



*Associative Memory R&D
in “extra dimension”*

Ted Liu
Fermilab

WIT 2012, Pisa
May 5, 2012

Detector design for triggering



Tracking Trigger Issues at hadron colliders (L1&L2)

Data transfer

Data formatting

Pattern Recognition

Associative Memory approach Others ...

Track Fitting

FPGA vs GPU


HLT

•AM in "extra dimension"


3D AM R&D: VIPRAM

= "Vertically Integrated Pattern Recognition Associative Memory"

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
 - ...

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





CDF SVX II





Channels used for SVT:
~ 0.2 millions



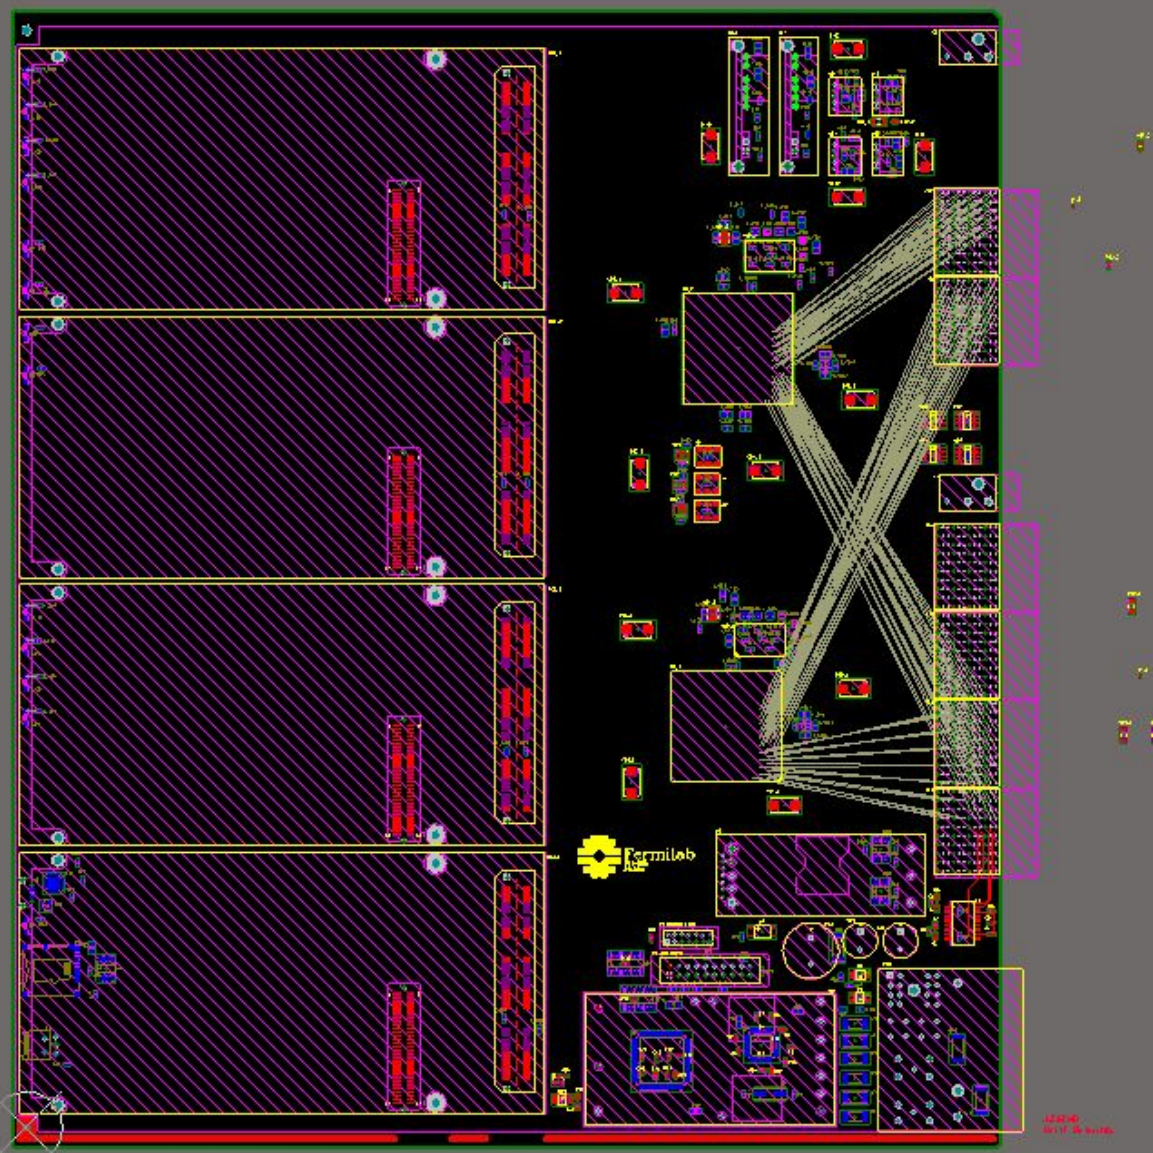
Data Formatting Challenges



3D models with ROD ID mapping
-- built by our new postdoc
Yasu Okumura



Data Formatting Challenges: the need for ATCA design



5 Data Formatter ATCA board design at Fermilab with full-mesh backplane for data sharing



Challenge in Tracking Trigger

- The PAST: hardware-based pattern recognition for fast track triggering has been very successful for HEP
 - CDF SVT: based on Associative Memory for pattern recognition
 - SVT ~ 400K patterns --> 6 Million patterns --> > 1 Billion at LHC at high luminosity
- **THE FUTURE: enormous challenges** in implementing pattern recognition for tracking trigger at LHC (L1&L2), due to
 - much higher occupancy (pile up) and event rates at the LHC
 - detectors much more massive, much larger number of channels
- There is a clear need to significantly improve the hardware-based pattern recognition to advance the state-of-the-art
 - **Associative Memory R&D for HEP: beyond Moore's law ...**
 - Or think harder to come up smarter/crazier ideas ...

October 24, 1988

VLSI STRUCTURES FOR TRACK FINDING

Mauro DELL'ORSO

Dipartimento di Fisica, Università di Pisa, Piazza Tarantini 2, 56100 Pisa, Italy

Luciano RISTORI

INFN Sezione di Pisa, Via Vecchia Livornese 58, 56100 Pisa, Italy

Received 24 October 1988

We discuss the architecture of a device based on the concept of *associative memory* designed to solve the track finding problem, typical of high energy physics experiments, in a time span of a few microseconds even for very high multiplicity events. This "machine" is implemented as a large array of custom VLSI chips. All the chips are equal and each of them stores a number of "patterns". All the patterns in all the chips are compared in parallel to the data coming from the detector while the detector is being read out.

