 2b

Fermilab ATCA board: PRB and data source for the demo

- Virtex 7 690T, GHT to RTM (40), FMC (12) & backplane (28)
- Version 2.0 RTM with 10 QSFP+ (version 1.0 shown)
- Used also in ATLAS FTK Data Formatter
- Production run for CMS demonstration underway, target summer 2016


Extensive I/O testing performed with several backplanes

- Best performance from 40G+ Comtel “AirPlane”
- Full shelf, with all GTH @ 10 Gbps, untuned: BER < 10^{-15}



Trigger tower demo with 2 ATCA shelves

- One to emulate front-end modules
 - 400 modules/fibers
 - No requirements on the DTC, pass through
- Other for Pattern Recognition Boards
 - Receive, format and deliver data to Pattern Recognition Mezzanines
 - X10 (x2) time MUX via backplane (PRM)



Data Sourcing Boards



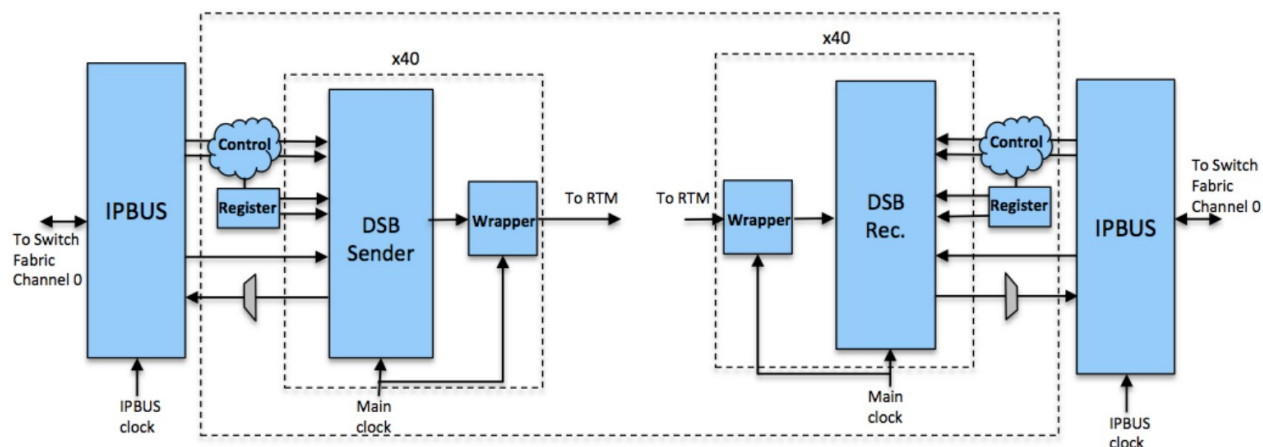


Pulsar 2b Firmware



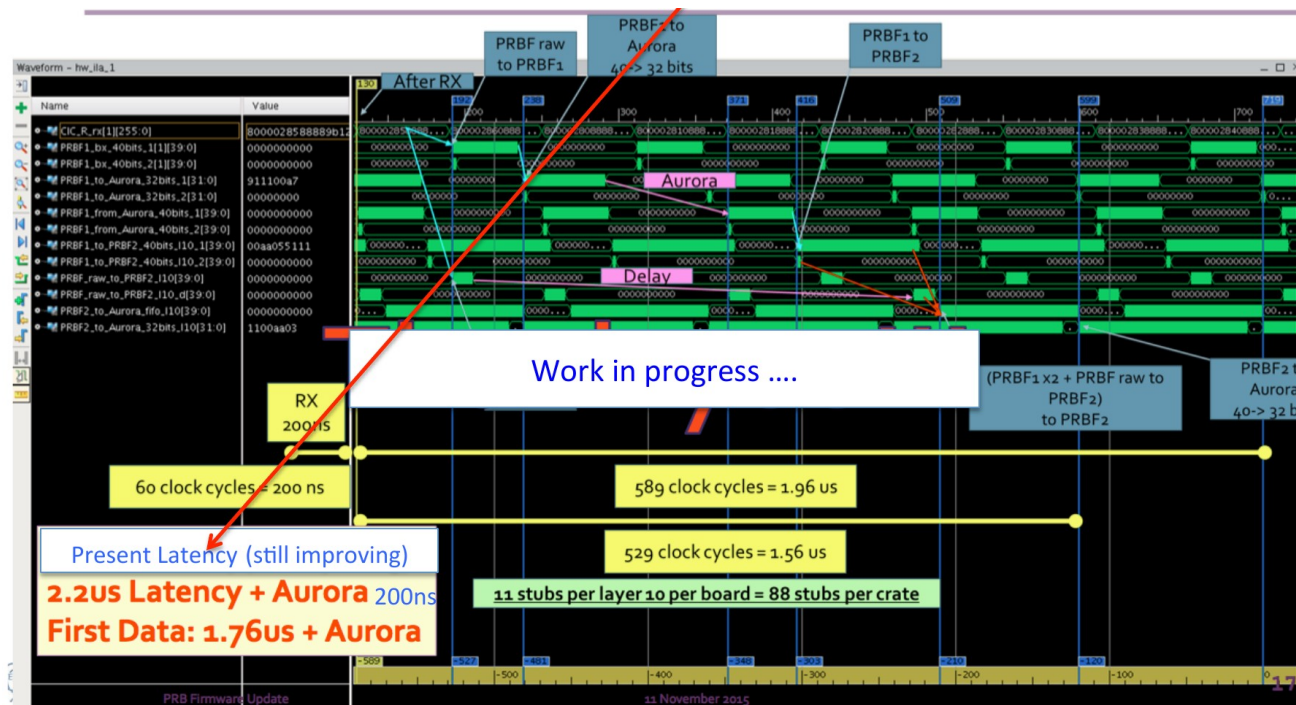
Data sourcing send & receive functionality operational

- Tested in single board loopback
- Data loading & retrieval via IPBus



1st version of data delivery firmware operational

- Tested in single board loopback
- Present latency ~ 2 μ s
- Optimization underway




Board-to-board integration tests will begin soon

AM for the Demonstration



Goal: > ~200k patterns / chip @ > 200 MHz

- Target of a long term, multi-prong R&D effort
- 28 nm (500k patterns), VIPRAM ...



2-prong, near-term approach for the demonstration

- **INFN AM06 / INFN PRM**
 - AM technology from ATLAS FTK
 - ~1.5M patterns per tower
