Tracking Trigger R&D for high luminosity LHC

-- the view from science and technology

Ted Liu

FNAL

June 1st, 2012

Chicago-Beijing Detector R&D Workshop

University of Chicago Beijing Center

Intro

It has become clear that track based trigger capability will be crucial to the ultimate physics reach at LHC in the future, as luminosity increases.

The current technology using fiber data transfer, FPGAs, custom chips and modern PCs cannot be scaled in a simple manner to accommodate all the tracking trigger demands.

Significant improvements, or breakthroughs, will be needed. In other words, aggressive R&D efforts will be required.

Ultimate physics at LHC in the future:

Go to Higher Energy: needs aggressive R&D in Magnet

Go to Higher Luminosity: needs aggressive Tracking Trigger R&D

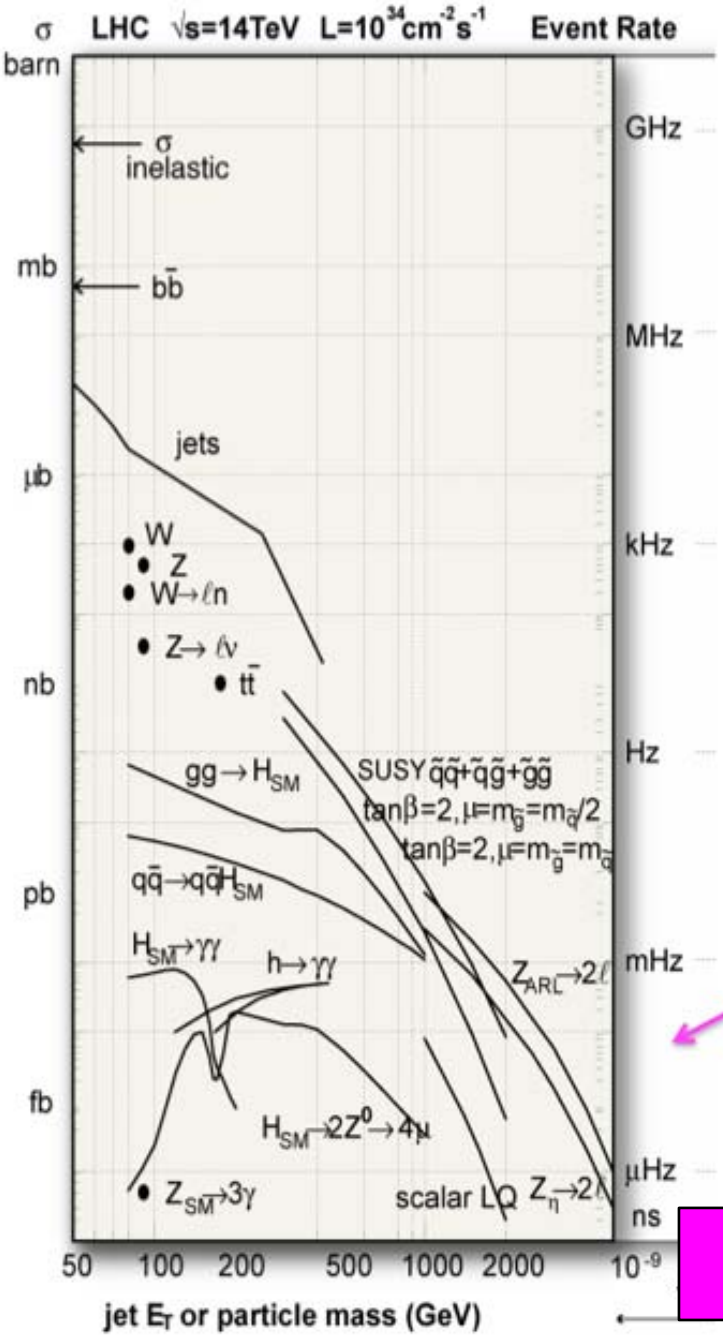

Hans Bethe's advice to a student who could not find a girlfriend for a long time:

"young man, if the cross section is so low, increase the luminosity!"

good advice for students who were interested in charm and beauty (the B Factory era…)

Hunting for new particle with cross section down to ~ 1 fb

100 fb^{-1} per year ~ 10^{-5} $fb^{-1}s^{-1}$ = $10^{34}cm^{-2}s^{-1}$

The need for Higher LHC Luminosity

Detector Luminosity Effects

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

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

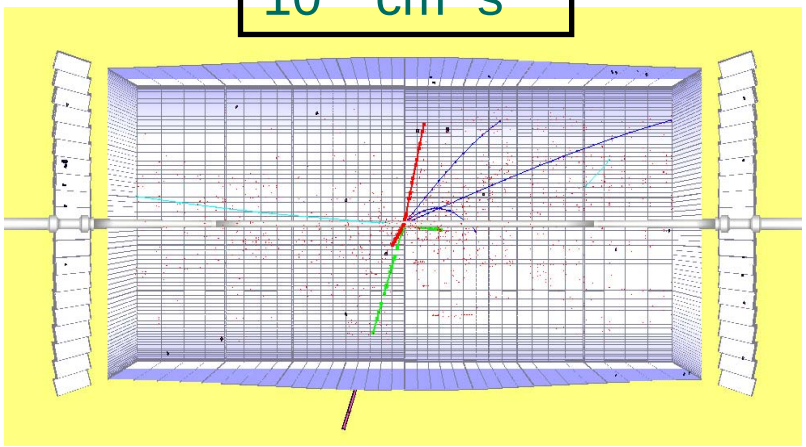

10^{33} $cm^{-2}s^{-1}$

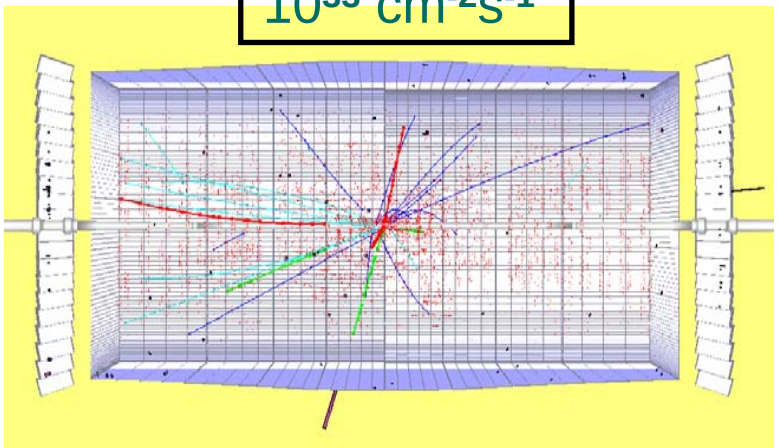

10^{34} $cm^{-2}s^{-1}$

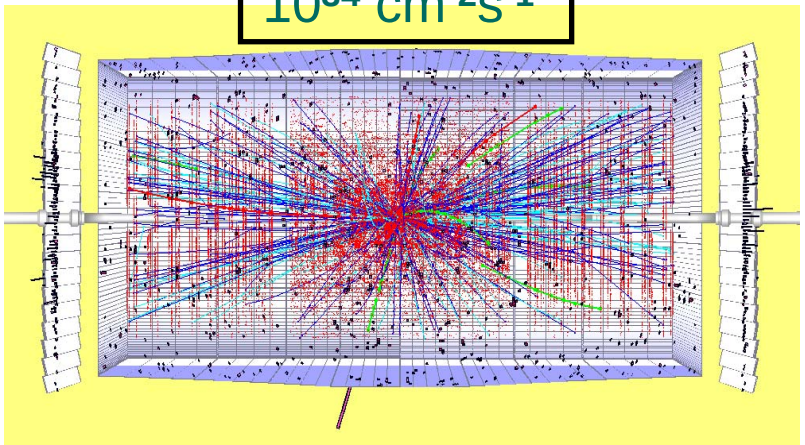

10^{35} $cm^{-2}s^{-1}$

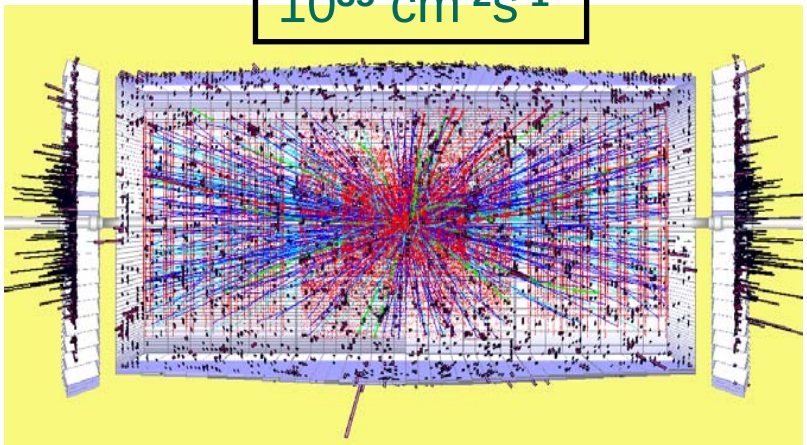

For hadron collider: increasing luminosity could get one into deep trouble …

Expected Pile-up at High Lumi LHC in ATLAS at 10^{35}

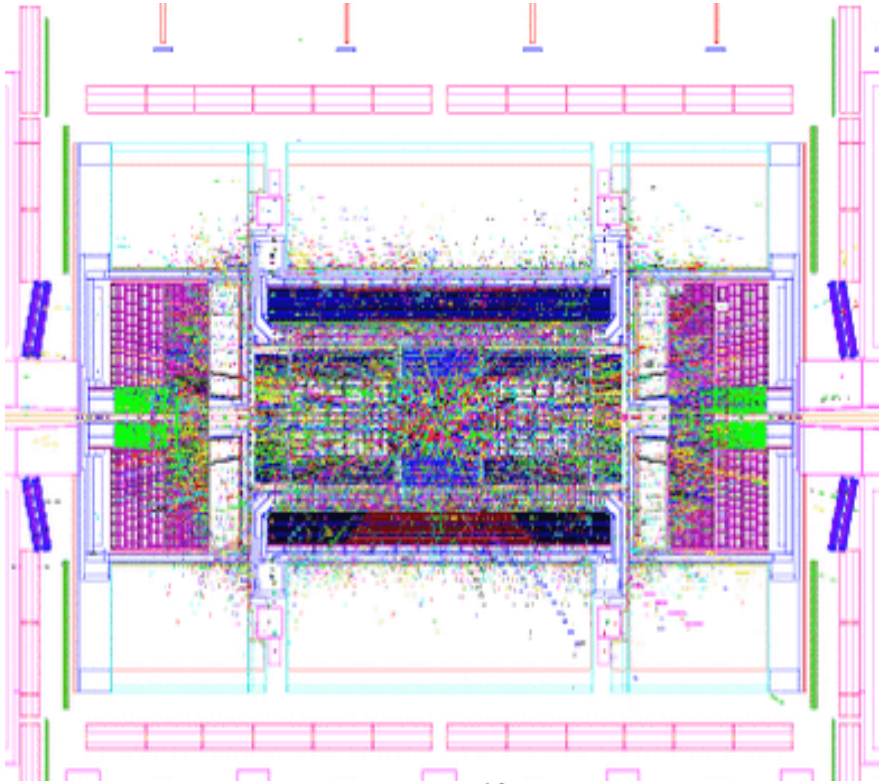

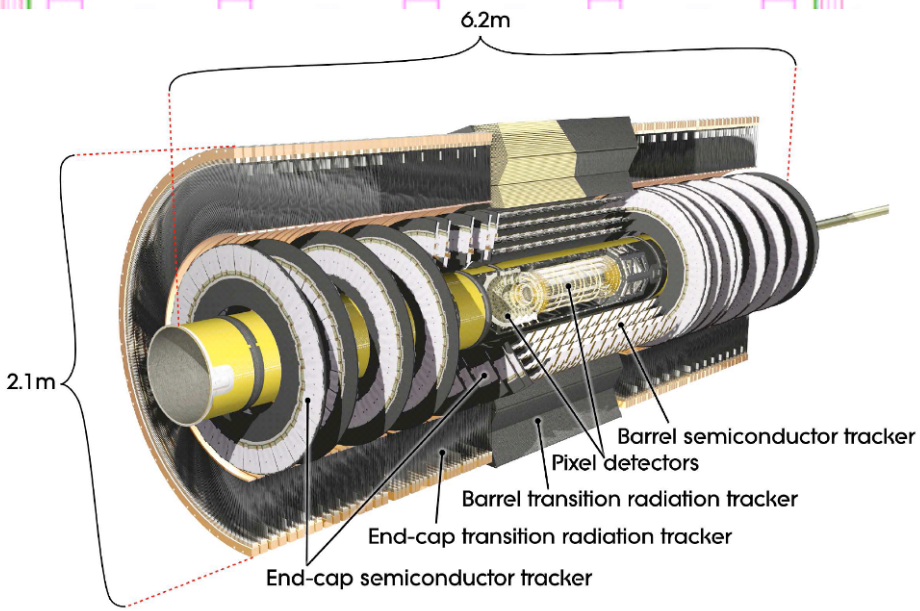

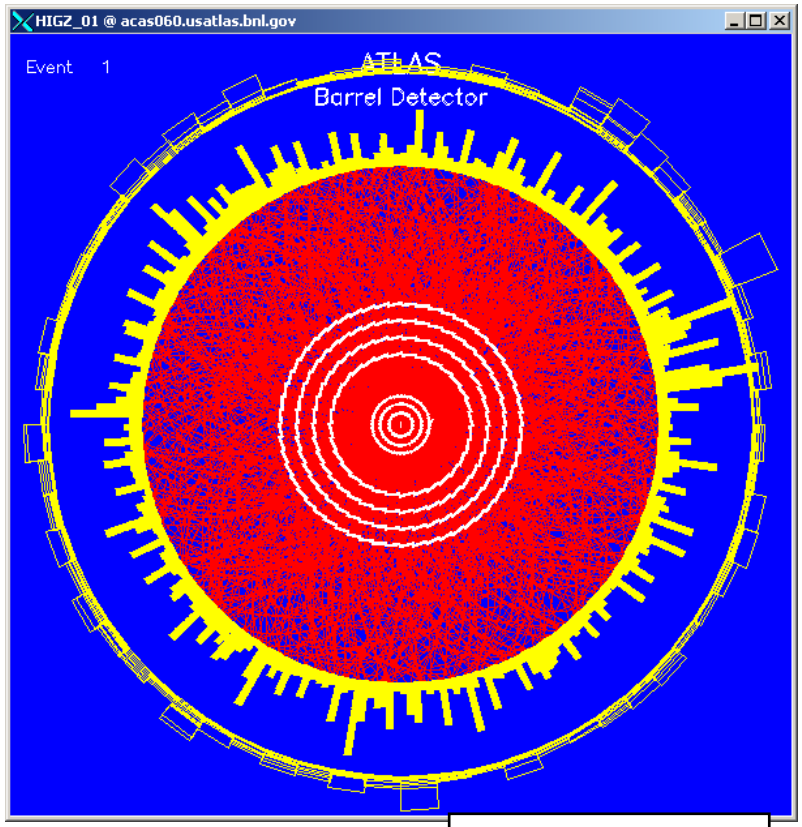

$N_{ch}(|y|\leq 0.5)$

Fermilab

The effect of pile-up on silicon detector occupancy

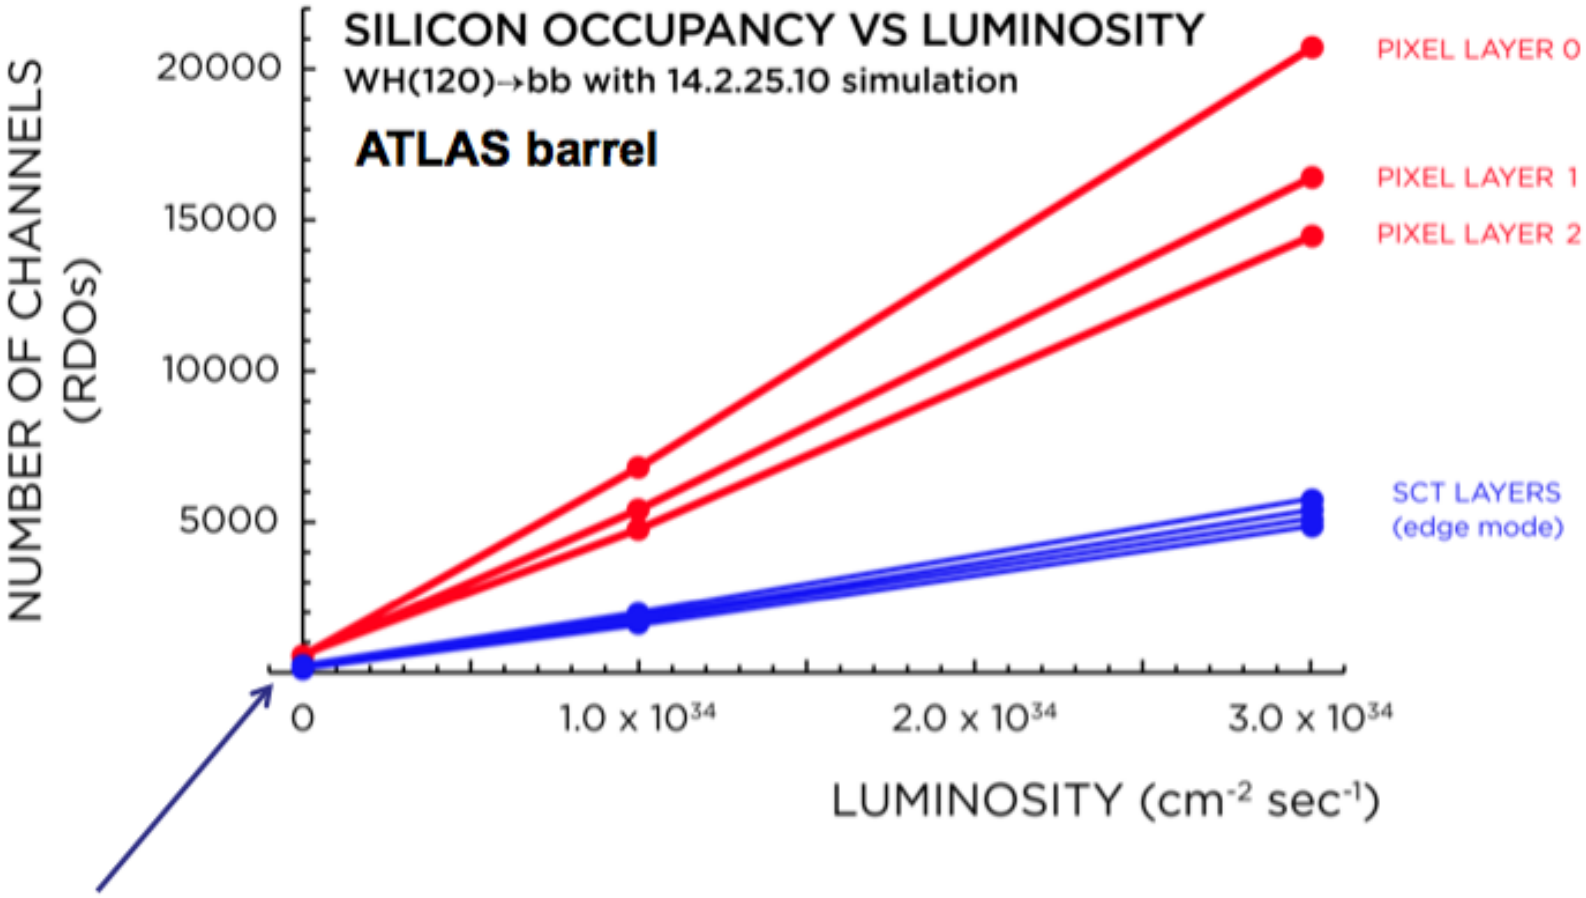

***y* intercept: all of the tracks from the primary interaction!**

ATLAS and CMS only have calorimeter and muon trigger at Level 1

Pattern recognition much easier on calo & muon:

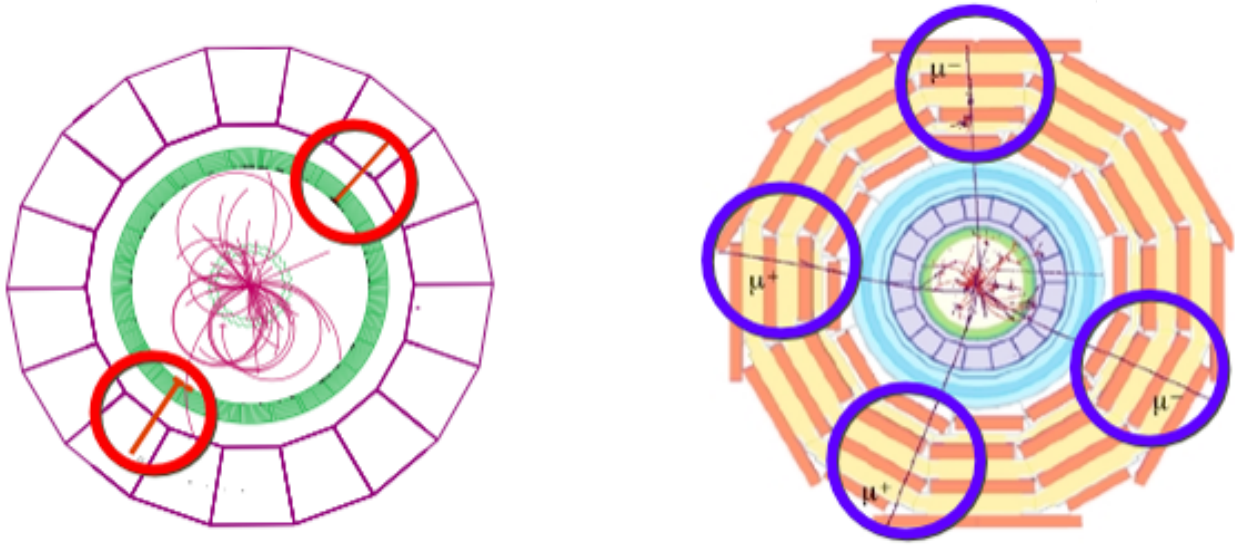

Complexity handled in software on CPUs

at high level trigger

The approach works well at low luminosity only …

Collisions (p-p) at LHC

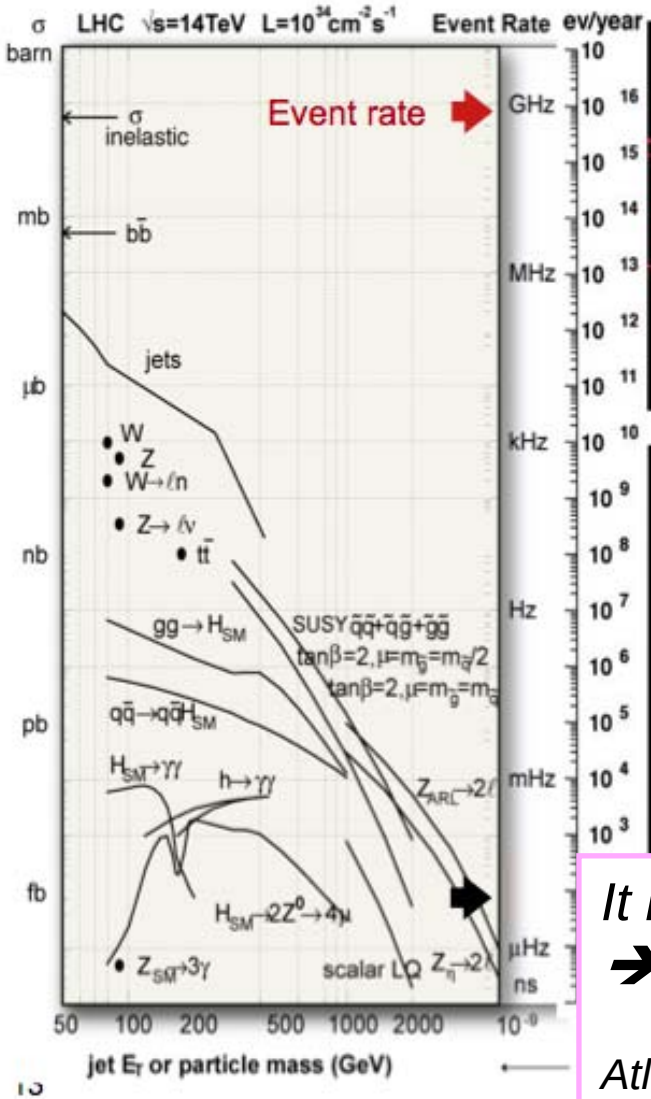

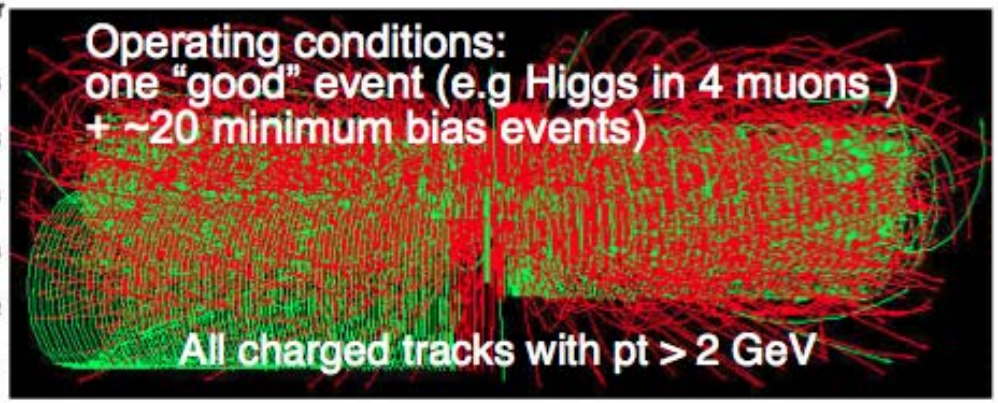

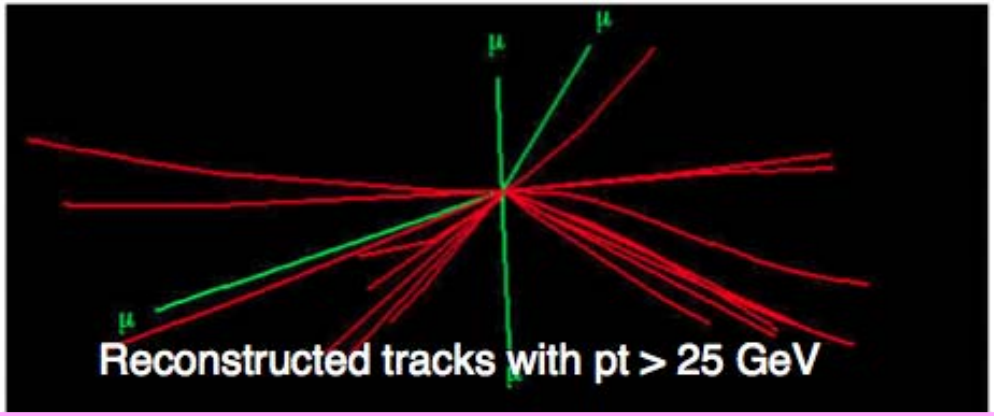

It is crucial to have tracking at early trigger stages ➔ enormous challenges !

Atlas FTK is the first step in this direction at LHC, and uses all silicon tracker info at L2. (UC/FNAL/ANL collaboration on FTK, with others).

Tracking Trigger Issues at hadron colliders (L1 & L2/HLT)

Detector design for triggering

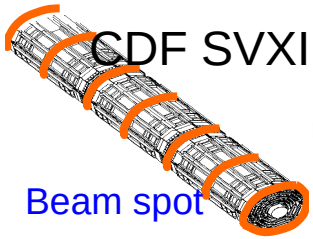

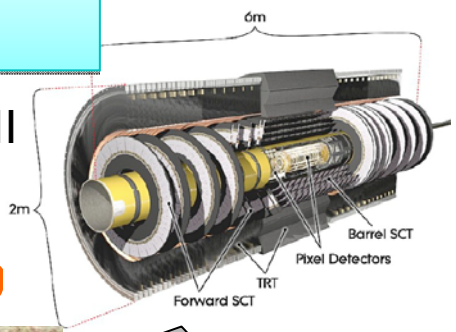

Data transfer

Data formatting

Pattern Recognition

Associative Memory
Others …

Track Fitting

FPGA vs CPU vs GPU

HLT

- *Three selected R&D topics at FNAL*

  (1) Data Formatter using ATCA technology
  (2) 3D AM R&D: VIPRAM
  = "Vertically Integrated Pattern Recognition Associative Memory"
  (3) GPU study for low latency tracking trigger applications

A sense of scale:

Atlas Silicon Tracker vs CDF SVX II

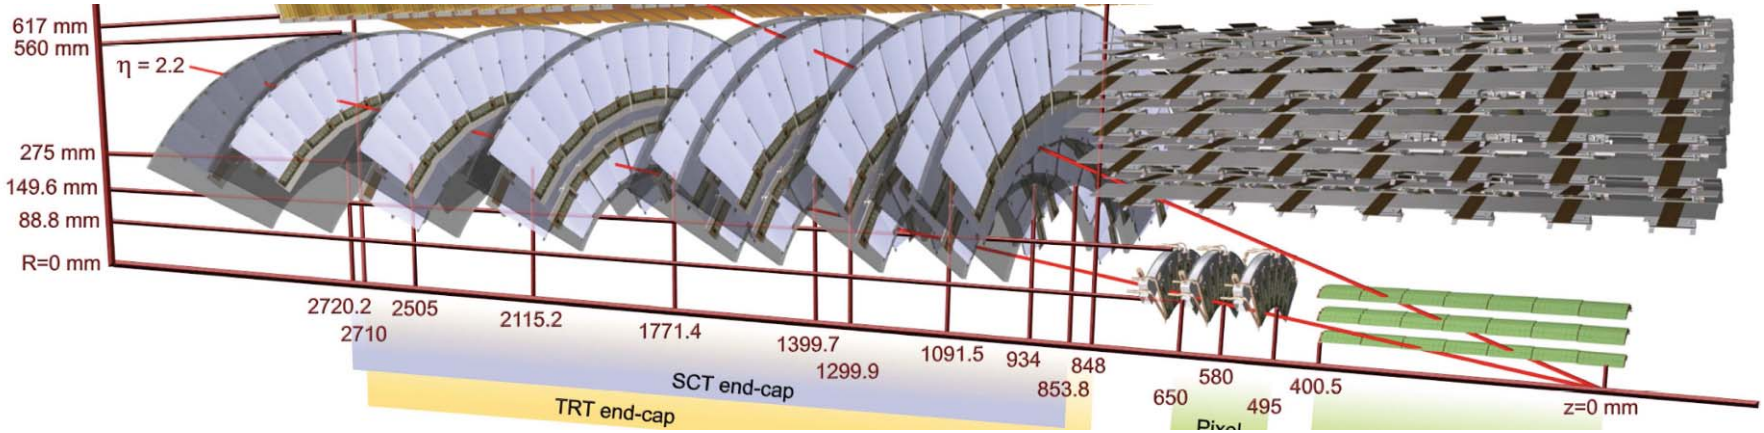

R-phi view of Barrel region:

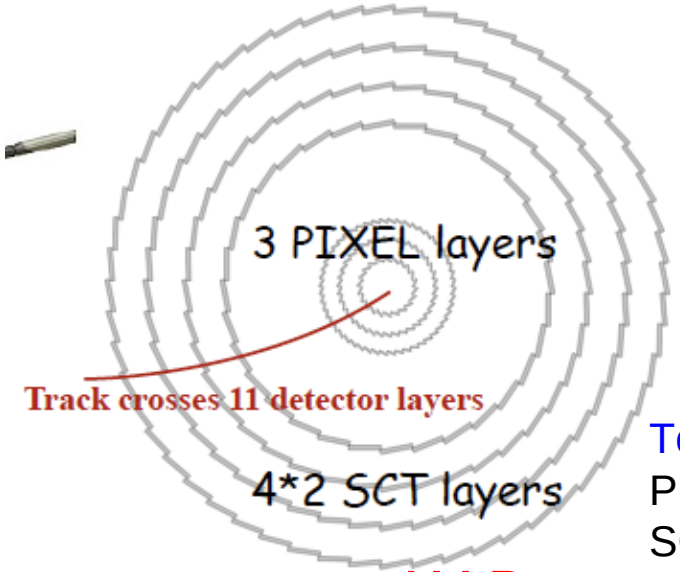

Other relevant aspects:
Collision energy/rate
Pileups/occupancy
Symmetrical design or not
Materials
Cabling map
…

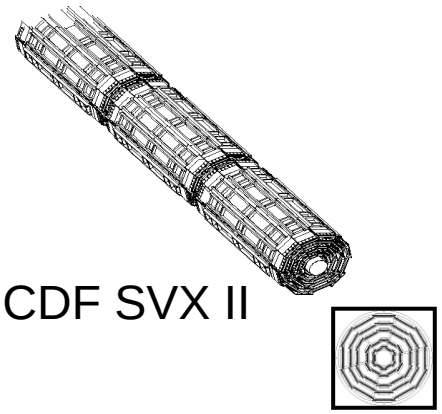

CDF SVX II

Total # of readout channels:
PIXELS: 80 millions
SCT: 6 millions

Channels used for SVT:
~ 0.2 millions

AM Patterns needed: > 1 Billion

patterns 1st used: ~400K

Fermilab

CDF original SVT system had ~400K patterns total … 128 patterns per AMchip -- commissioned around ~2001.

Question: Can we put the entire SVT system into one chip? This is what it takes to handle high luminosity LHC tracking trigger pattern recognition

~400K patterns per chip ➔ few x 1000 chips ➔ Billion patterns…

Data Formatting Challenges

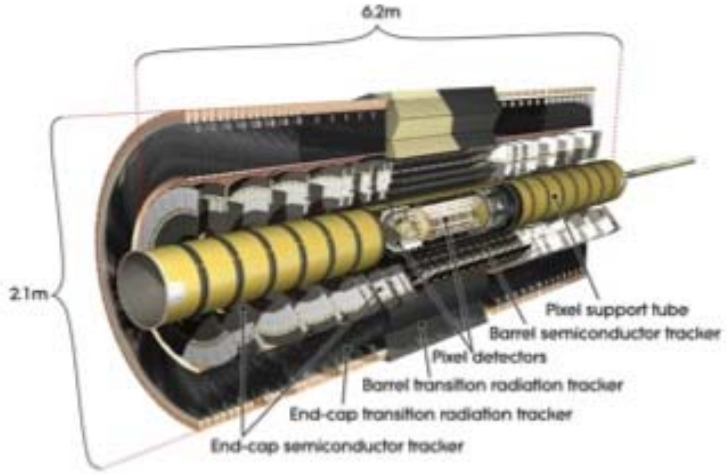

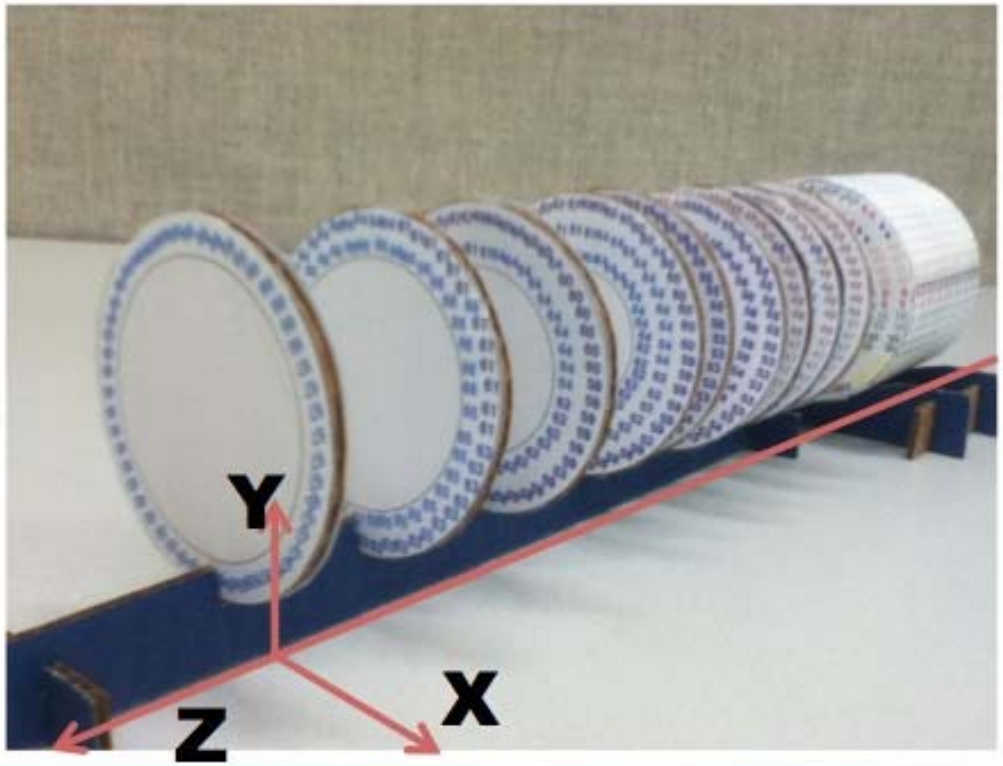

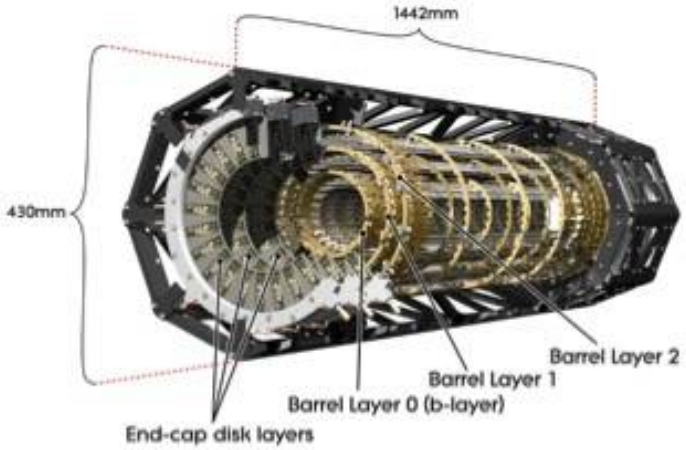

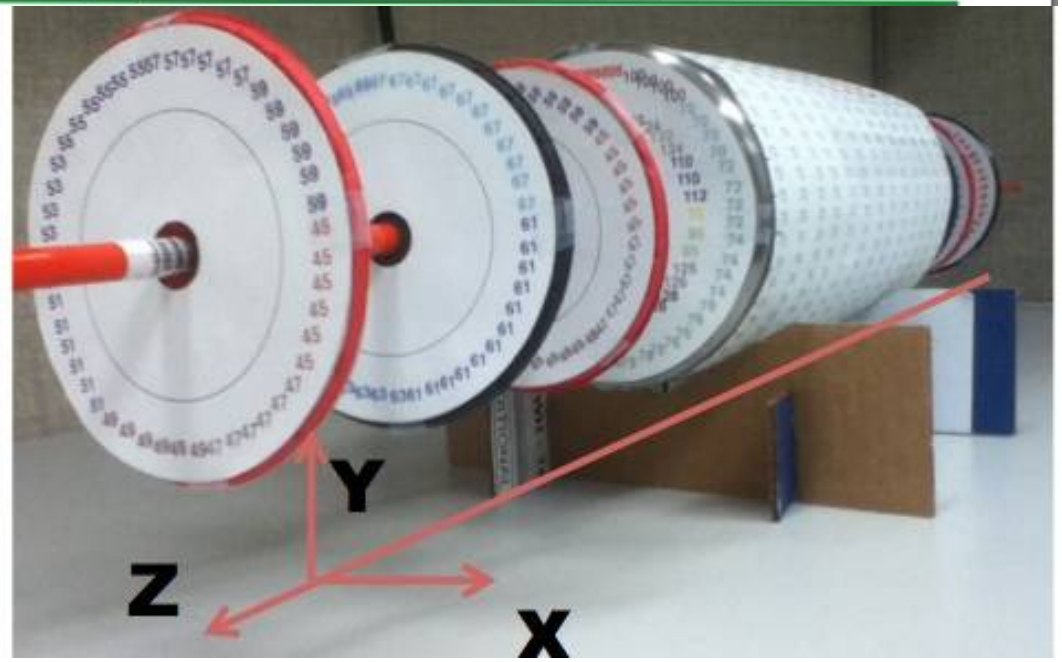

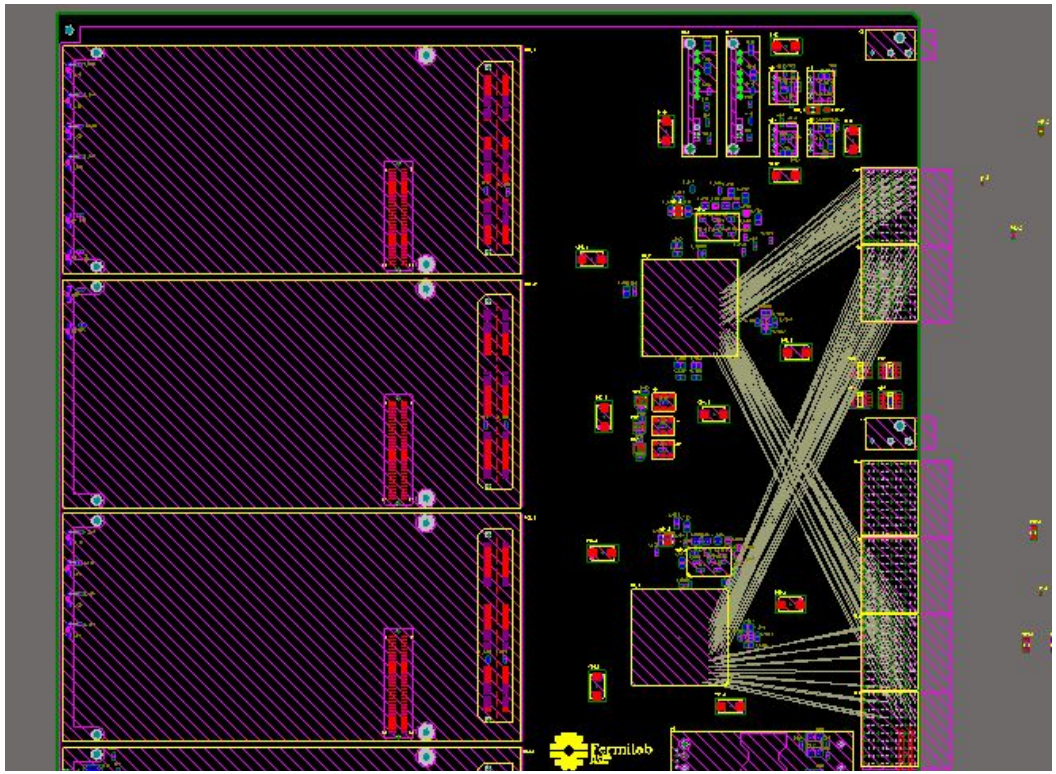

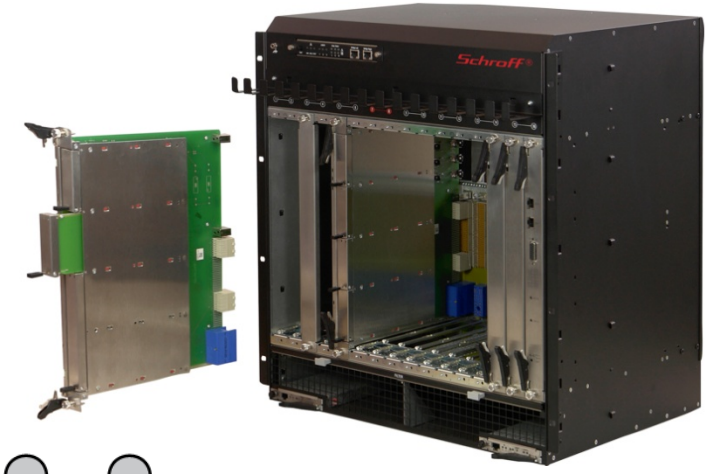

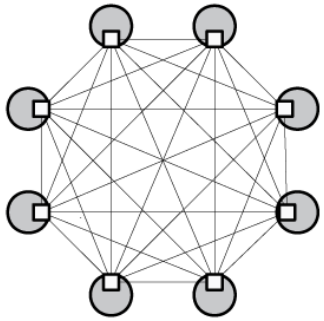

Data Formatter ATCA board (Pulsar-IIa) design at Fermilab with
full-mesh backplane for data sharing (one FPGA per trigger tower, 64 trigger towers total)

16 DF boards

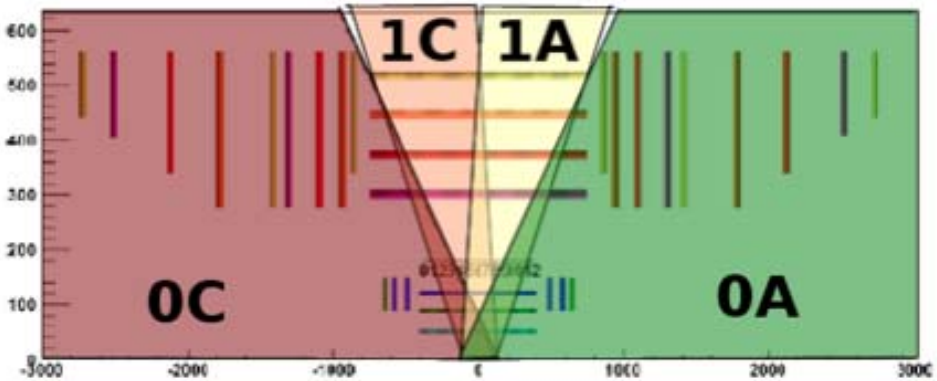

16 DF boards

Figure 2: The Four FTK η regions. Note the significant overlap in the high occupancy central barrel regions.