1. Introduction

The quality of results from present and future high energy physics experiments depends to some extent on the implementation of fast and efficient track finding algorithms. The detection of *heavy flavor* production, for example, depends on the reconstruction of secondary vertices generated by the decay of long lived particles, which in turn requires the reconstruction of the majority of the tracks in every event.

Particularly appealing is the possibility of having detailed tracking information available at trigger level even for high multiplicity events. This information could be used to select events based on impact parameter or secondary vertices. If we could do this in a sufficiently short time we would significantly enrich the sample of events containing heavy flavors.

Typical events feature up to several tens of tracks each of them traversing a few position sensitive detector layers. Each layer detects many hits and we must correctly correlate hits belonging to the same track on different layers before we can compute the parameters

2. The detector

In this discussion we will assume that our detector consists of a number of layers, each layer being segmented into a number of *bins*. When charged particles cross the detector they *hit* one bin per layer. No particular assumption is made on the shape of trajectories: they could be straight or curved. Also the detector layers need not be parallel nor flat. This abstraction is meant to represent a whole class of real detectors (drift chambers, silicon microstrip detectors etc.). In the real world the coordinates of each hit will actually be the result of some computation performed on "raw" data: it could be the center of gravity of a cluster or a charge division interpolation or a drift-time to space conversion depending on the particular class of detector we are considering. We assume that all these operations are performed upstream and that the resulting coordinates are "binned" in some way before being transmitted to our device.

A brief history ...
All started here at Pisa

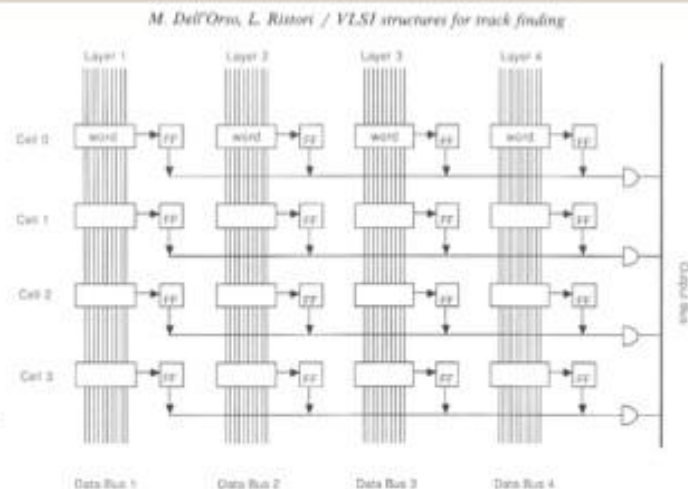


Fig. 3. Associative memory architecture.