 - Actual ASIC, realistic multichip I/O
 - but not designed for L1 application (latency)
- **FPGA emulation / FNAL PRM**
 - 2-tier AM design *emulated on Ultrascale* Kintex, ~4k patterns per mezzanine
 - Low latency, approx. clock accurate to 2-tier VIPRAM design



Currently with AM05 chips


- Double width FMC, **Kintex 7 FPGA**
- ~32k patterns

AM06 PRM in development

- ~1.5M patterns with 12 AM06 chips, enough for all trigger towers
- **Ultrascale Kintex FPGA**
- Design frozen, target mounting of 1st prototype in May


AM06 status

- First batch (ATLAS/FTK) arrived in Jan. 2016
- **No issues observed**
- Qualification for CMS underway
- First chips for CMS ~May ...




Tested / qualified using evaluation board and Pulsar 2b @ FNAL


- ✓ Electrical and mechanical
- ✓ Power dissipation
- ✓ FMC LVDS electrical connection, slow rate communication, high rate (400 MHz DDR) tested for a few lines
- ✓ High speed serial links between AM05 ↔ FPGA tested up to 2 Gbps, BER < 10^{-14}
- ✓ High speed serial links between host FPGA ↔ PRM FPGA tested at 8 Gbps, BER < 7×10^{-15}
- ✓ JTAG channels (FPGA & AM05)
- ✓ AM05 configuration (serdes & CAMs)




AM->FPGA @ 2Gbps



FPGA->AM @ 2Gbps



Kintex7->Ev.board @ 8Gbps



Ev.board->Kintex7 @ 8Gbps

Double width FMC


2 QSFP+

2 Ultrascale FPGAs

- KU040/KU060 – 2FFVA1156C
- GTH up to 16.3 Gbps
- One to emulate AM functionality before chip arrives
- Other to executing Data Organizer functionality and track fitting