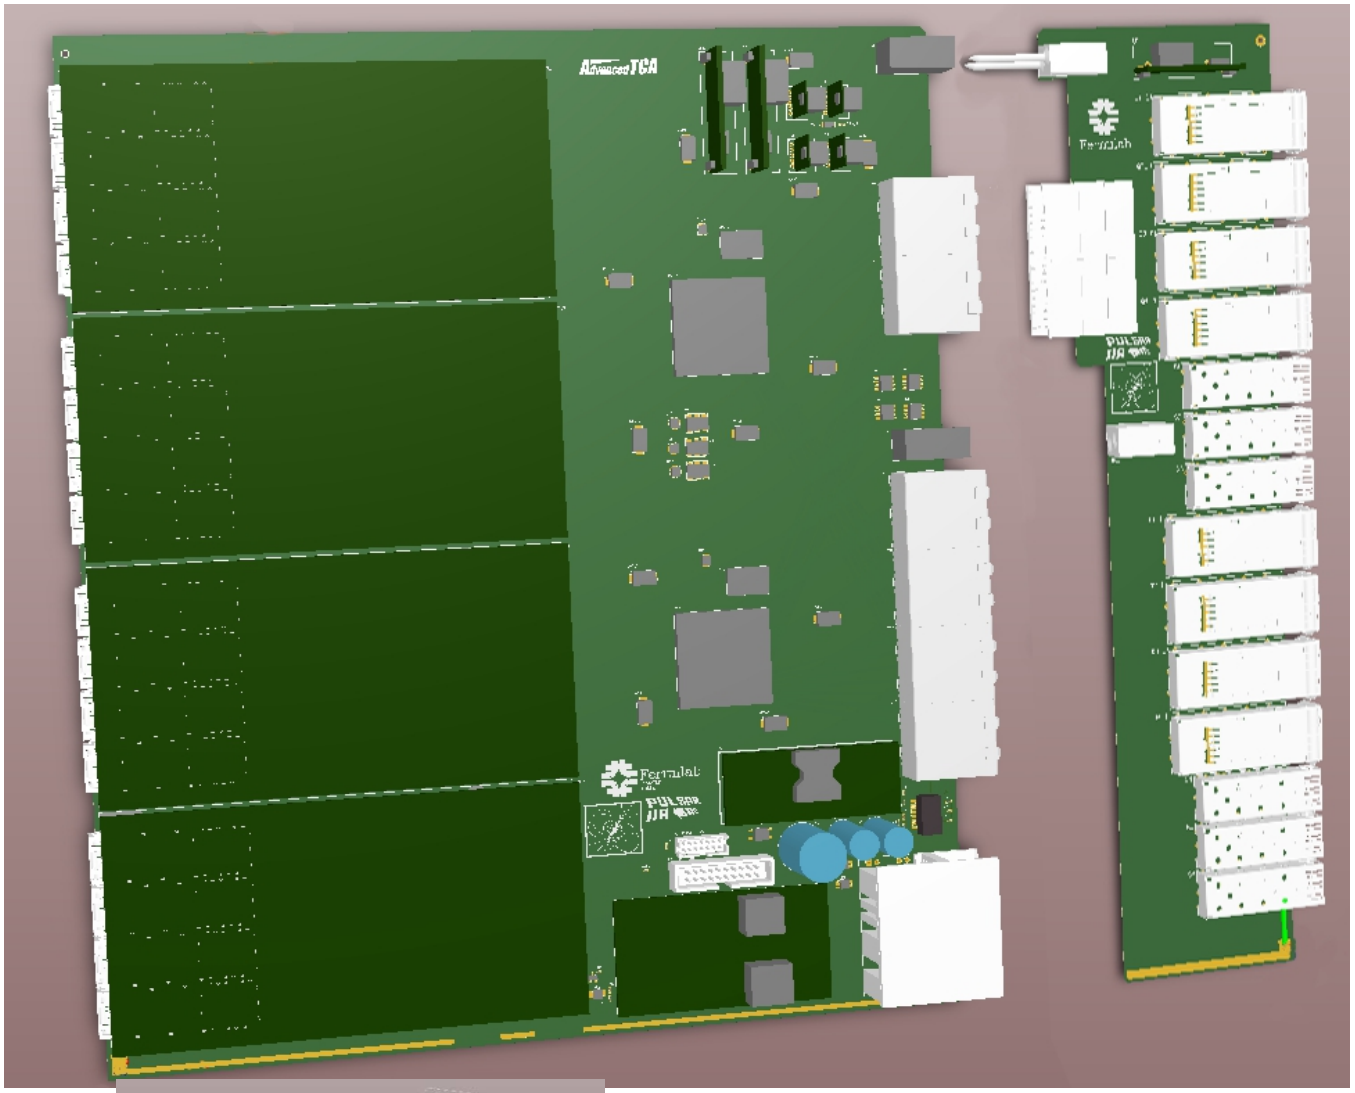

Pulsar-IIa

I/O bandwidth approaching:

~ 1 Tbps

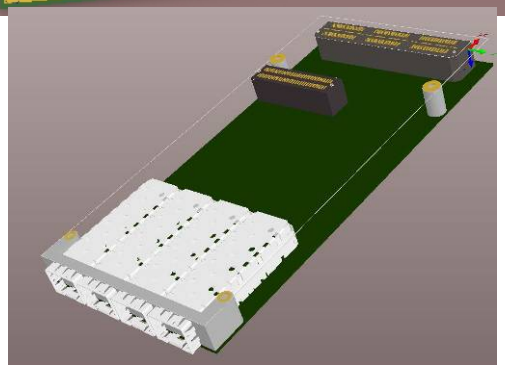

Challenging Implementation

Extensive bandwidth requirements study done using real beam data

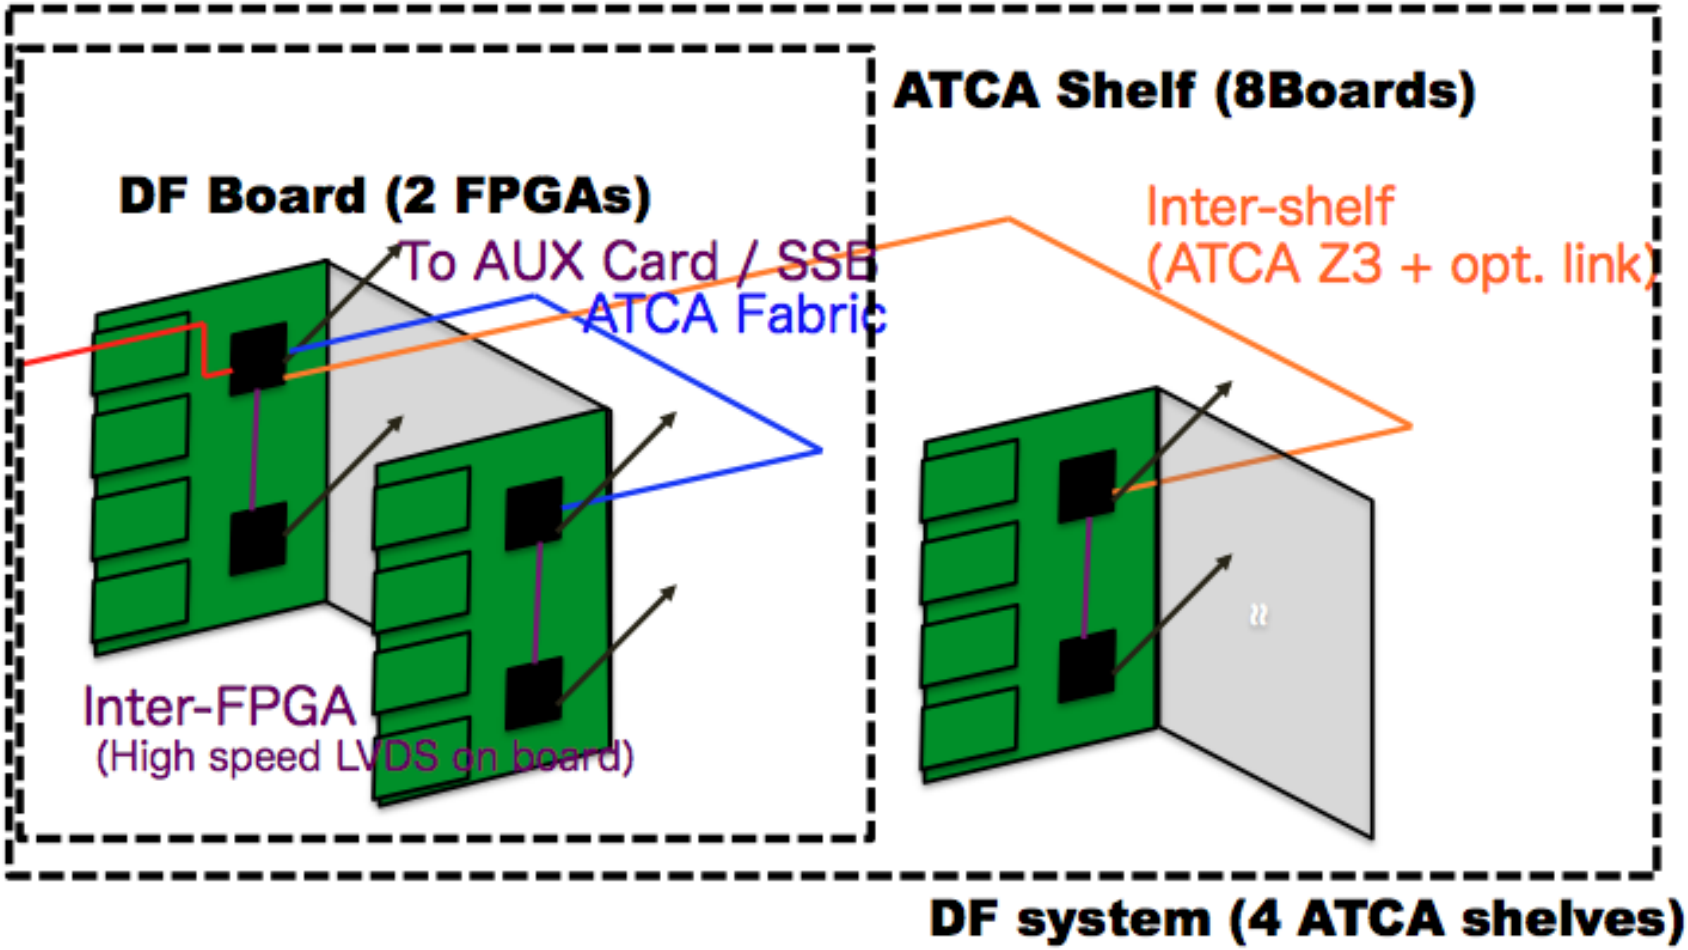

Tracking Trigger Issues at hadron colliders

Detector design for triggering

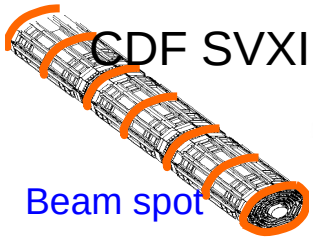

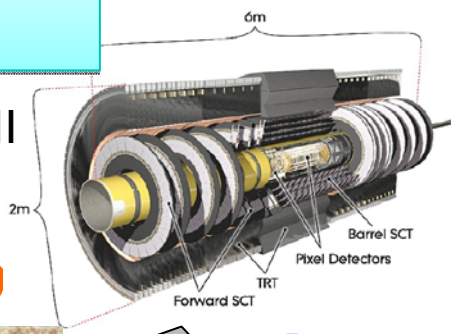

Data transfer

Data formatting

Pulsar-IIa ATCA design

Pattern Recognition

Associative Memory approach

Others …

Track Fitting

FPGA vs CPU vs GPU

HLT

- *New ideas/approaches in four areas*
- *Three of them fully developed*

(1) Data Formatter using ATCA technology

(2) 3D AM R&D: VIPRAM

= "Vertically Integrated Pattern Recognition Associative Memory"

(3) GPU study for low latency trigger applications

CDF Silicon Vertex Trigger (SVT) for RunII

Pattern Matching using Associative memory
(M.Dell' Orso and L. Ristori: initial idea in 1985)

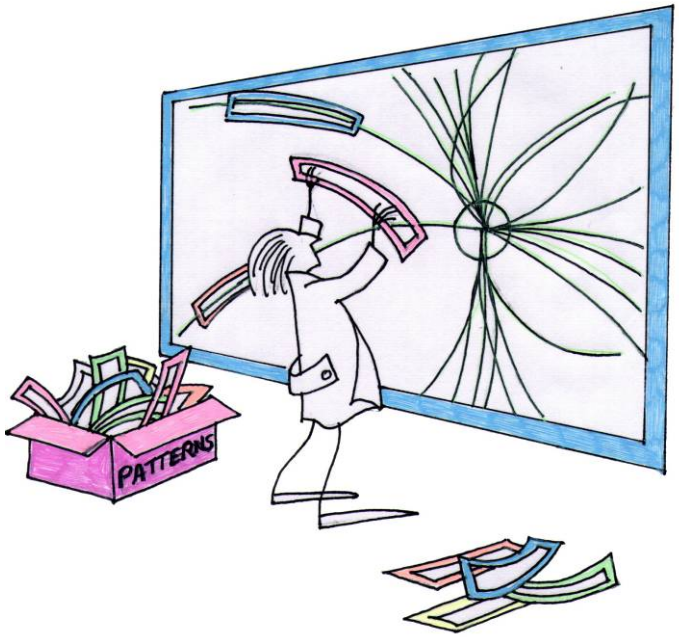

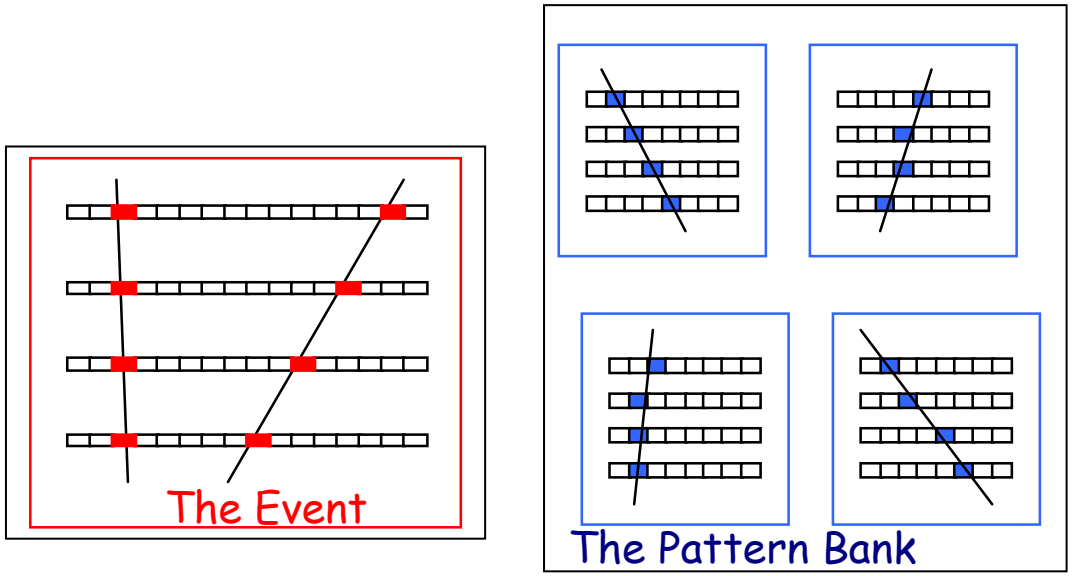

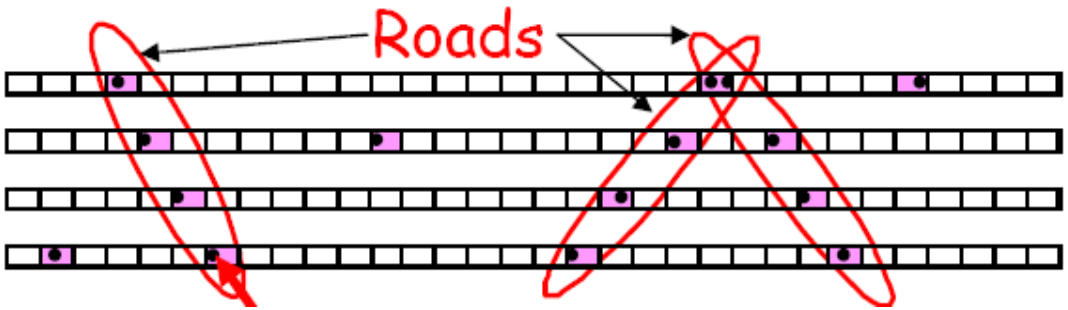

Finds low resolution track candidates first (roads)

Then perform track fitting downstream

Tracking in 2 steps

Pattern recognition and track fitting done separately and pipelined

1. Find low resolution track candidates called "roads".

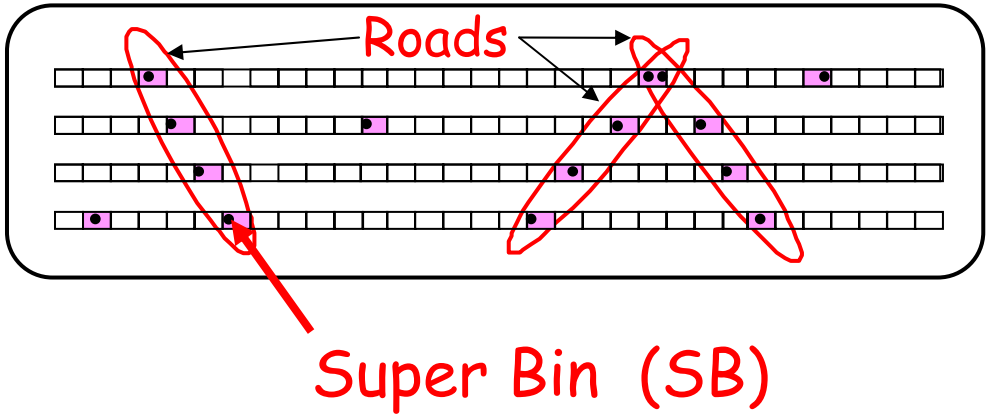

2. Then fit tracks inside roads.

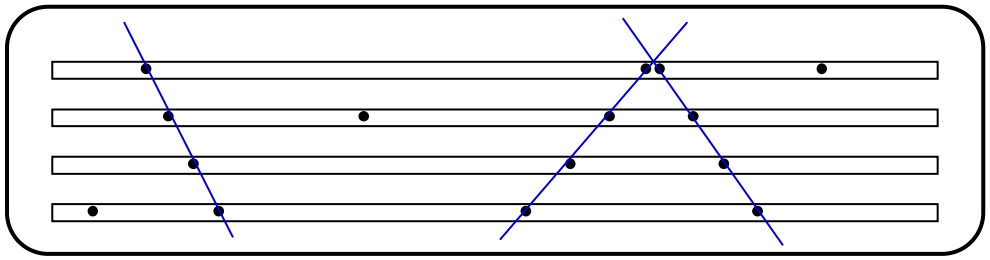

A very successful approach at CDF for RunII, based on Associative Memory, which in turn, based on CAM

Original CDF SVT system

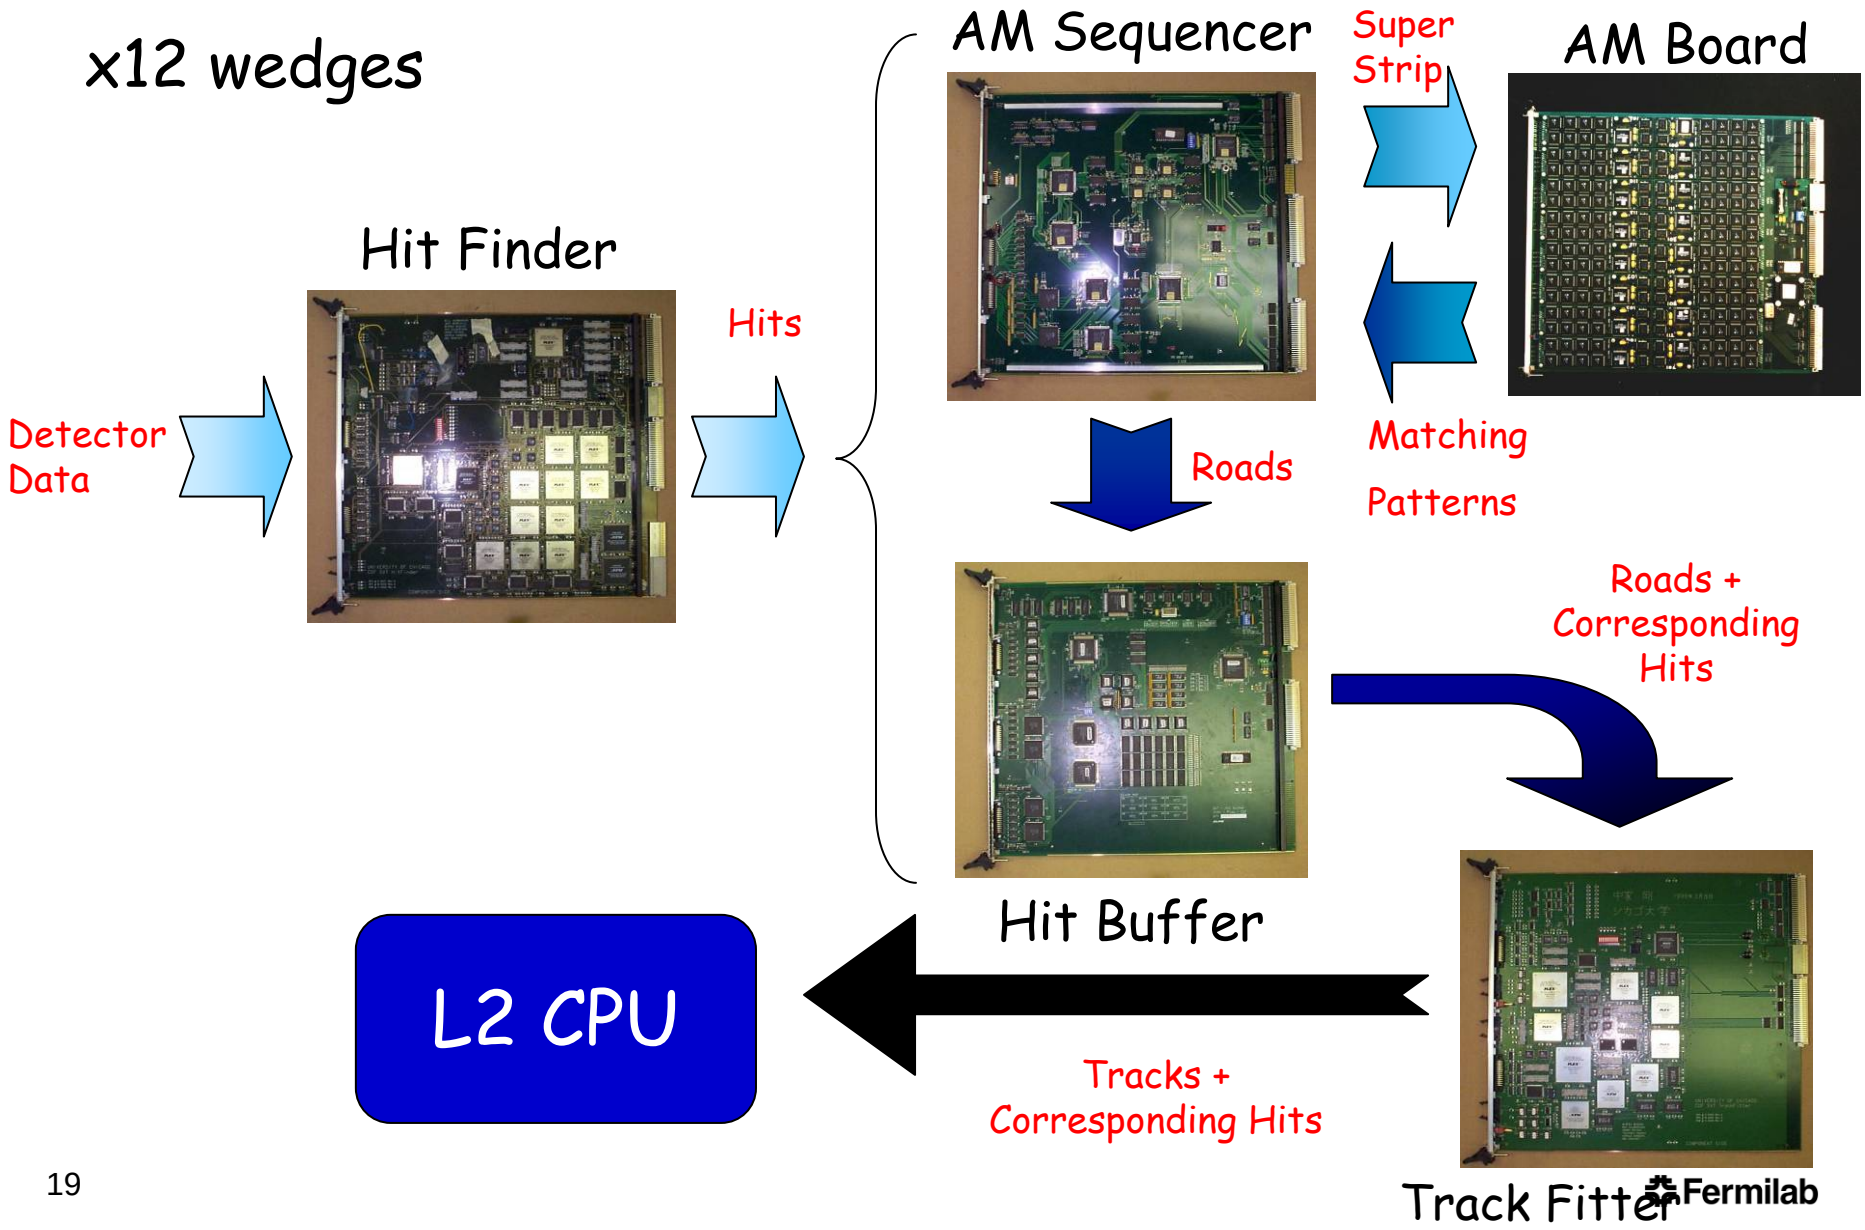

Fermilab

Content Addressable Memory (CAM)