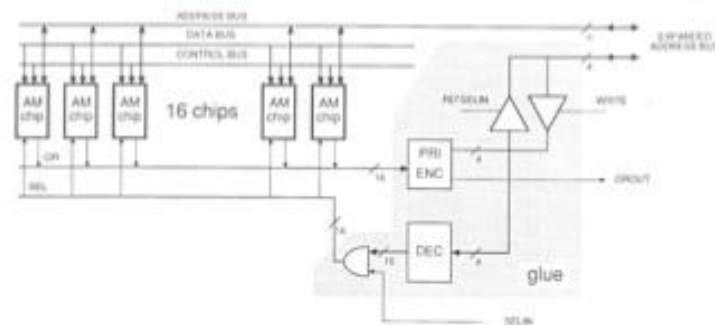


Fig. 5. 16 AM chips tied by the "glue".

We discuss the architecture of a device based on the concept of *associative memory* designed to solve the track finding problem, typical of high energy physics experiments, in a time span of a few microseconds even for very high multiplicity events. This "machine" is implemented as a large array of custom VLSI chips. All the chips are equal and each of them stores a number of "patterns". All the patterns in all the chips are compared in parallel to the data coming from the detector while the detector is being read out.

CDF original SVT system had 384K patterns total ...
128 patterns per AMchip -- commissioned around ~2001.




Question:

Can we put the entire SVT system into one chip?

.... the rest of this talk...

How CAM works

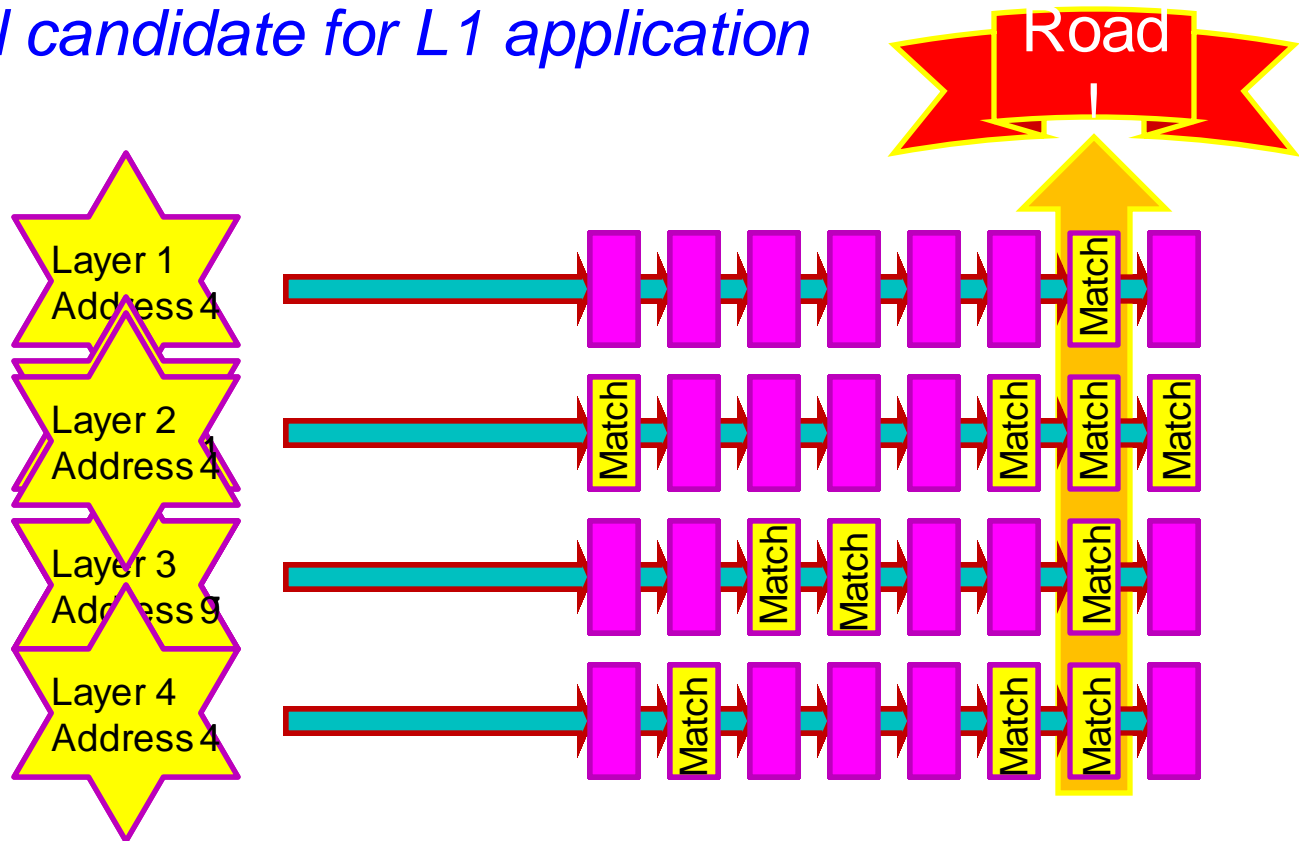
- CAM: inverse of RAM
 - user supplies a data word and it searches its *entire* memory *in a single operation* to see if that data word is stored anywhere in it