36 Mb low latency DDR 2+ SDRAM



Socket for next generation VIPRAM



3 KU040 cards in hand, several KU060 FMCs available in spring 2016

Integration with Pulsar2b and GTH testing

- FMC LVDS
- QSFP+ GTH
- PRM US ↔ PRB V7
- PRM US1 ↔ PRM US2
 - le: local bus

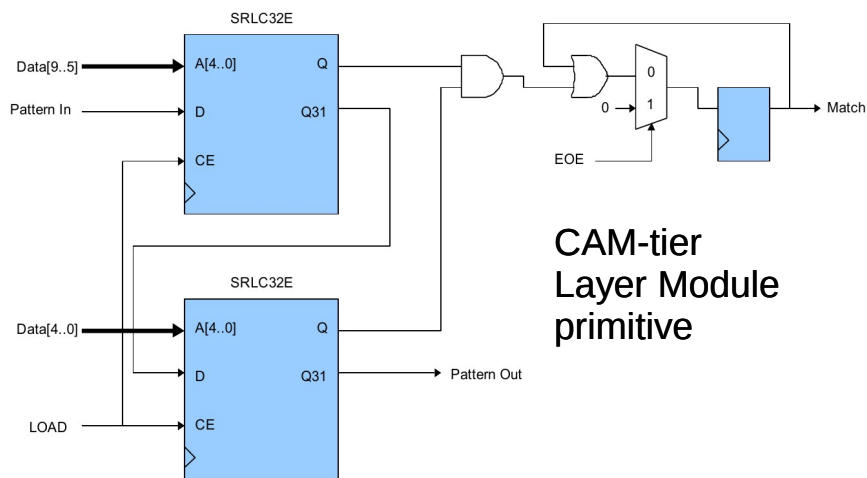


All lines error free!

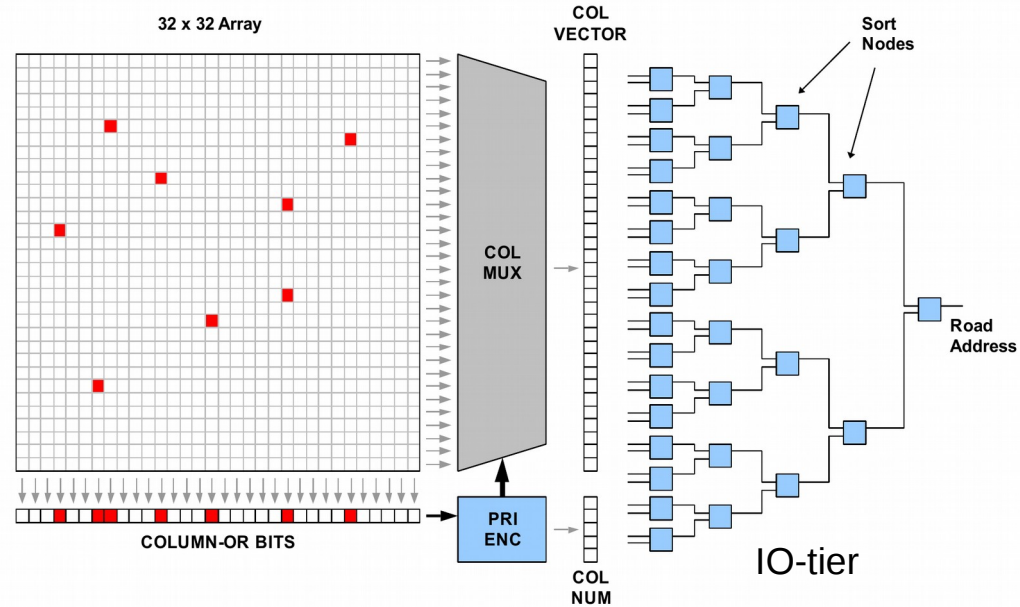
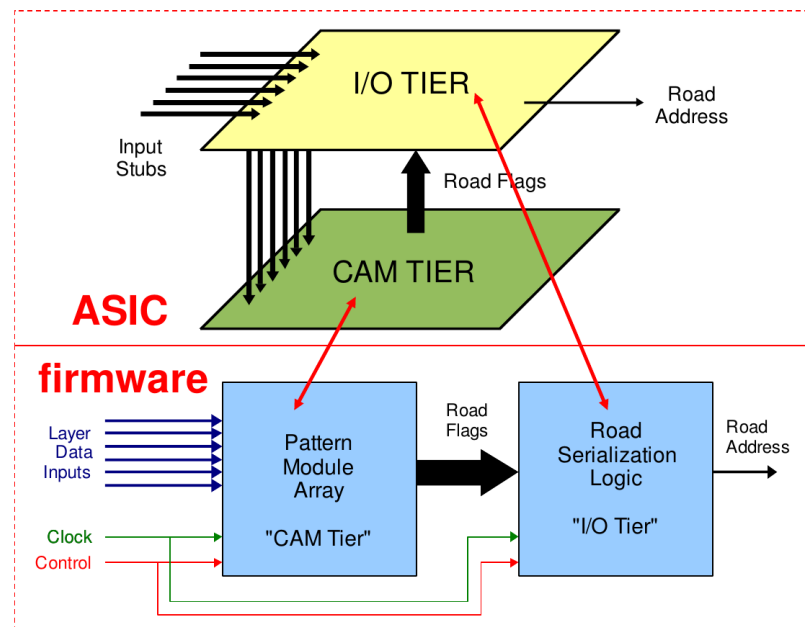
- All GTH up @ 10 Gbps, local bus up to 15.6 Gbps