CAM: inverse of RAM

- user supplies a data word and it searches its *entire* memory *in a single operation* to see if that data word is stored anywhere in it

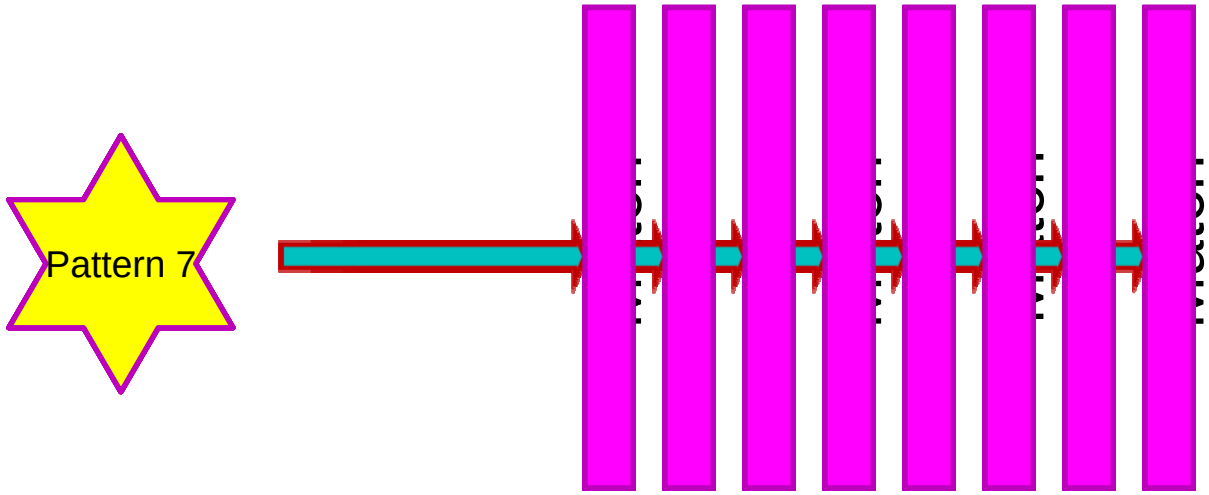

- One pattern at a time
- There is no memory of previous matches

How PRAM works

Pattern Recognition Associative Memory (PRAM)

- *Pattern recognition finishes as soon as all hits arrive*
- *Potential candidate for L1 application*

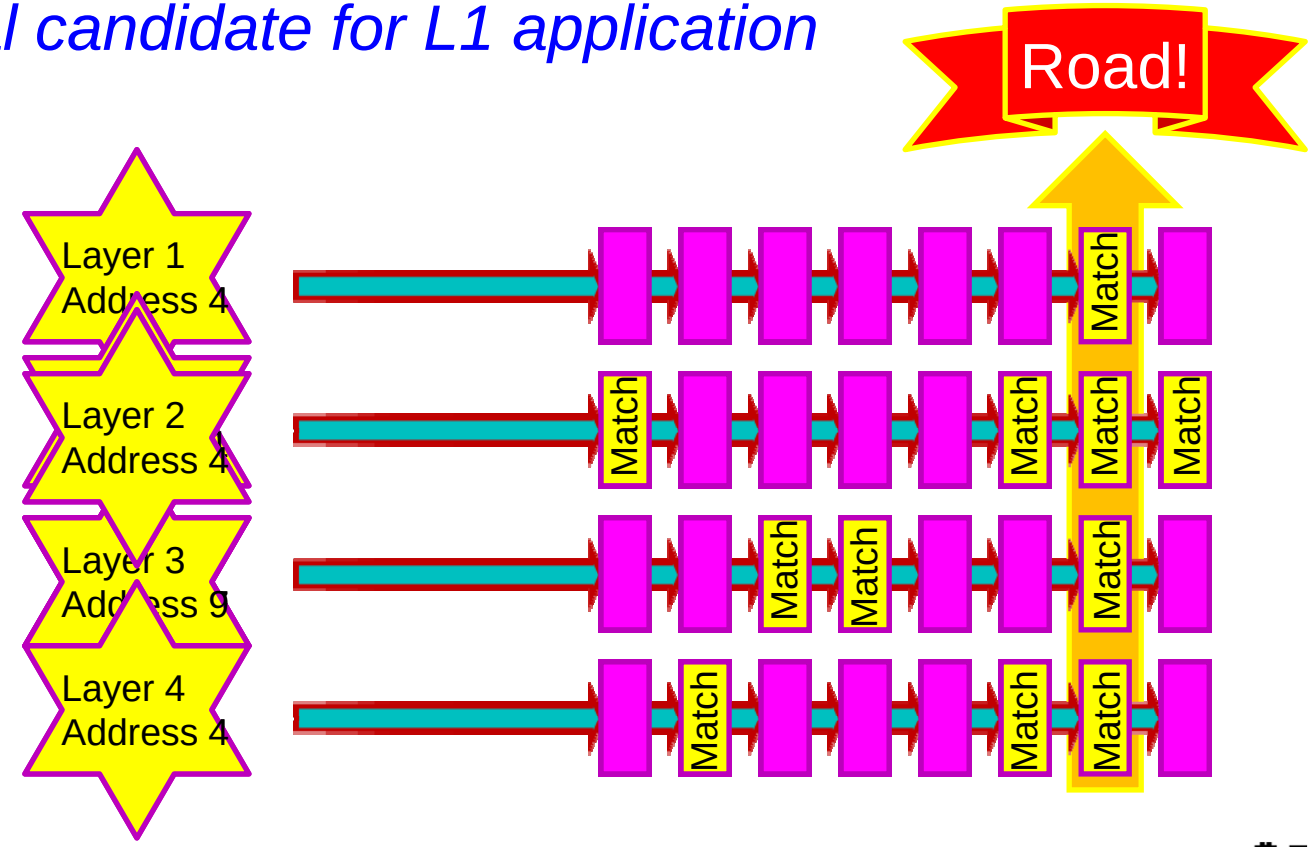

Anatomy of a PRAM
(Pattern Recognition Associative Memory)

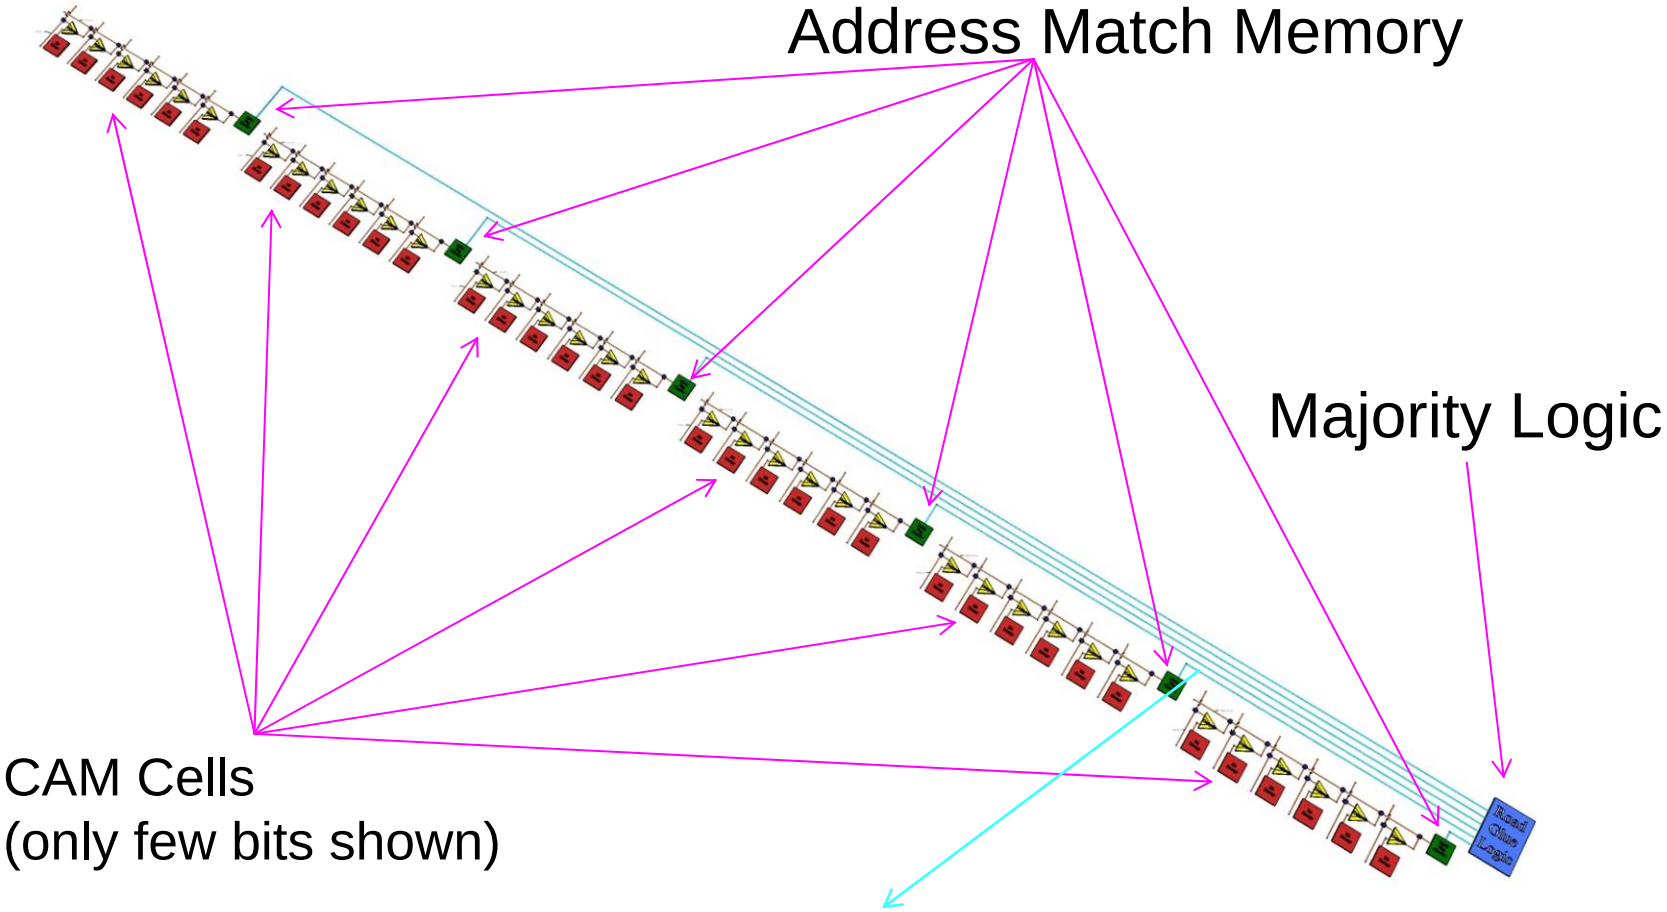

Trace Length -> Capacitance -> Power Consumption or Reduced Speed

More detector layers, or more bits involved, design more spread out in 2D

→ less pattern density, higher power consumption …

Comments on Associative Memory

Based on *CAM cells to match and majority logic to associate* hits in different detector layers to a set of pre-determined hit patterns

Critical figures of merit for an AM based system:

(higher) pattern density & speed and (lower) power density

- However, at chip level, more detector layers means more CAM cells are needed for a given pattern, the layout are more spread out in two dimensions (for a given technology node) resulting in decreasing pattern density and increasing driving load capacitance or power consumption, which in turn reduces the maximal speed of operation
- Performance limited by Moore's Law

This is the main limitation of an otherwise very powerful and proven approach for its future applications within and beyond HEP.

The Challenge of future AM design

Increase the patterns density by 2 orders of magnitude;

and

increase the speed by a factor of >~ 3,

while

keeping the power consumption more or less the same

Much higher Patten Density & higher Speed
Yet much less Power Density
almost *too good to be true*

New idea: could go to "extra dimension" to achieve this
➔ *generic R&D effort at Fermilab*

From 2D to 3D

VIPRAM (Vertically Integrated Pattern Recognition Associative Memory)

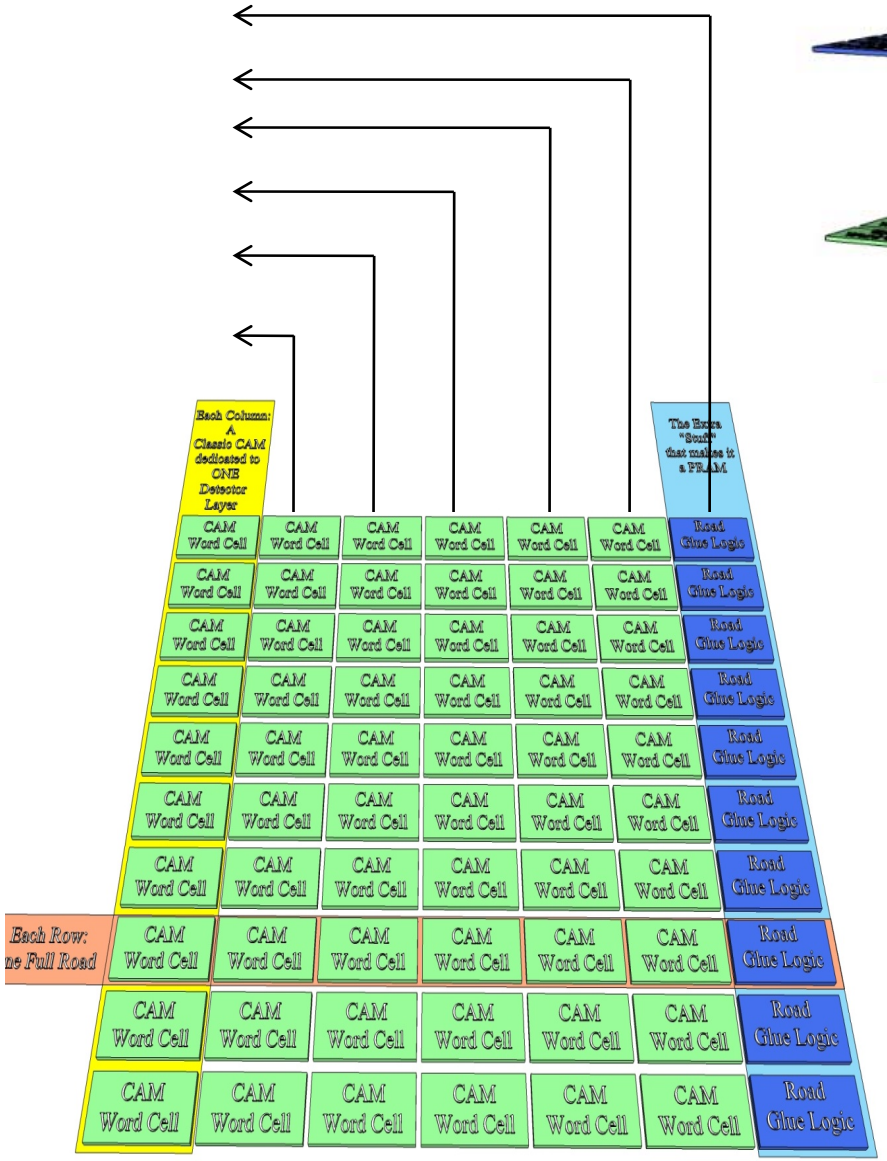

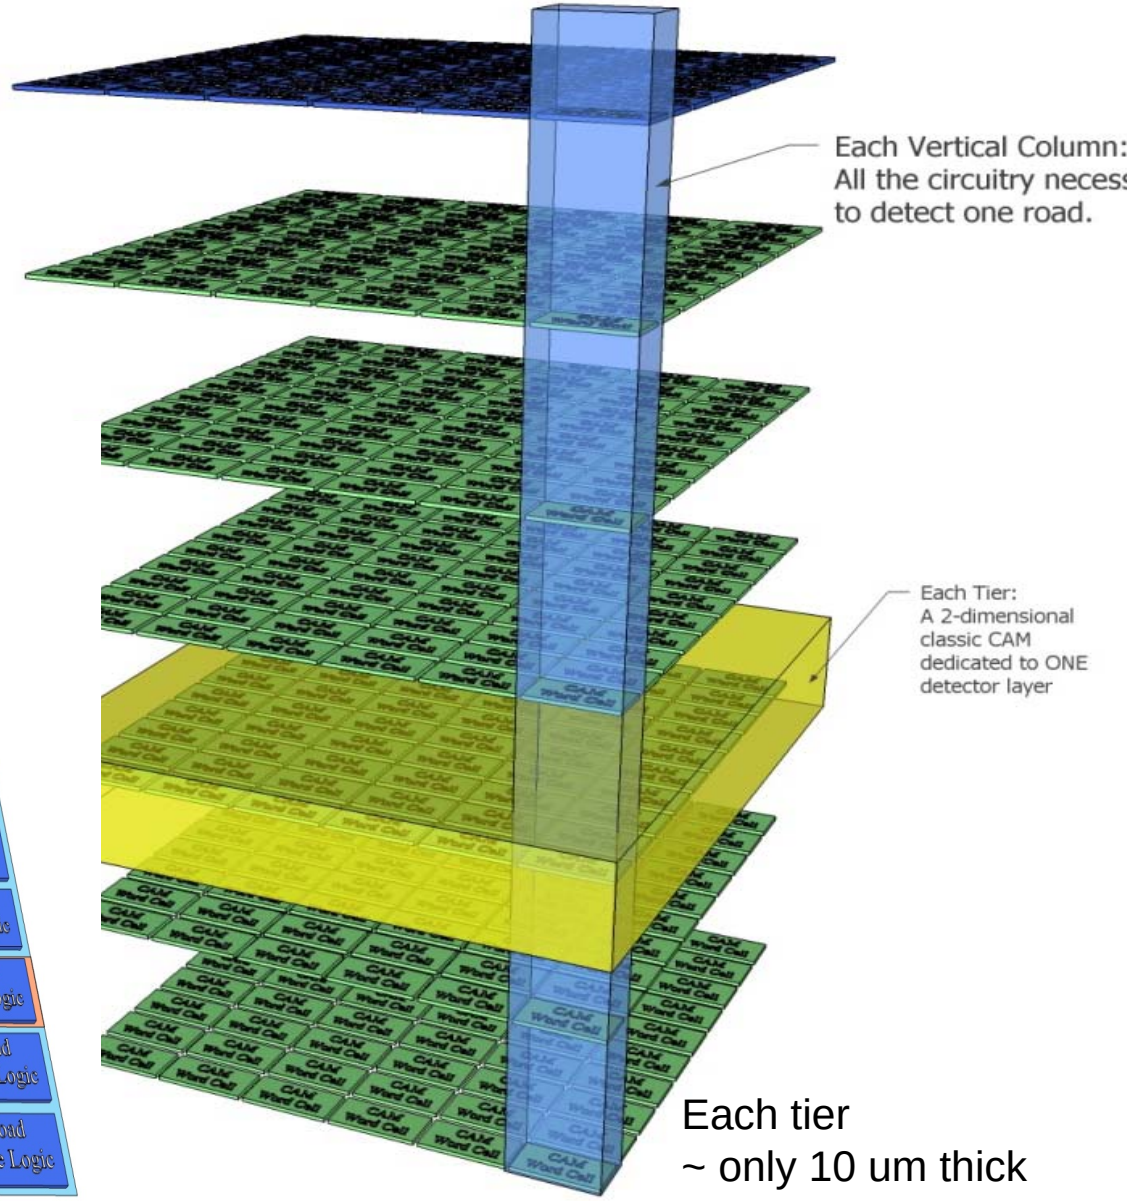

Each tier
~ only 10 um thick

VIPRAM

(Vertically Integrated Pattern Recognition Associative Memory)