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

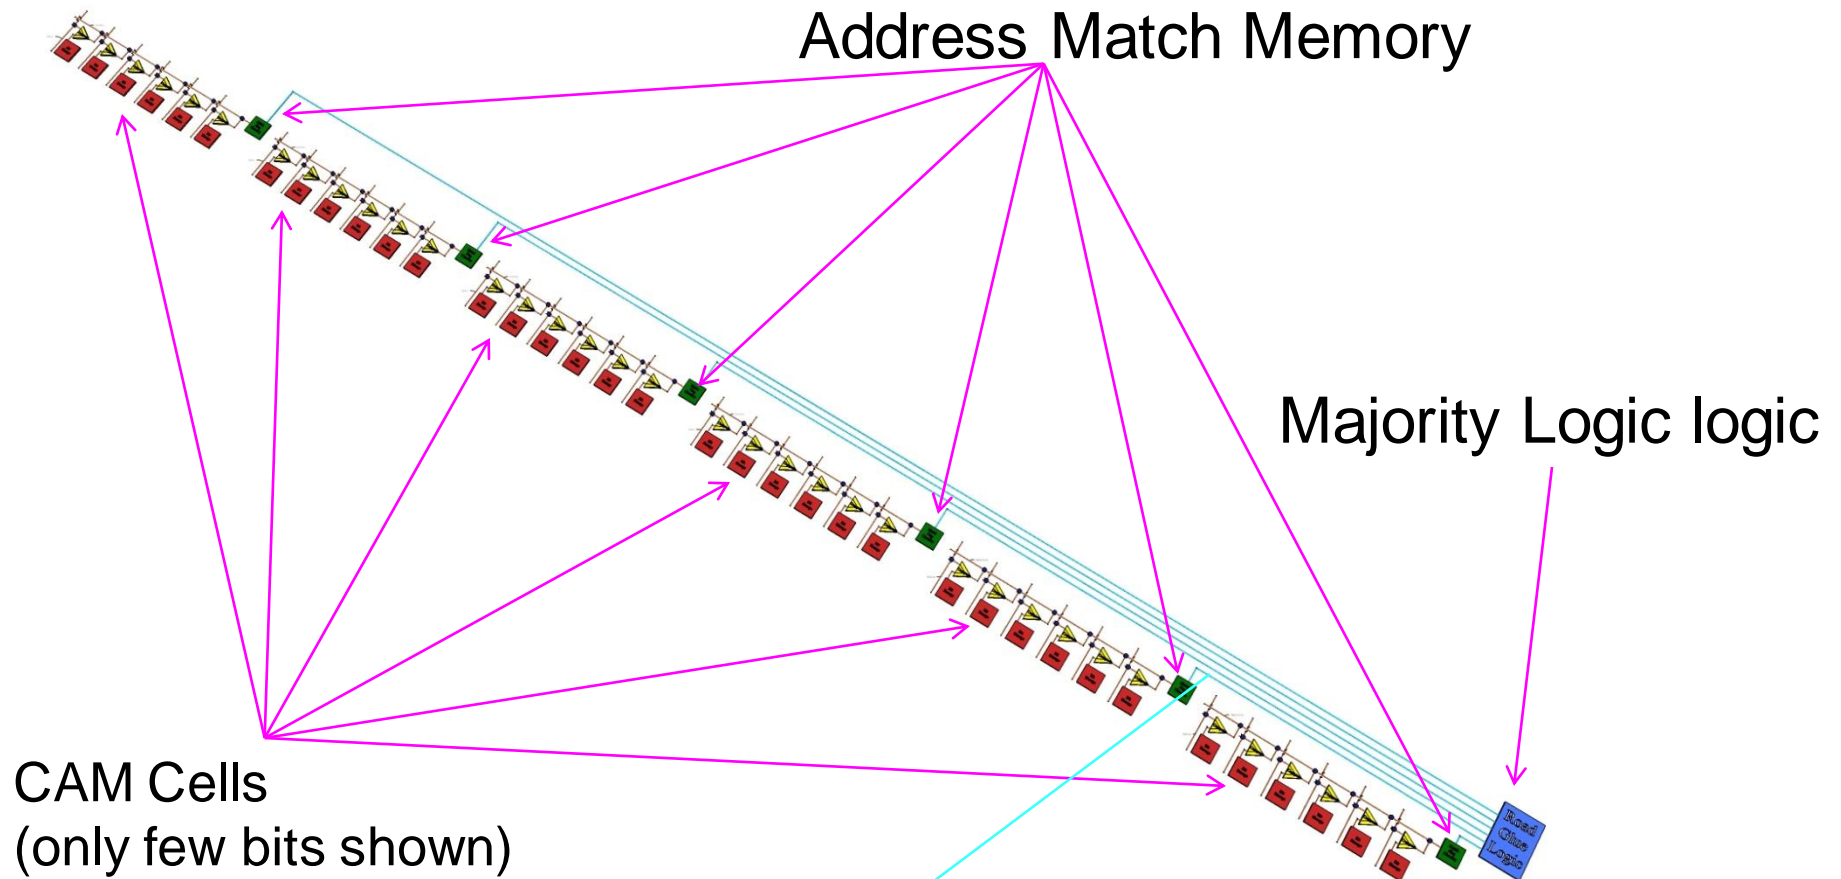
How PRAM works

- 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
 - As such, it contains large arrays of CAM cells and majority logic units that are reproduced many times and ordered in a fashion that is periodic in two dimensions
- 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.
- *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 $>\sim 3$,
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

One has to go to “extra dimension” to (possibly) achieve this
→ generic R&D effort at Fermilab

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

VIPRAM (Vertically Integrated Pattern Recognition Associative Memory)






Fig. 4 - A 3D PRAM

VIPRAM


(Vertically Integrated Pattern Recognition Associative Memory)

Pattern recognition for tracking is naturally a task in 3D




track

Side view



Top view



VIPRAM concept (developed at Fermilab):


http://hep.uchicago.edu/~thliu/projects/VIPRAM/TIPP2011_VIPRAM_Paper.V11.preprint.pdf

Advantages of VIPRAM approach (I)

- A VIPRAM cell can process n layers of a road pattern in about the size of just one CAM cell (*pattern density increased by $\sim n$*)
- 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.
- The top tier: a 2D array of signals that indicate whether or not a road has been flagged
 - Can be readout like a Pixel detector
 - Uniform TSVs distribution across



Advantages of VIPRAM approach (II)

- The VIPRAM 3D 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
 - N CAM tiers can handle N or 2xN detector layers... (flexibility in Majority Logic)
 - Would facilitate design reuse,
 - Would allow integration of different types of pattern recognition algorithm
 - Would allow integration of different types of detectors (such as Muon or CAL)
 - More fault tolerance with design for redundancy
 - Radiation tolerance issues
 - Possible use inside detector (also reduce latency)
- VIPRAM can also be used to match a single, broad word, thus acting as a conventional CAM with configurable width and less power consumption