VHDL model of 2-tier VIPRAM

- “CAM-tier” processes stubs for event N
 - Store 32b pattern in SLRC32E shift reg
 - Input data as 5b address



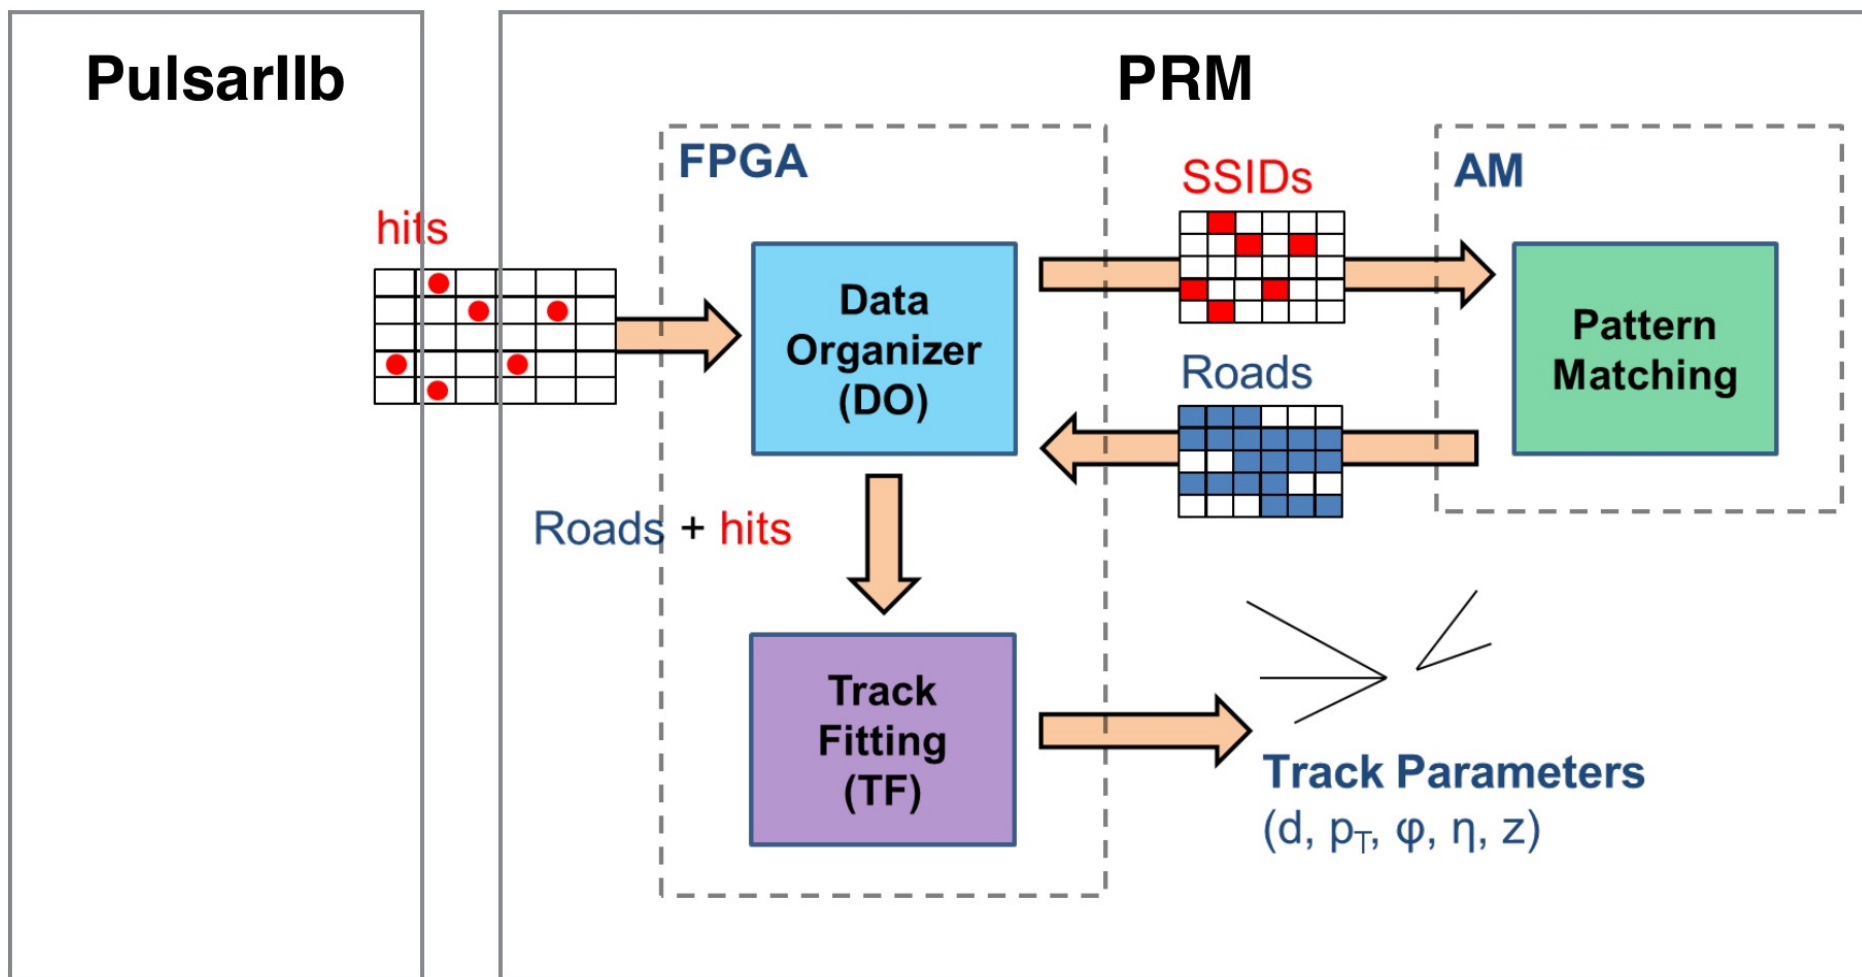
- “I/O-tier” captures road flags, outputs road addresses for event N-1
- 1-4k patterns, fully programmable
 - 1st road 7 clks after EOE