Pattern recognition for tracking is naturally a task in 3D

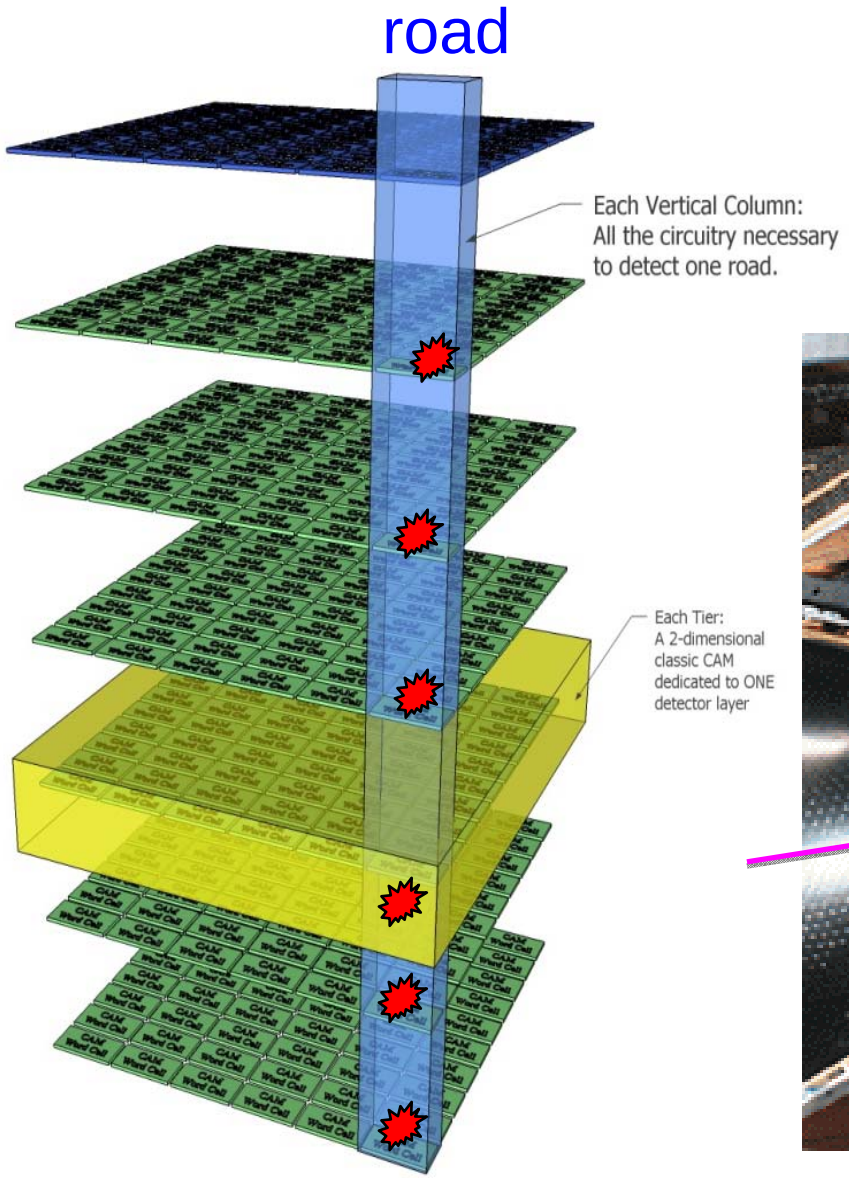

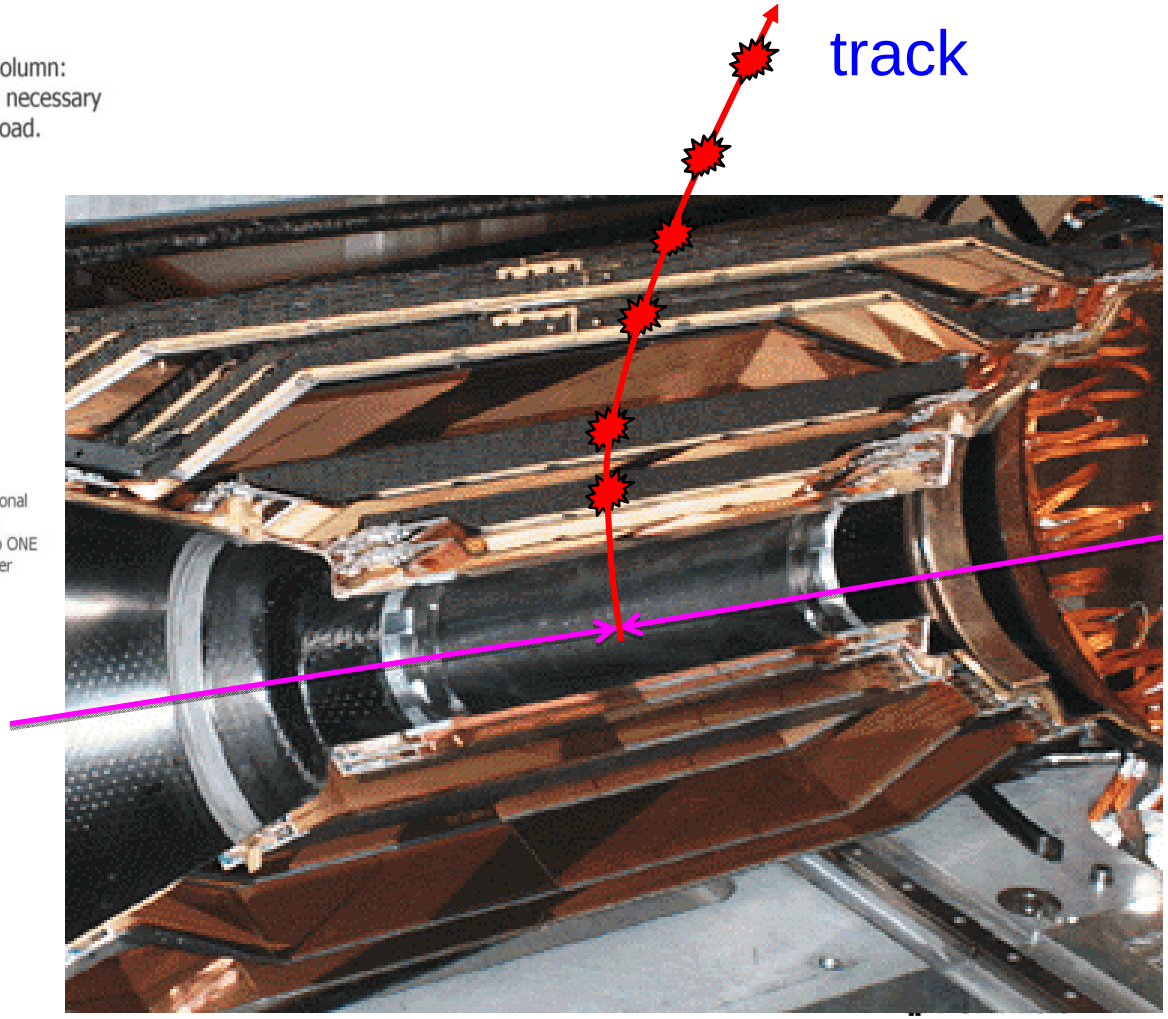

Side view

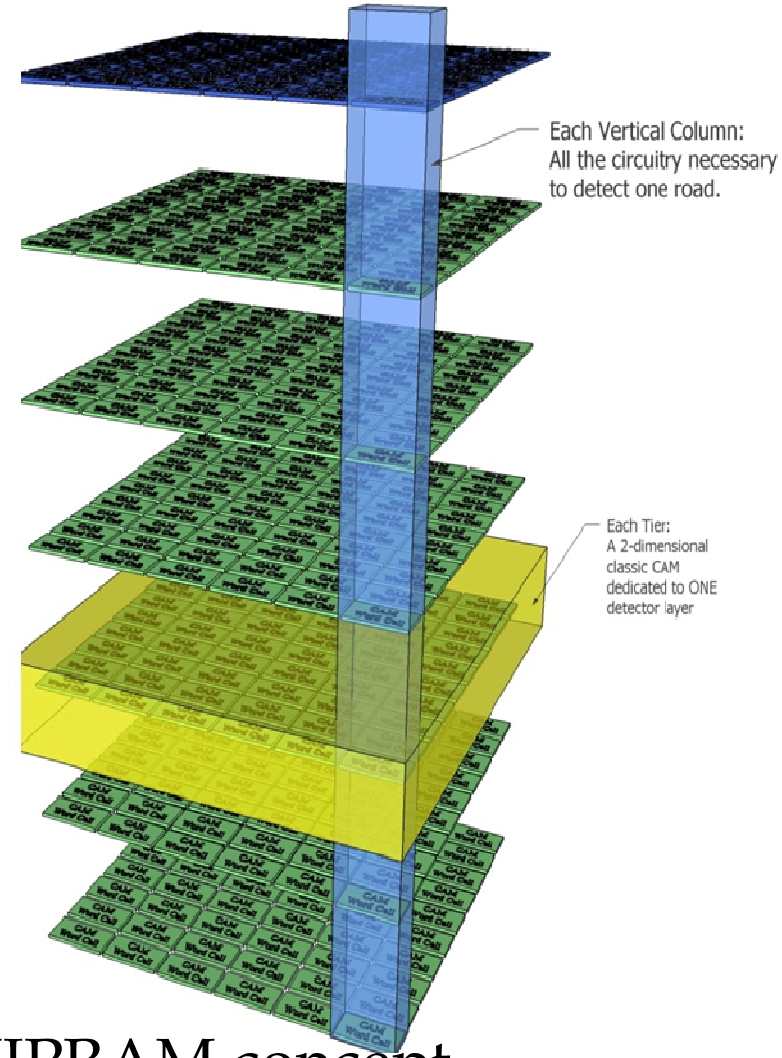

Top view

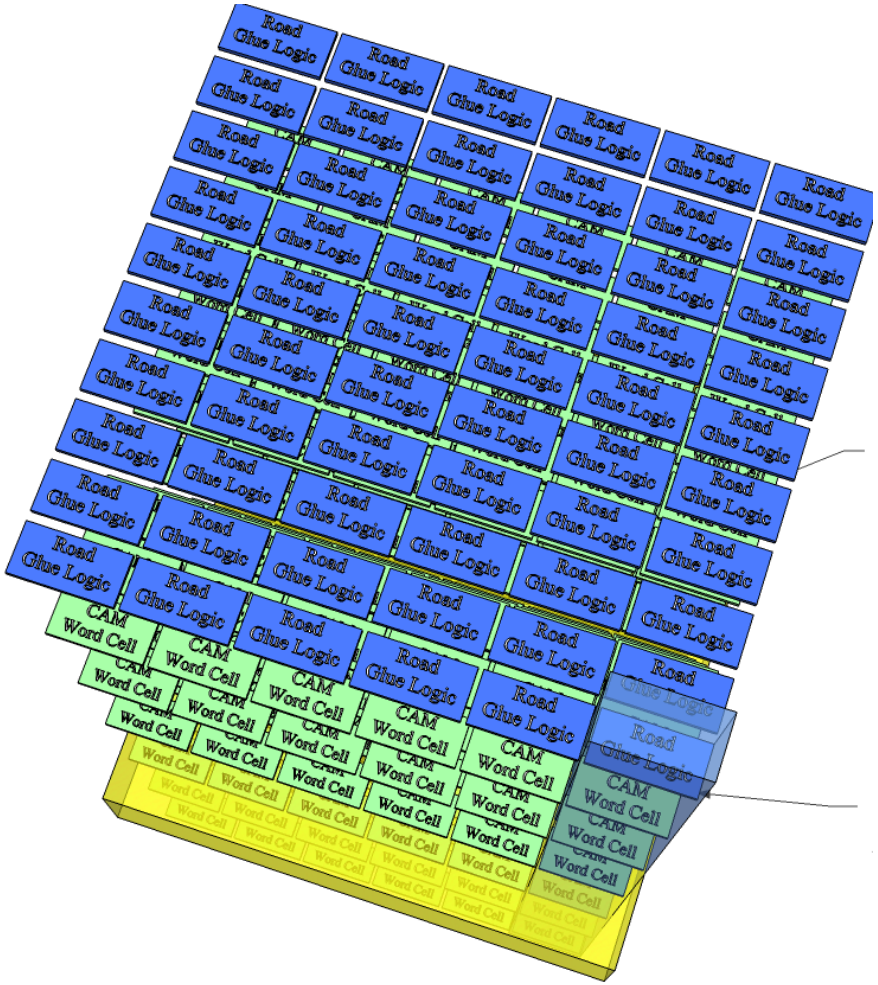

VIPRAM concept

* Even with 130 nm, could reach ~200K patterns/cm^2

Fermilab

VIPRAM is almost an ideal case for the application of 3D vertical integration technology

A VIPRAM cell can process n layers of a road pattern in about the size of just one CAM cell (*pattern density increased by ~ n*)

Directly shortens the longest of the driving lines in the pattern recognition cell (address match lines).

As these lines repeat throughout the chip, expect a significant impact on performance (*reduced power density or higher speed*)

Makes the layout of the CAM cells, Majority Logic cells, as well as the input/output busses simpler, more uniform/efficient. The new 3D structure allows much more freedom in layout.

VIPRAM architecture is inherently open and flexible, making possible the design of more general purpose fast pattern recognition devices far beyond the original AM used for HEP

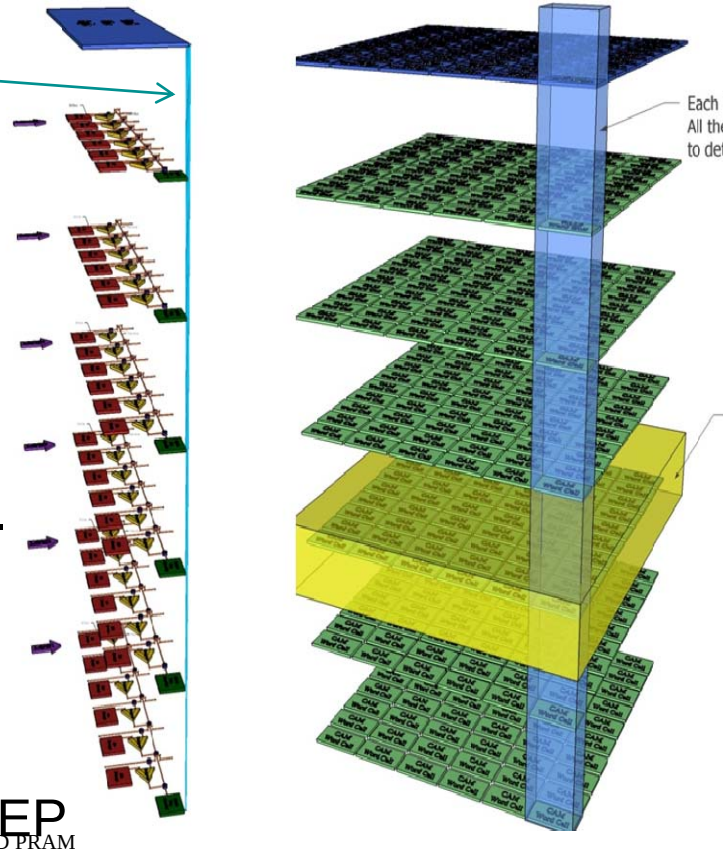

Fig. 4 - A 3D PRAM

VIPRAM

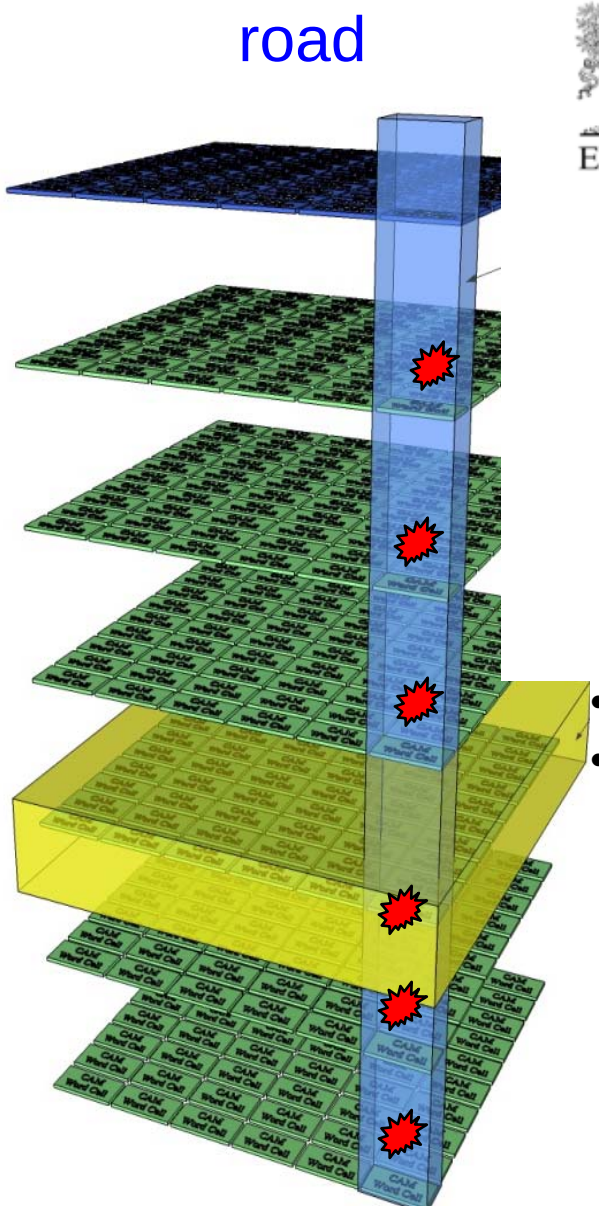

ELSEVIER

Available online at www.sciencedirect.com
ScienceDirect

Physics Procedia 00 (2011) 000–000

Physics Procedia

www.elsevier.com/locate/procedia

TIPP 2011 – Technology and Instrumentation for Particle Physics 2011

A New Concept of Vertically Integrated Pattern Recognition Associative Memory

Ted Liu*, Jim Hoff*, Grzegorz Deptuch, Ray Yarema

Particle Physics Division, Fermilab, P.O. Box 500, Batavia, IL 60510 USA

- Fermilab US Patent paper work (ROI)
- DOE CDRD award (over 3 years)

Design Work involved

Control/interface/readout design

Majority Logic cell design

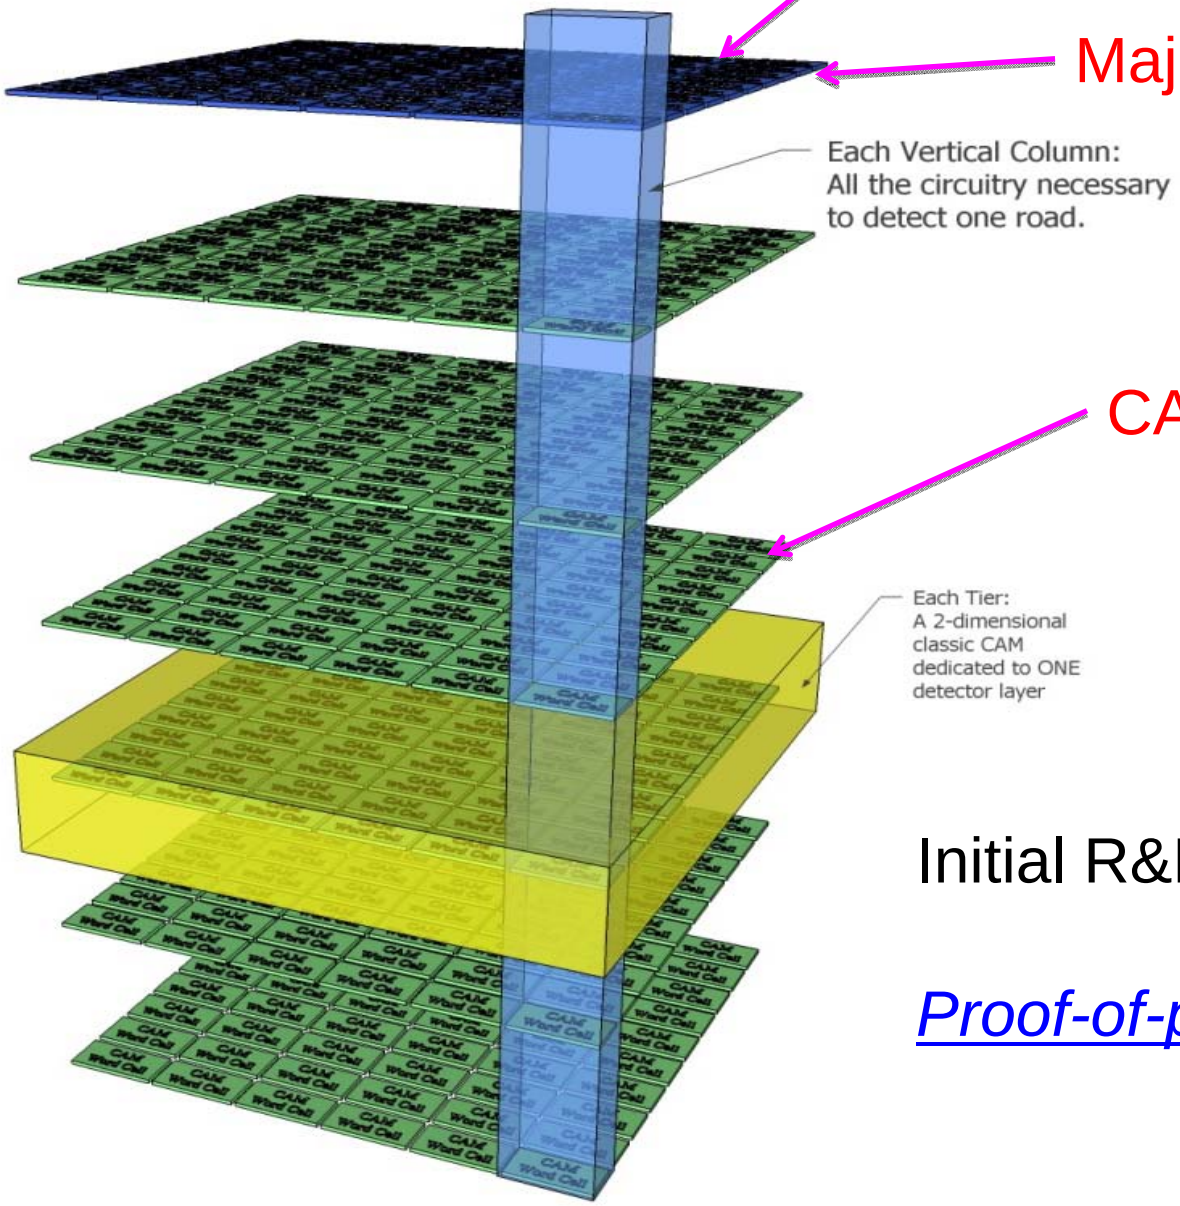

CAM cell design

Initial R&D goal:

Proof-of-principle demonstration

Initial Goal of R&D

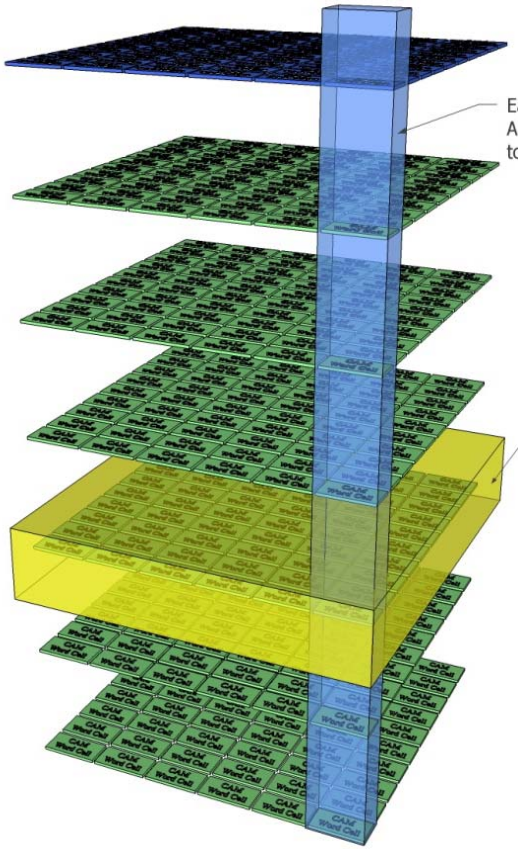

- *proof-of-principle* demonstration
- density: ~ 200K patterns/cm**2
- 130 nm technology node
- CDRD funding over 3 years
- *FNAL/ANL/UC*

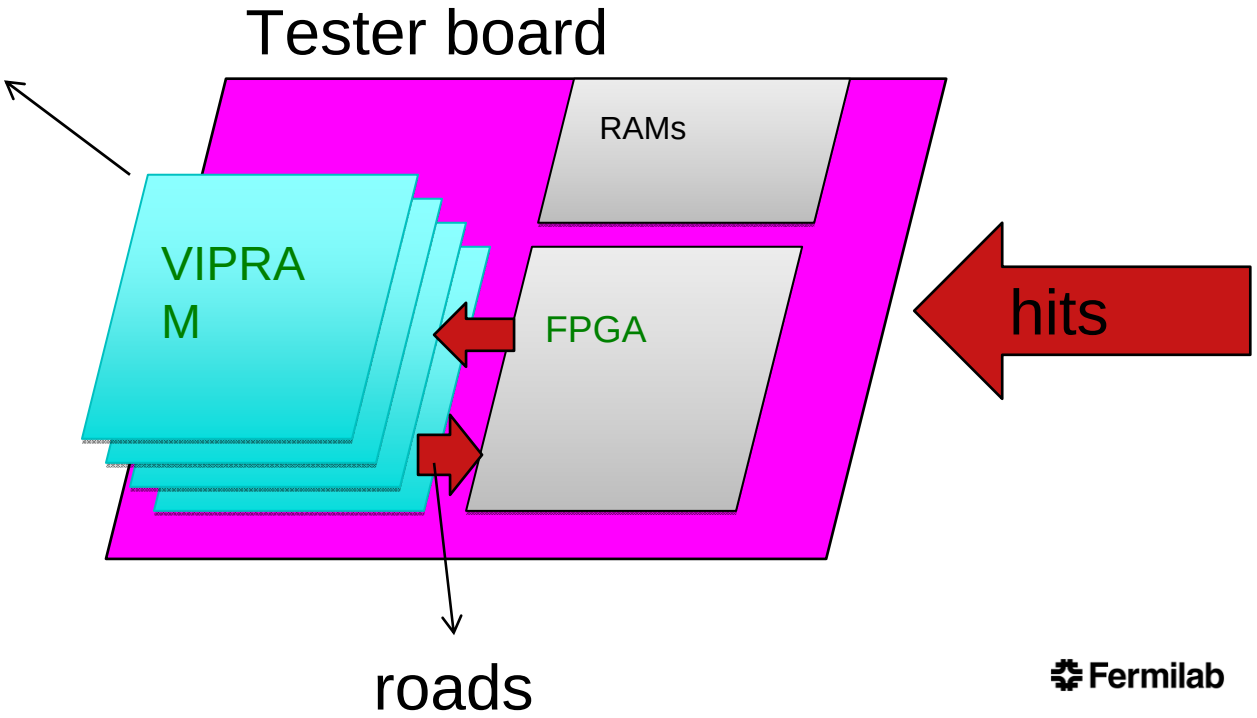

Integrate AM and TF/FPGA stages into one chip

Bandwidth between AM stage and Track Fitting stage could be another challenge

- needs to transfer large number of fired roads and associated full resolution hits into the TF stage
- The larger the AM pattern size per chip, the more demand
- *Highly desirable if the two stages can be integrated*
- High speed serial I/O on FPGA can be used for input data IO
- *Board & system level design could be much simplified*

3D Technology could help here (in the future)

- Example: silicon interposer approach for Xilinx Virtex-7 FPGA
- *Would make the chip much more flexible (within & outside HEP)*

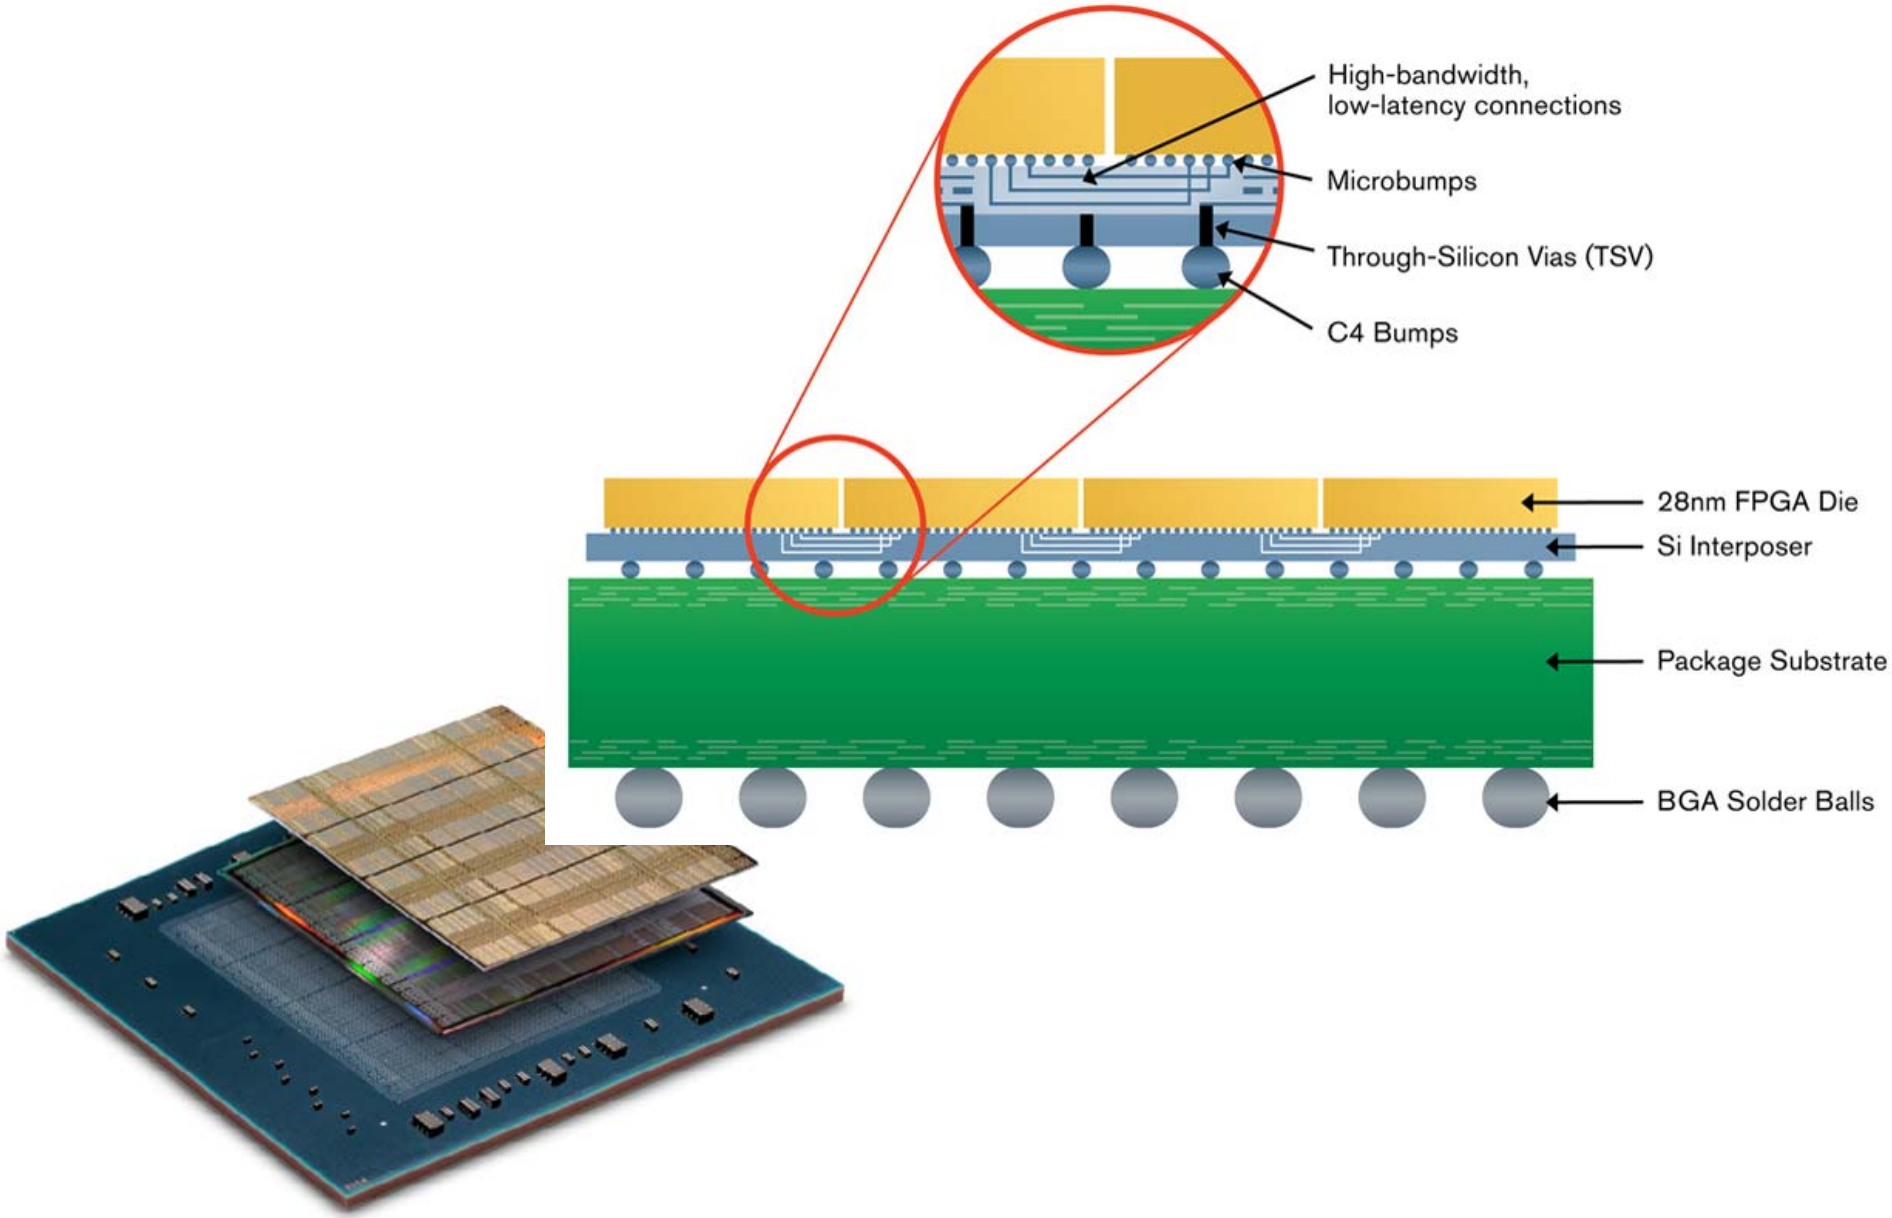

Virtex-7 2000T FPGA Utilizing
Stacked Silicon Interconnect Technology

Fermilab

SVT in one chip: 2nd phase of VIPRAM project

➔ general purpose chip

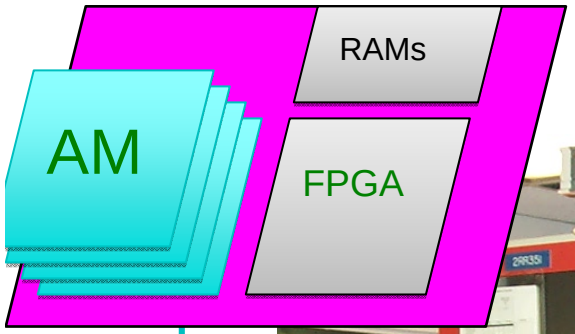

Original SVT system had ~400K patterns total

*Aim to reach ~500K per cm**2 for VIPRAM chip …*

Tracking Trigger Issues at hadron colliders

Detector design for triggering

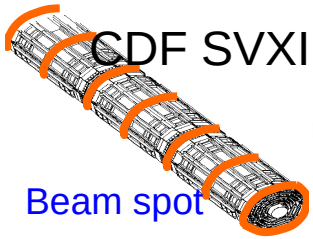

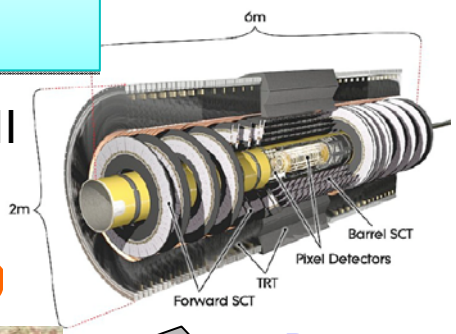

Data transfer

Data formatting

Pulsar-IIa

Pattern Recognition

3D Associative Memory

Others …

- *New ideas/approaches in four areas*
- *Three of them fully developed*

Track Fitting

FPGA vs CPU vs GPU → HLT

(1) Data Formatter using ATCA technology

(2) 3D AM R&D: VIPRAM

= "Vertically Integrated Pattern Recognition Associative Memory"

(3) GPU study for low latency trigger applications

As luminosity/occupancy increase, track fitting task becomes challenging

GPUs (Graphic Processing Units) have evolved into highly parallel, multi-threaded, multicore processors with remarkable computational power and high memory bandwidth. The GPUs also come with software environment that allows developers to use C/C++ as a high-level programming language, making it highly accessible to the general user. The combination of highly parallel architecture, high memory bandwidth as well as the user-friendly software environment makes GPUs a potentially promising technology for effective real-time processing for future high energy physics experiments.

GPU vs CPU evolution

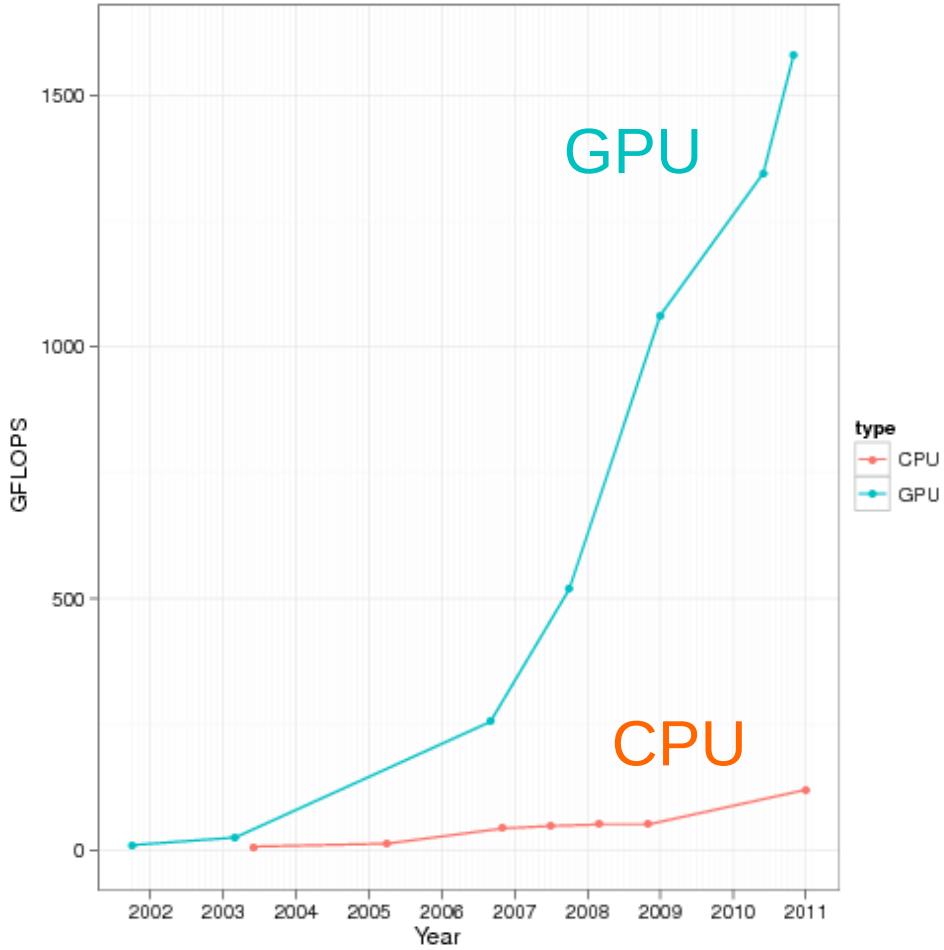

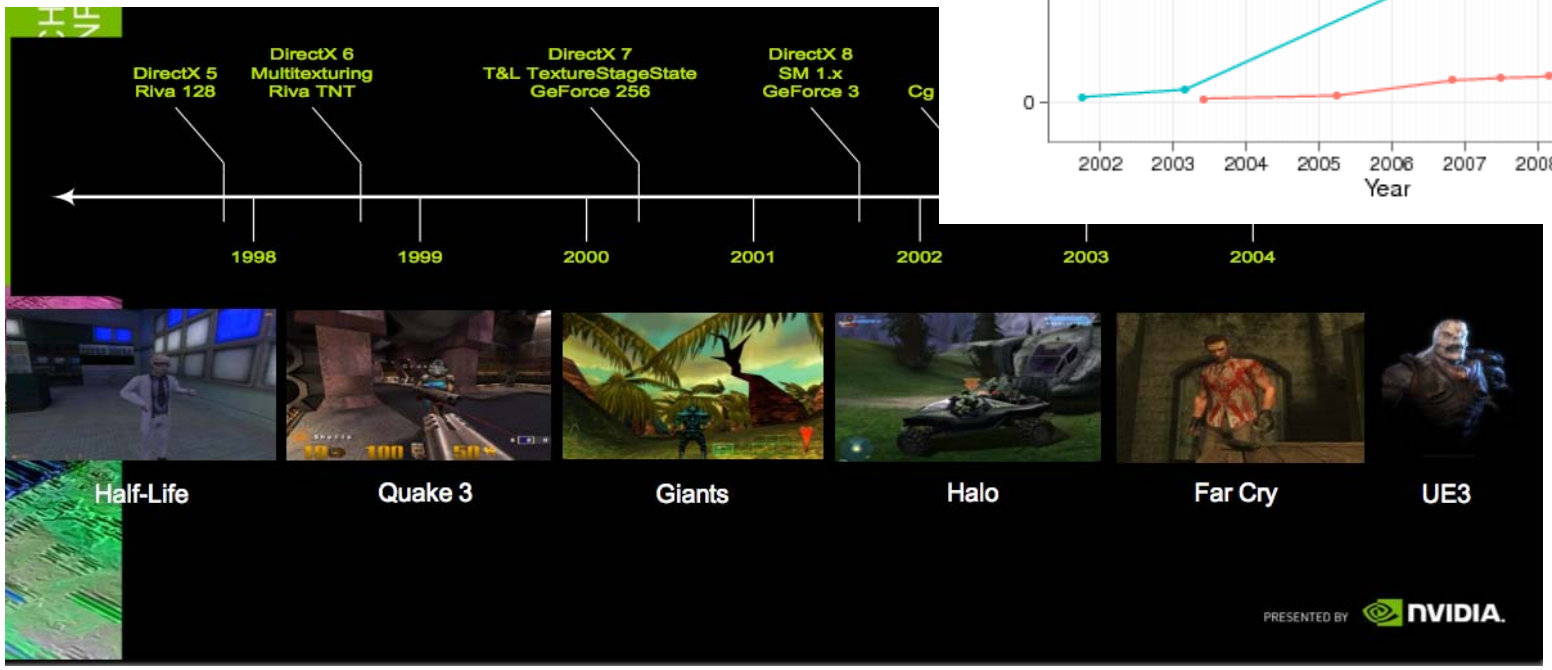

Fermilab

GPU performance study
fast track fitting

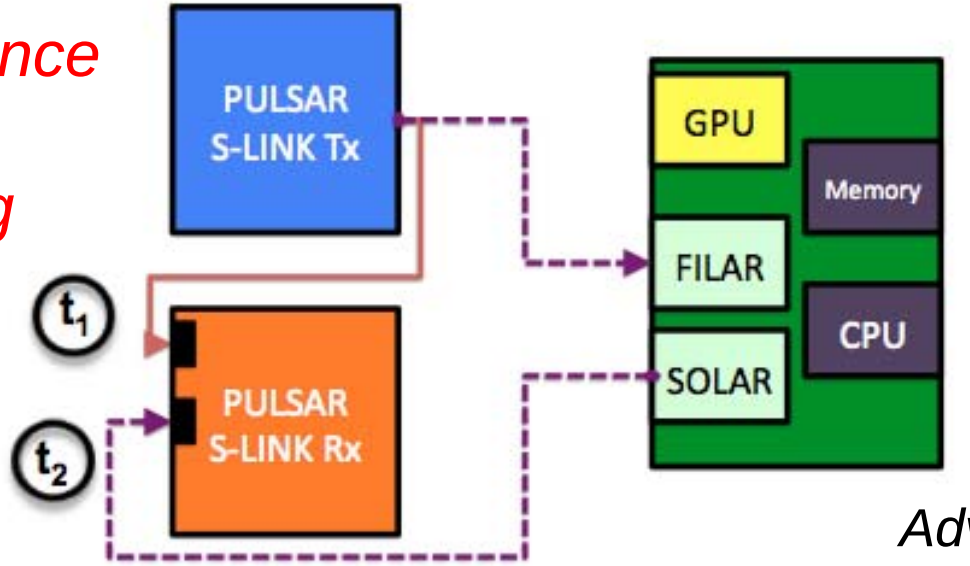

Advanced GPU lab

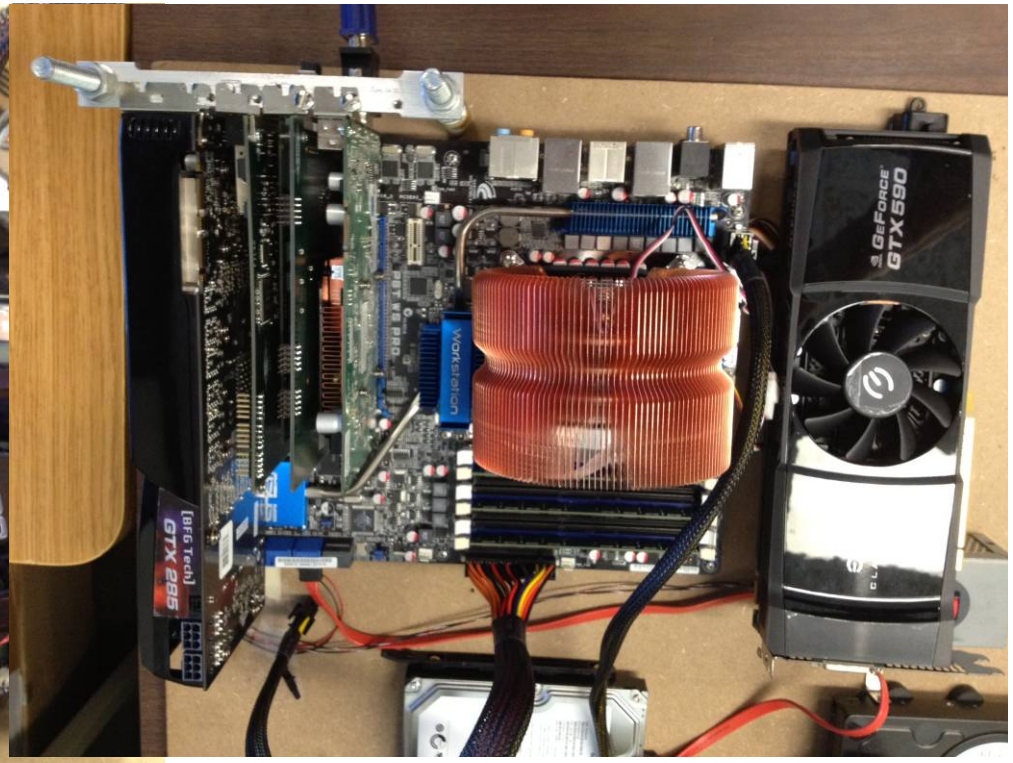

ELSEVIER

Available online at www.sciencedirect.com

ScienceDirect

Physics Procedia 00 (2012) 1–8

Physics Procedia

www.elsevier.com/locate/procedia

TIPP 2011 – Technology and Instrumentation in Particle Physics 2011

Performance Study of GPUs in Real-Time Trigger Applications for HEP Experiments

W. Ketchum[a], S. Amerio[b], D. Bastieri[b,c], M. Bauce[b,c], P. Catastini[d], S. Gelain[c], K. Hahn[d], Y. K. Kim[d,a], T. Liu[d], D. Lucchesi[b,c], G. Urso[e]