PRM Data Flow

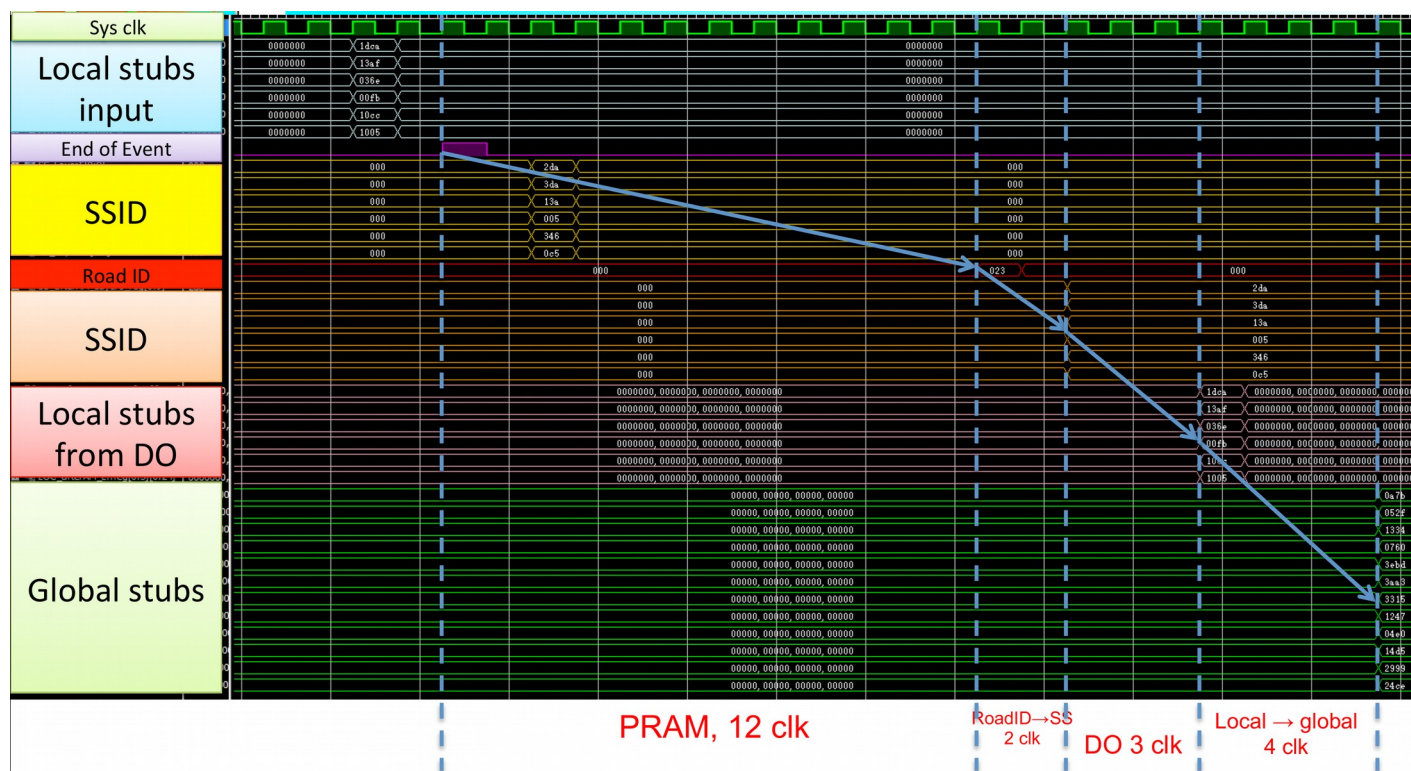
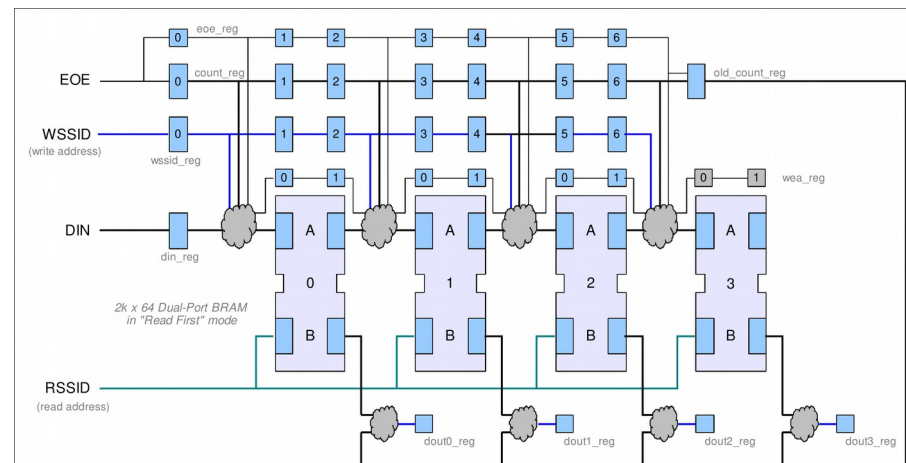
FPGA data flow equivalent in both mezzanines

- Implementations presently differ, optimized for different architectures



A smart database for stub storage

- Determines SS for input stubs
- Stores stubs in RAM at addresses indexed by SSID
- Sends SS's to AM for matching
- Sends stubs from SS's in a match to Fitter
- Optimized for fast stub retrieval