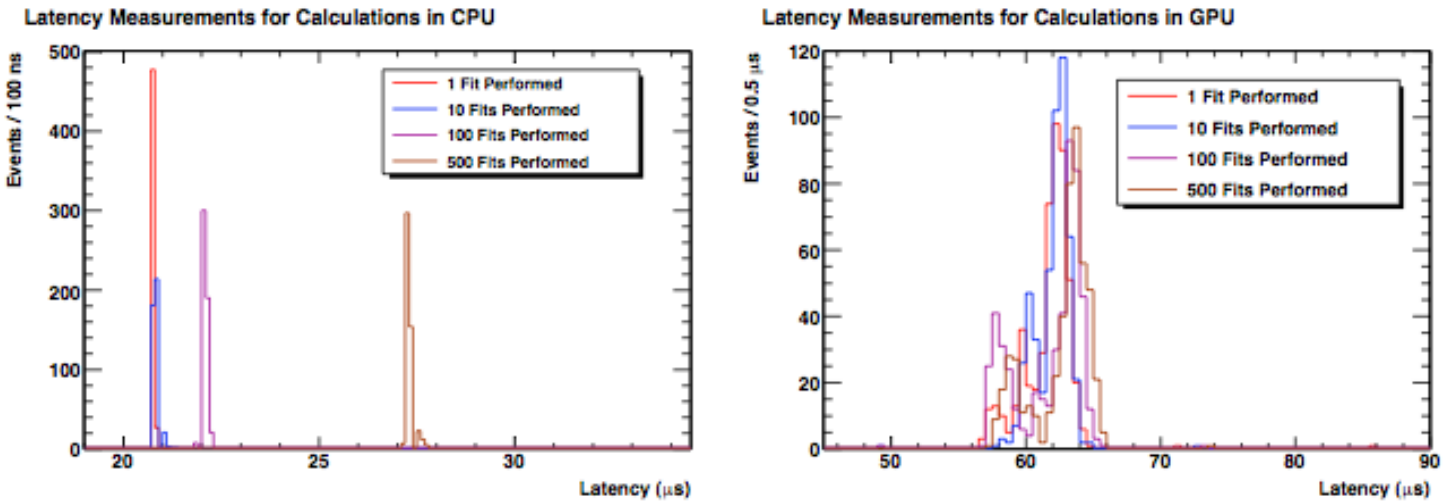

Figure 4: Latency when performing our track-fitting algorithm in the CPU (*left*) and GPU (*right*), varying the number of fits performed.

Initial results promising even with an old low-end gaming GPU from Best-Buy

Tracking Trigger Issues

Detector design for triggering

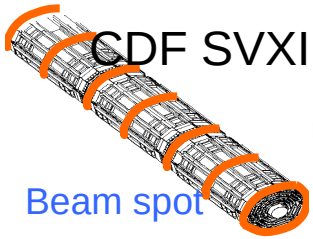

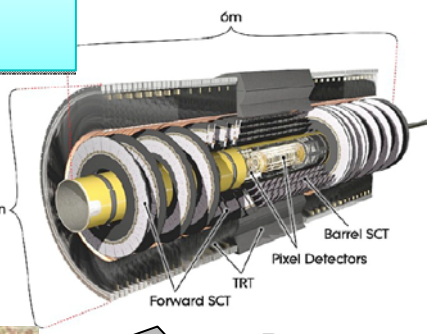

Data transfer

Data formatting

Pulsar II-a

Pattern Recognition

3D VIPRAM / SVT in one chip

Track Fitting

GPU

HLT

Combining the three developments, one could have a very powerful, general purpose modern trigger/pattern recognition/DAQ system … far beyond the HEP tracking trigger applications…

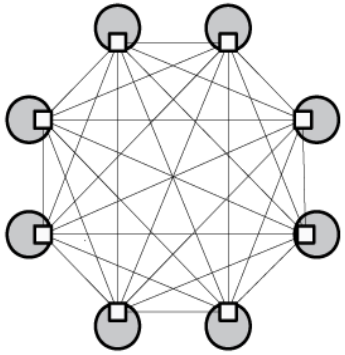

AdvancedTCA®

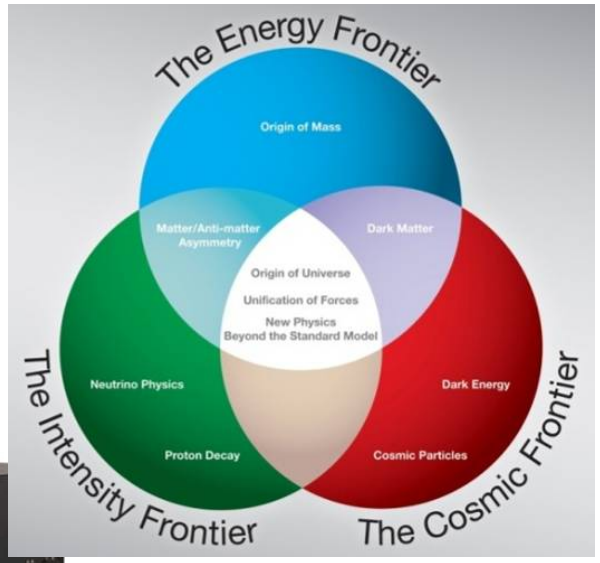

Fiber inputs

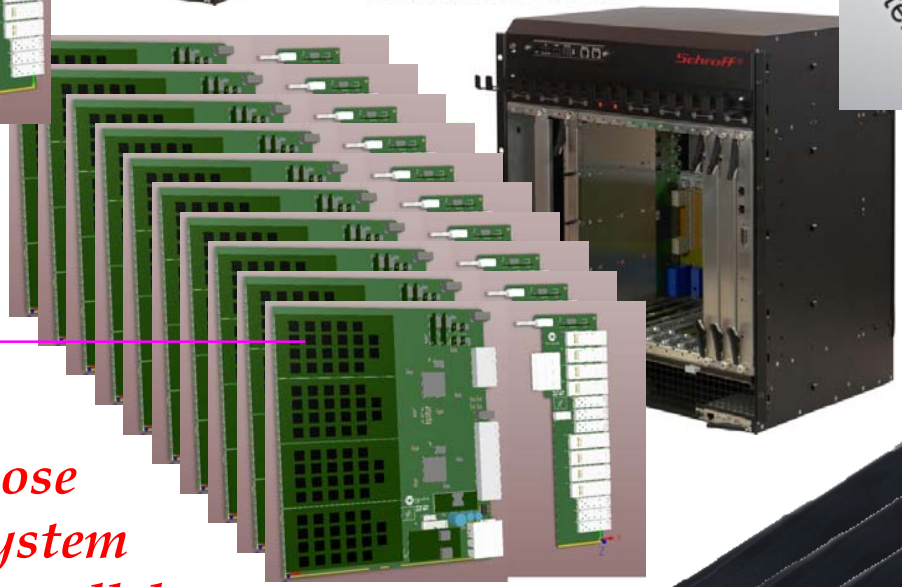

A bank of VIPRAM or SVT in a chip

General Purpose Trigger/DAQ System with MASSIVE parallel processing & pattern recognition power

GPU farm

Patterns ~ Billion / crate/shelf

Fermilab

Summary

Detector design for triggering

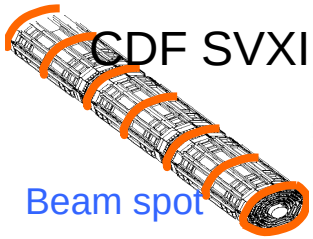

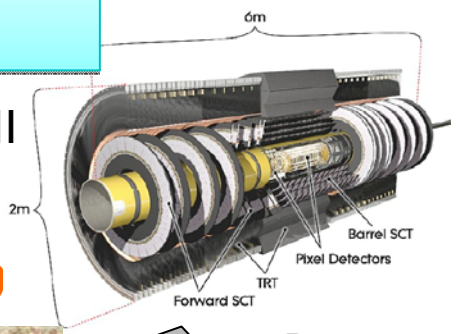

Data transfer

Data formatting

Pulsar II-a general purpose ATCA

Pattern Recognition

3D VIPRAM / SVT in one chip

Track Fitting

GPU

HLT

Ultimate physics at LHC in the future:

Go to Higher Energy: needs aggressive R&D in Magnet

Go to Higher Luminosity: needs aggressive Tracking Trigger R&D

Backup

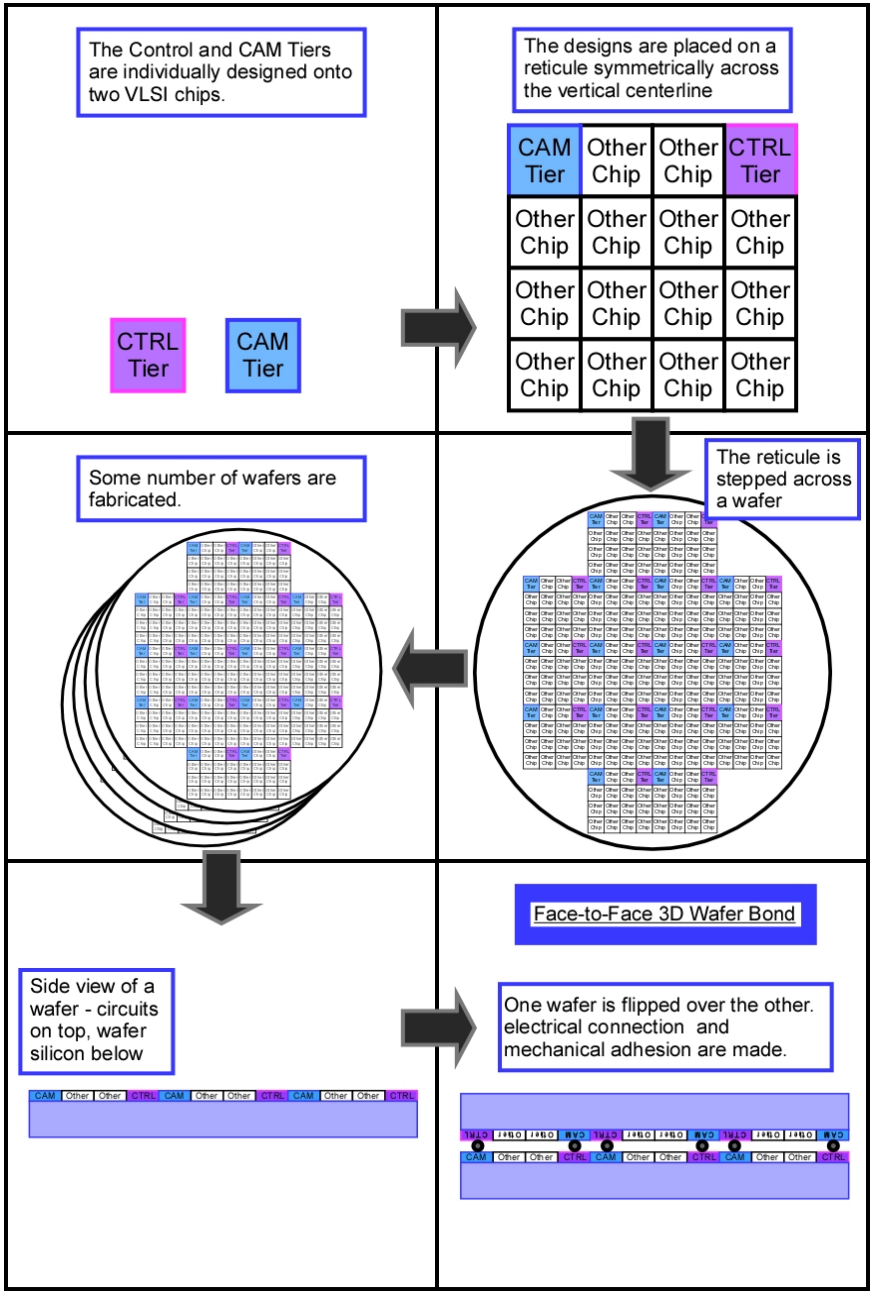

Figure 5 - A two -tier, Single Mask Set 3D MPW process

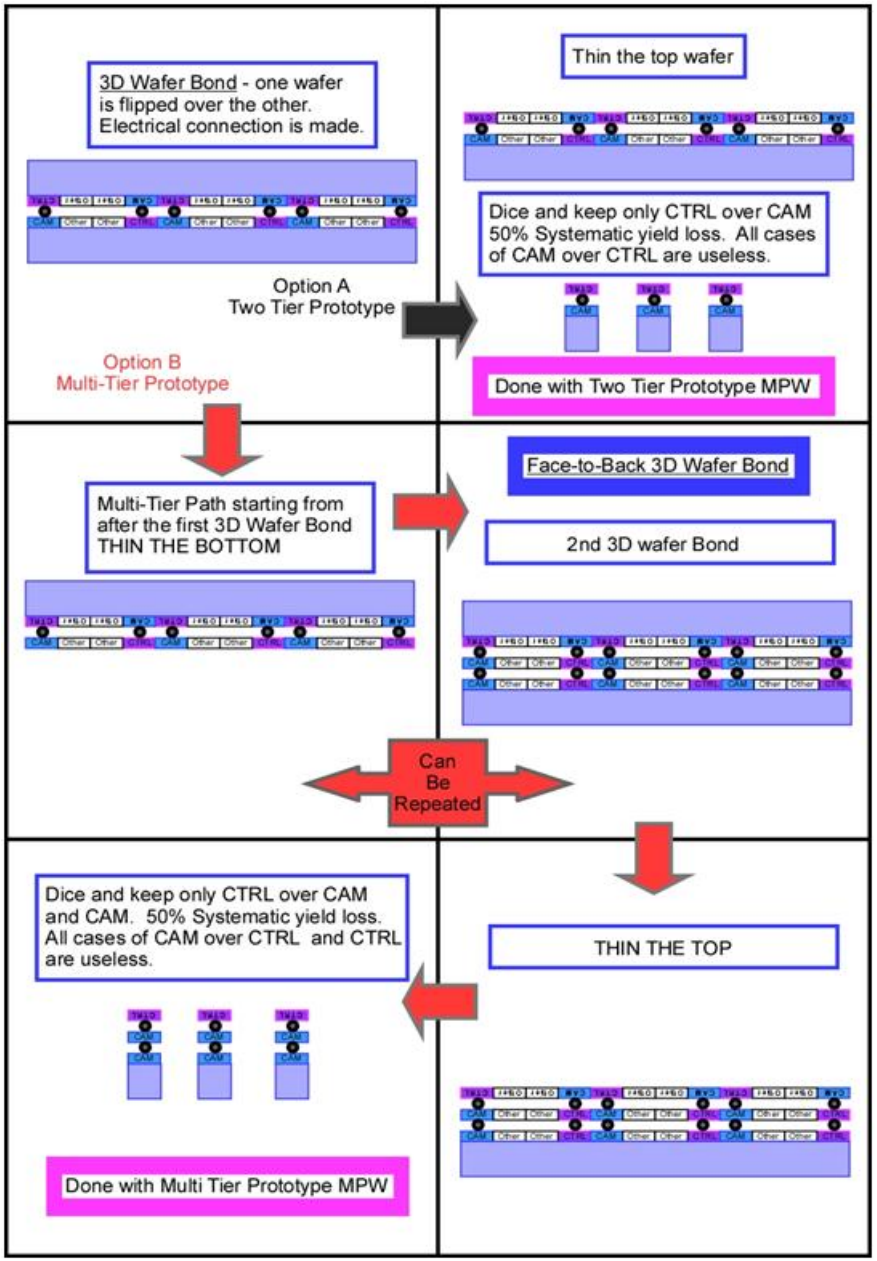

Figure 6 - The conclusion of a typical 3D MPW process OR an alternate process available to the VIPRAM.

Through Silicon Via (TSV)

A Solution Without a Problem for… Half of a Century

W. Shockley…

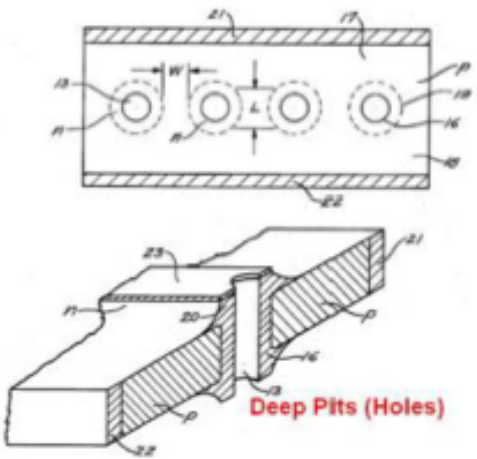

U.S. Patent # 3,044,909: "Semiconductive Wafer and Method of Making the Same", 1958

…and M. Smith & E. Stern

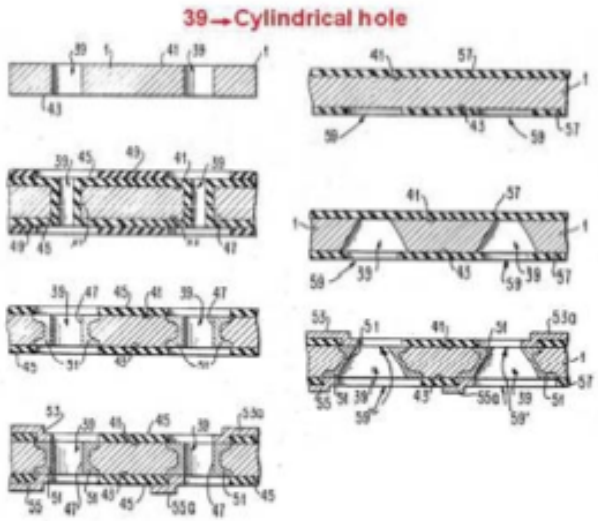

U.S. Patent #: 3,343,256 "Methods of Making Thru-Connections in Semiconductor Wafers", 1964

© Synopsys 2010

SYNOPSYS® Predictable Success

Fermilab

3D Technology in 30 seconds

3D technology: the integration of thinned and bonded silicon integrated circuits with vertical interconnects between IC layers

- Vertical interconnects: Through-Silicon-Vias (TSVs)
- Applications: memories, pixel arrays, microprocessors & FPGAs

Performance can be improved by reducing interconnect R/L/C for higher speed and density…

Freedom to divide functionality among tiers to create new designs that are simply not possible in 2D

- Useful when a task can be partitioned into multiple sections that are physically and logically separable, and the interconnects among them are straightforward

Moore's law is approaching severe limitations

3D could be the next scaling engine

Not just as merely an extension of Moore's law,

also provides novel design opportunities

Examples of commercial applications of 3D Technology

Increase density dramatically

- Example: 3D DRAM stacking (control/interface tier + memory cell tiers)
- Footprint or size reduction has been the main driving factor
- Available commercially in embedded, wireless, and memory devices

Increase memory access bandwidth dramatically

- 3D integration of memory layers onto processor chip
- Eliminate the slower and higher-power off-chip buses (tens of ~ mm) by replacing them with high-bandwidth and low-latency short vertical interconnections (~ tens of um)
- Potential to remove some “fundamental bottlenecks” in computing

Both examples are relevant to AM R&D in 3D (see later)

Routing in 3D can be efficient, esp. if functional elements are arranged such that the interconnects among tiers are mostly vertical

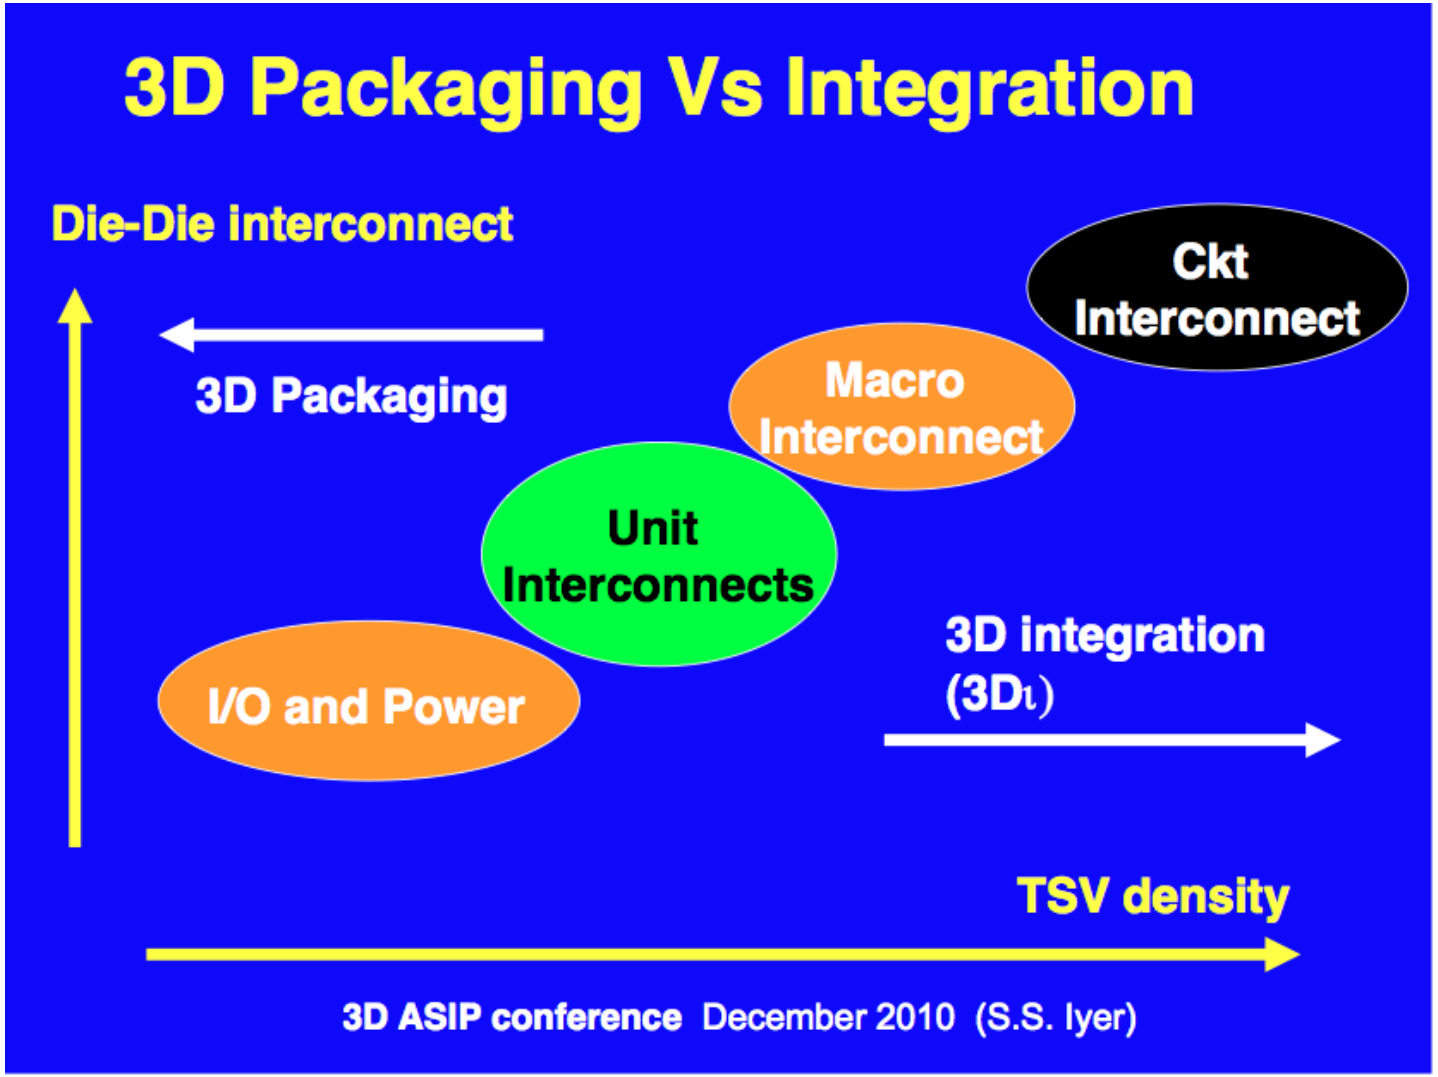
3D Packaging Vs Integration
Die-Die interconnect
3D Packaging
Ckt Interconnect
Macro Interconnect
Unit Interconnects
I/O and Power
3D integration (3Dι)
TSV density
3D ASIP conference December 2010 (S.S. Iyer)

How Real is 3D?