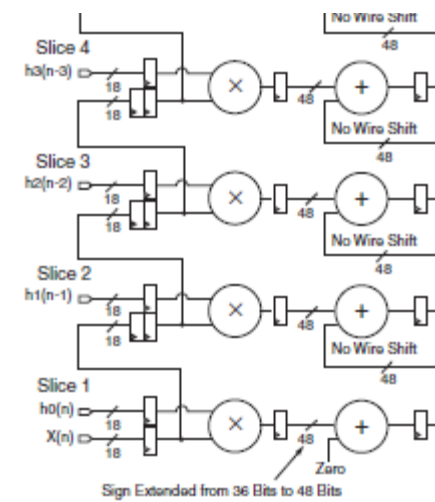
Latency: 21 clks
 (~100 ns @ 250 MHz) for a single muon from stub input to output at DO

Uses 20-30% of KCU040 BRAM

PRM : Track Fitting

Both PRMs use PCA-based, linearized track fit

- Heavy use of DSP48 for multiply/accumulate (MACC)
- Working implementations of both cascading and multicycle MACC
- Single fitter runs up to 500+ MHz



Hardware fit consistent with (emu)simulation

- Focus now on scaling number of fitter instances

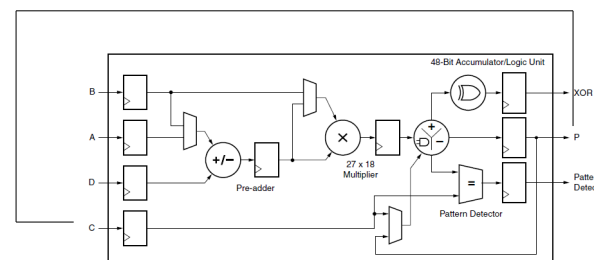

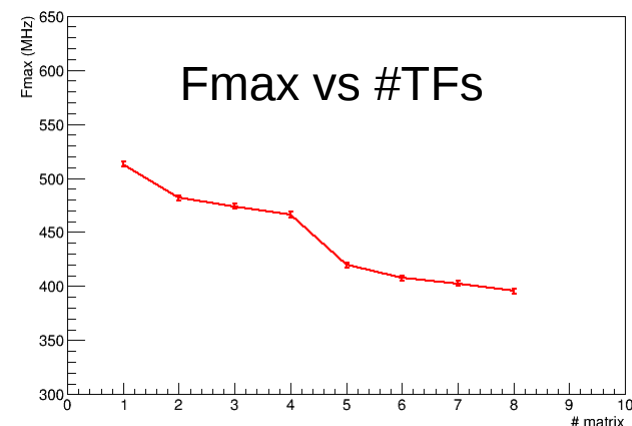




Figure 1-1: Basic DSP48E2 Functionality



8 cascading TFs



Outlook



Q4'16 technical demonstration!

- CMS L1 track finding part of the Phase-2 Tracker Upgrade project
- Successful demonstration crucial for the Tk TDR
- At least one approach must be shown to be feasible ...



Outlook (2)



More on demonstration goals:

- Full slice integration and performance measurements
- Explore all aspects of system operation
 - eg: I/O, formatting, pre-processing, efficiency/rate/resolution vs latency
- Scale to footprint/cost of a full system

Continued R&D ... :

- ATCA board design
- high-speed links
- Evolution of AM technology
- Demonstration experience will guide future development

L1Tk Interface with L1 trigger

- L1Tk ↔ L1 interface specification, design interplay under discussion
- See talk E. Perez on Thursday for discussion of downstream L1

Summary

CMS L1 tracking ...

- Ⓐ Is a key enabler of the CMS HL-LHC physics program
- Ⓑ Involves significant challenges wrt rate, combinatorics & latency
- Ⓒ Has three associated R&D efforts exploring various approaches to solving these problems
 - These efforts are building upon many years of design and development work
 - Now working full steam toward technical demonstration
- Ⓓ Feasibility will be assessed & performance characterized for the 2017 Tracker TDR

Summary

CMS L1 tracking ...

- ✓ Is a key enabler of the CMS HL-LHC physics program
- ✓ Involves significant challenges wrt rate, combinatorics & latency
- ✓ Has three associated R&D efforts exploring various approaches to solving these problems
 - These efforts are building upon many years of design and development work
 - Now working full steam toward technical demonstration
- ✓ Feasibility will be assessed & performance characterized for the 2017 Tracker TDR

Thank You!