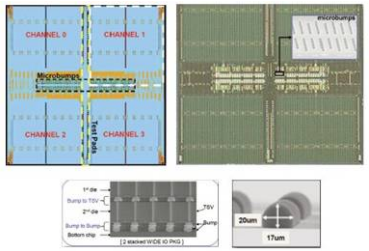

Samsung

16Gb NAND flash (2Gx8 chips),
Wide Bus DRAM

Micron

Wide Bus DRAM

Intel

CPU + memory

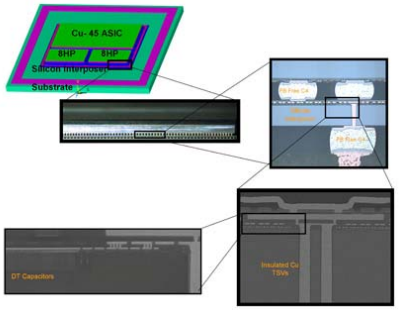

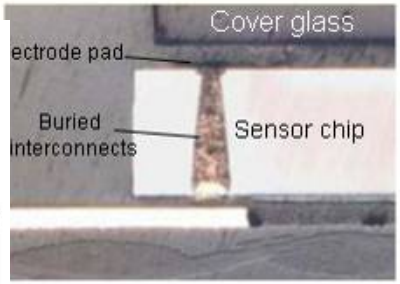

OKI

CMOS Sensor

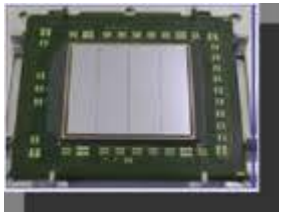

Xilinx

4 die 65nm interposer

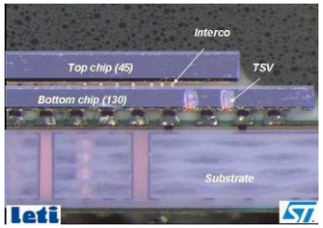

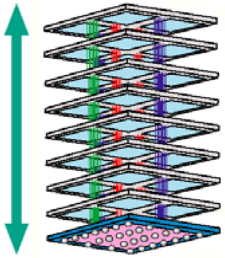

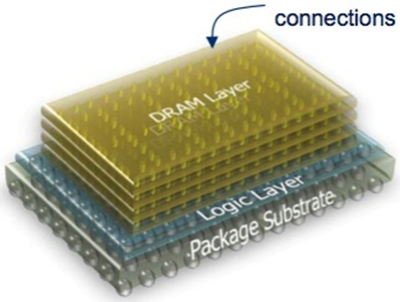

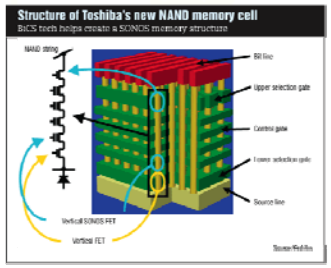

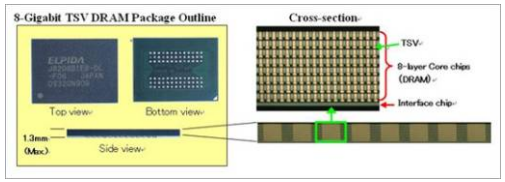

Raytheon/Ziptronix

PIN Detector Device

IBM

RF Silicon Circuit Board / TSV
Logic & Analog

Toshiba

3D NAND

Fermilab

Most recent successful 3D project

3D MAPS Processor V1

3D MAssively Parallel processor with Stacked memory

Architecture and Memory Model

- number and type of cores: 64, 5-stage, in-order, 2-way VLIW
- memory capacity: 256KB SRAM
- 3D stacking: 2 tiers face-to-face bonded (= core + memory)
- memory model: dedicated 4KB SRAM tile per core
- memory latency: 1 clock cycle, 1 read per every instruction
- memory bandwidth achieved: 61.3GB/sec peak achievable

Technology, Performance, and Power

- technology: Chartered Semiconductor 130nm
- footprint area: 5mm x 5mm
- clock frequency: 277MHz
- operating voltage: 1.5V
- maximum power consumption: up to 4W

Reliability

- maximum IR-drop: up to 78mV
- maximum coupling noise: 574 mV
- clock skew/slew: skew = 82ps, slew = 117ps
- maximum temperature: coming up

TSVs, Face-to-face (F2F) Vias, and IOs

- TSV diameter and pitch: 1.2um, 5um (Tezzaron)
- F2F via diameter and pitch: 3.4um, 5um (Tezzaron)
- total TSV count: 2240 (= 35x64) dummy, 47940 (= 204x235) IO
- total F2F via count: 7424 (= 116x64) signal, 43776 (= 684x64) P/G
- total IO count: 14 signal, 205 P/G (1.5V), 16 P/G (2.5V)

Georgia Tech/US DOD, 3D-MAPS Processor V1

- arguably the FIRST many-core 3D processor from academia
- designed to demonstrate memory BW/power benefit of 3D processor

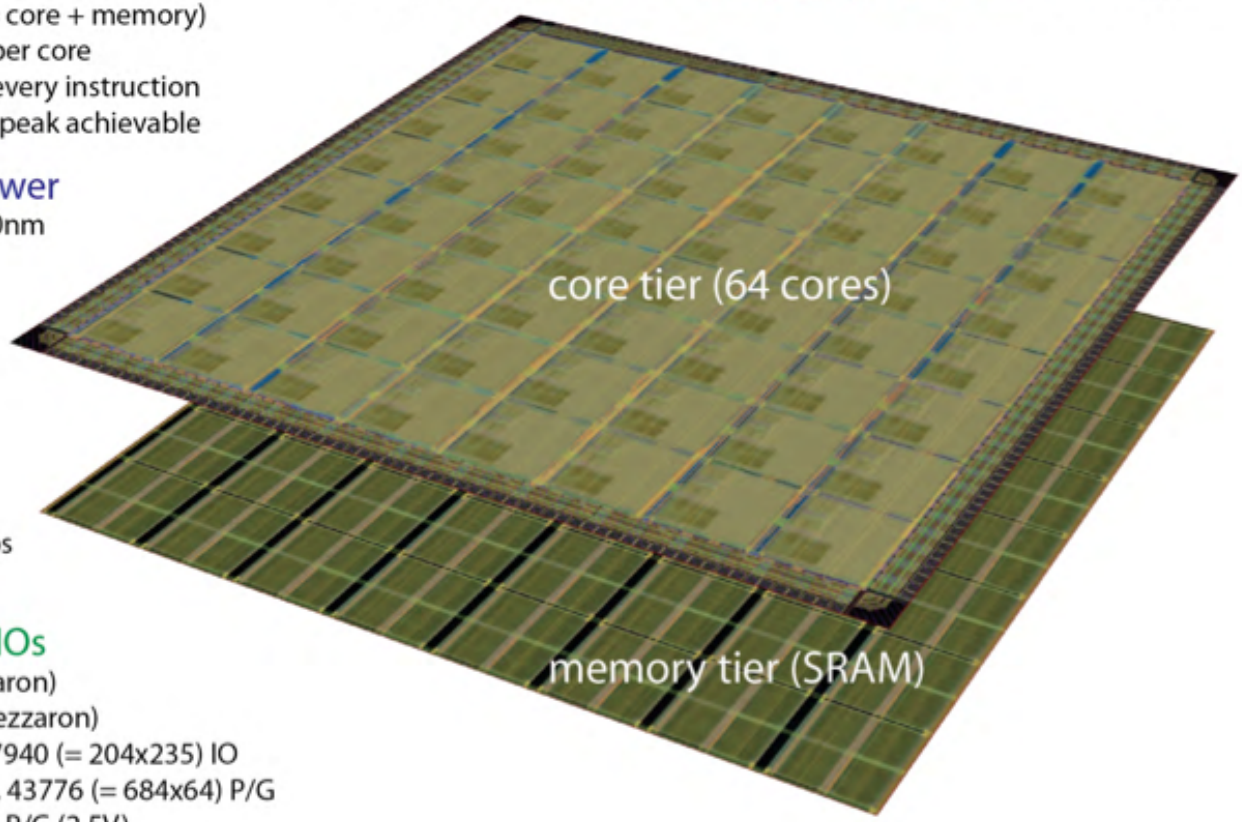

From http://www.gtcad.gatech.edu/3d-maps/

Next version (V2) is coming…

3D-MAPS V1 vs V2

- **MOSIS/Tezzaron 3D IC MPW (taped out: Oct 2011)**

	3D-MAPS V1	3D-MAPS V2
# of tiers	2 (1 logic, 1 SRAM)	5 (2 logic, 3 DRAM)
# of cores	64	128
Memory capacity	256KB SRAM	256MB DRAM & 512KB SRAM
Logic footprint	5mm X 5mm	10mm X 10mm
DRAM footprint	-	20mm X 12mm
Bonding style	F2F	F2F and F2B
TSV/F2F usage	~ 50K / ~50K	~ 150K / ~185K
Memory access*	2048 bit/cycle SRAM	1024 bit/cycle DRAM
freq / power	277MHz / 4.0W	175MHz / 10.4W

* Wide-I/O allows 512 bit/cycle DRAM access

The following slides are from:
http://web.eecs.umich.edu/~dfick/files/fick_isscc2012_slides.pdf

Centip3De: A 3930 DMIPS/W Configurable Near-Threshold 3D Stacked System With 64 ARM Cortex-M3 Cores

David Fick, Ronald G. Dreslinski, Bharan Giridhar,
Gyouho Kim, Sangwon Seo, Matthew Fojtik,
Sudhir Satpathy, Yoonmyung Lee, Daeyeon Kim,
Nurrachman Liu, Michael Wieckowski, Gregory Chen,
Trevor Mudge, Dennis Sylvester, David Blaauw

University of Michigan

Michigan Integrated Circuits Laboratory

Centip3De System Overview

- 7-Layer NTC system
- 2-Layer system completed fabrication with measured results
- Full 7-layer system expected mid 2012

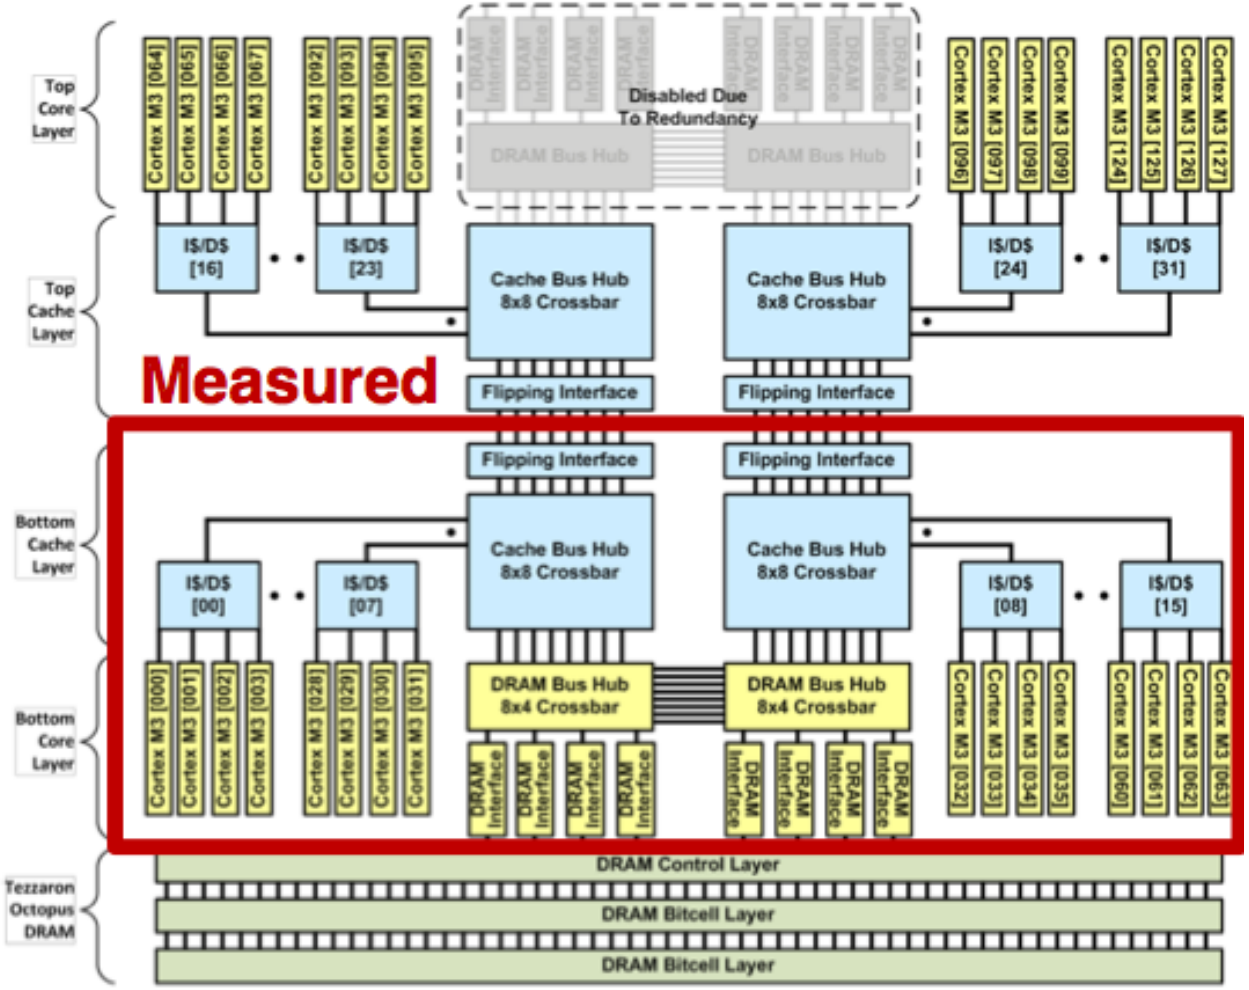

Centip3De System Overview

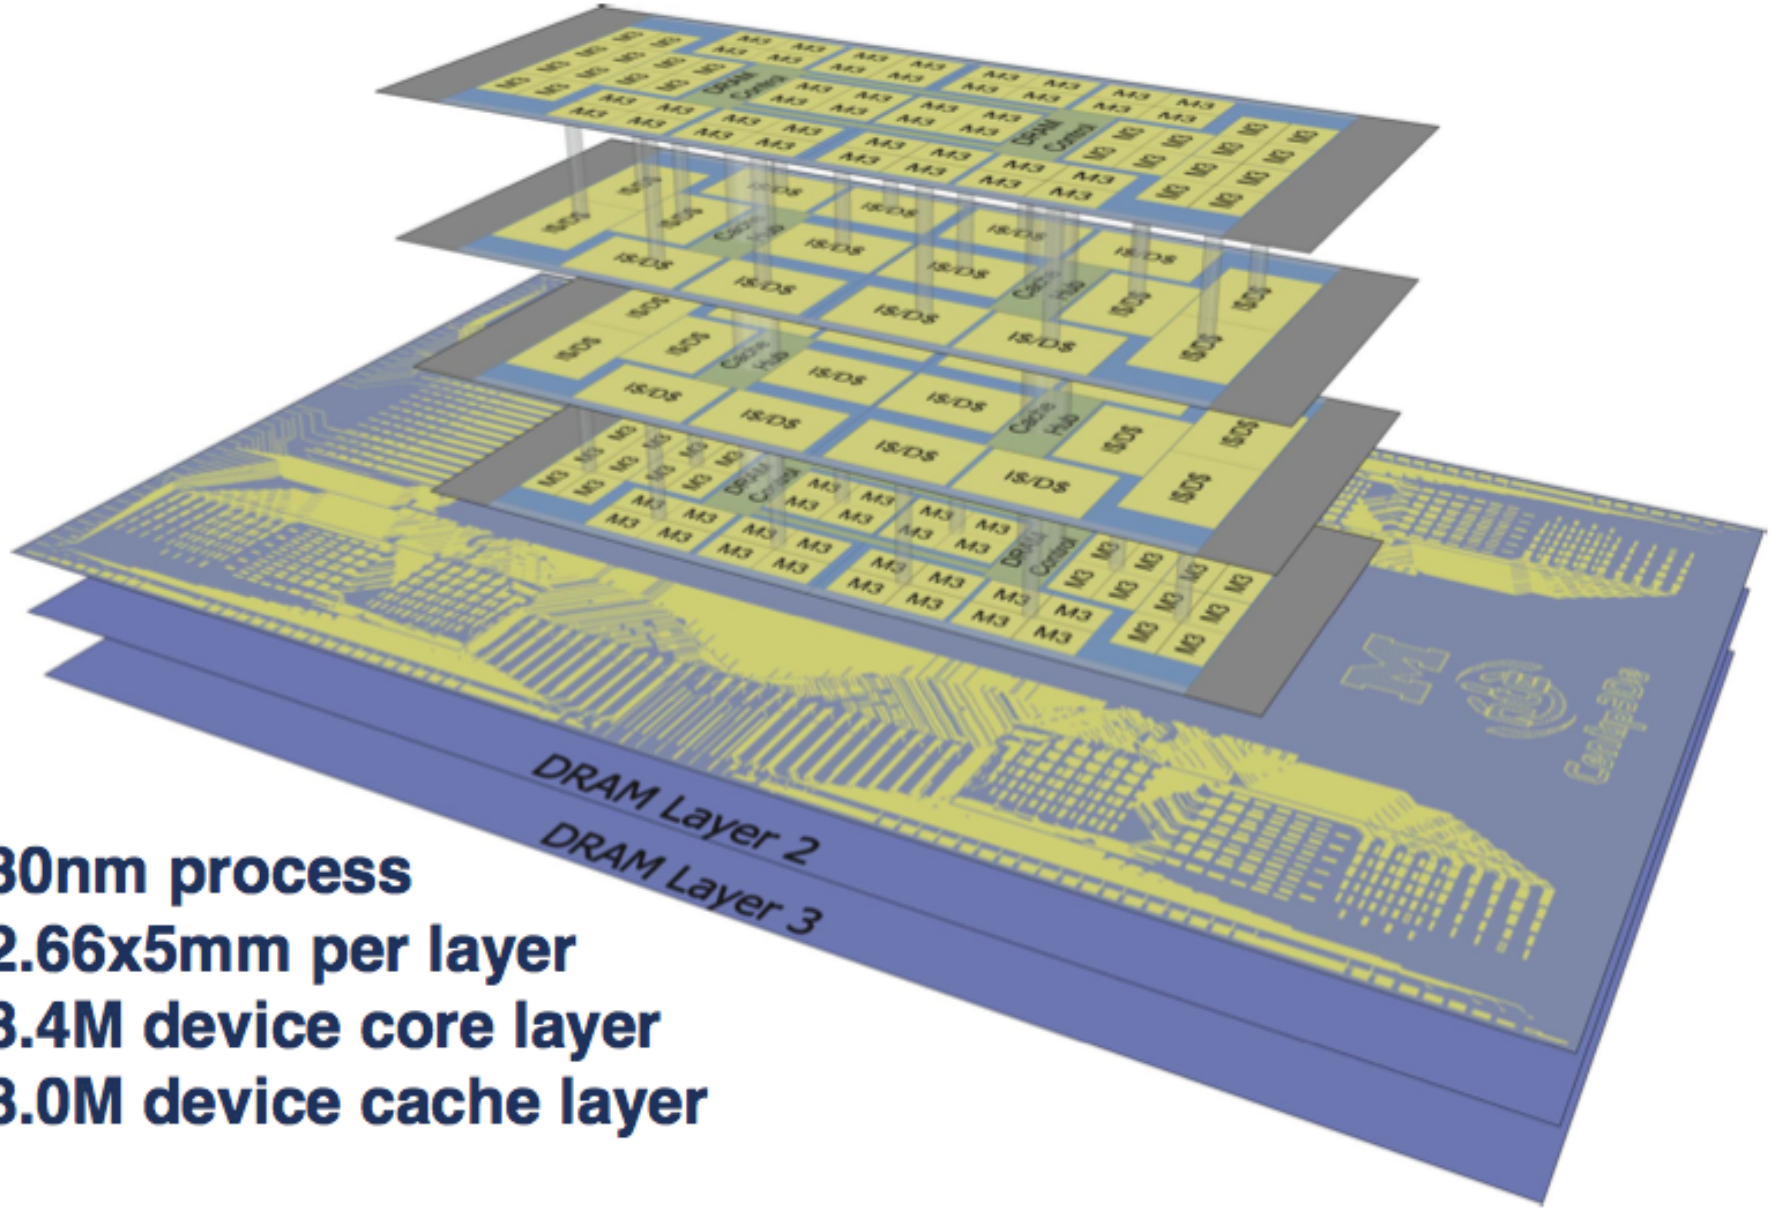

130nm process
12.66x5mm per layer
28.4M device core layer
18.0M device cache layer