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

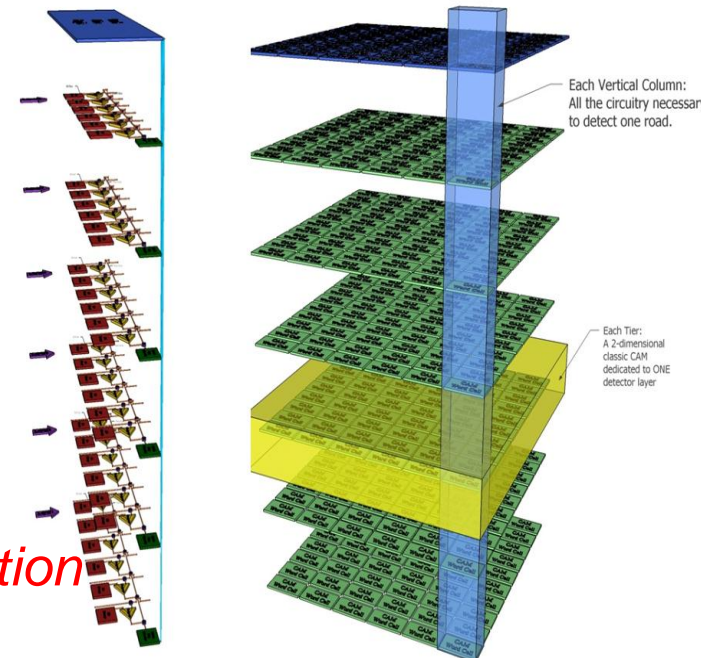



Fig. 4 - A 3D PRAM

Other interesting aspects of VIPRAM



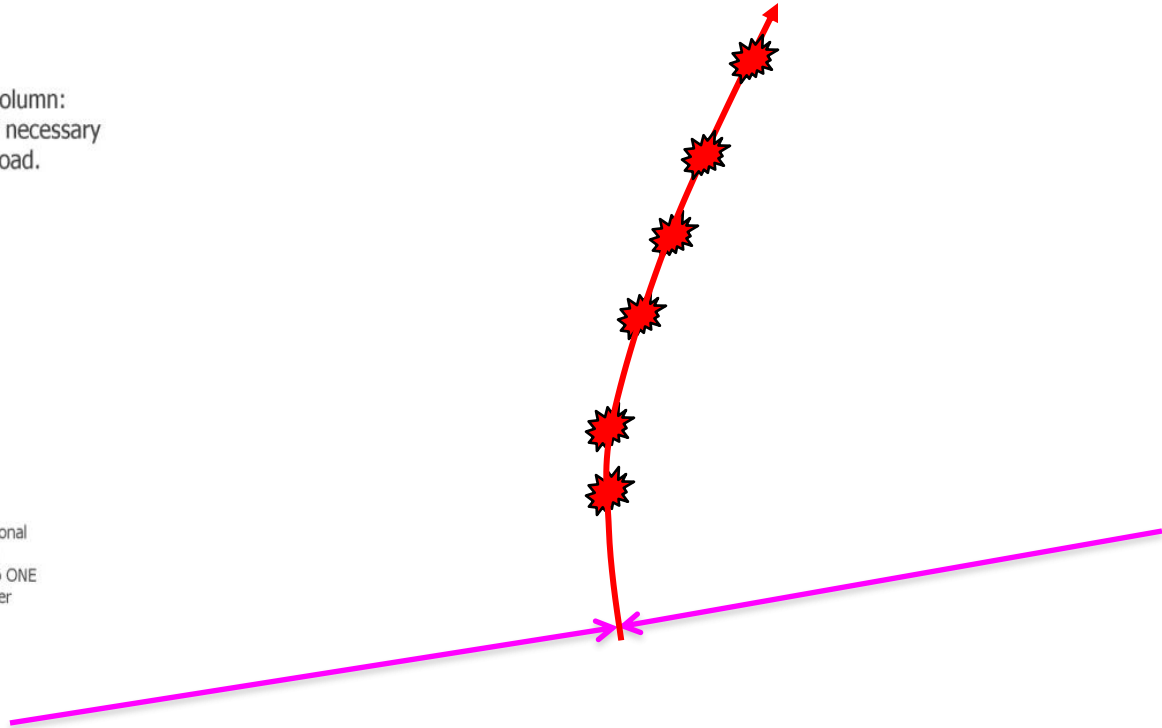

- *Power & Thermal modeling and analysis*
 - *CAM cell simple yet power hungry*
 - *VIPRAM 3D structure uniform*
- *Fault tolerance*
 - *Intrinsically forgiving*
 - *Redundancy design*
- *Radiation tolerance issues*
 - *How things scale with 3D*
 - *How to improve with 3D*
 - ...

VIPRAM generic R&D proposal to DOE

http://hep.uchicago.edu/~thliu/projects/VIPRAM/VIPRAM_DOE_LAB11-438-V2-submit.pdf

road

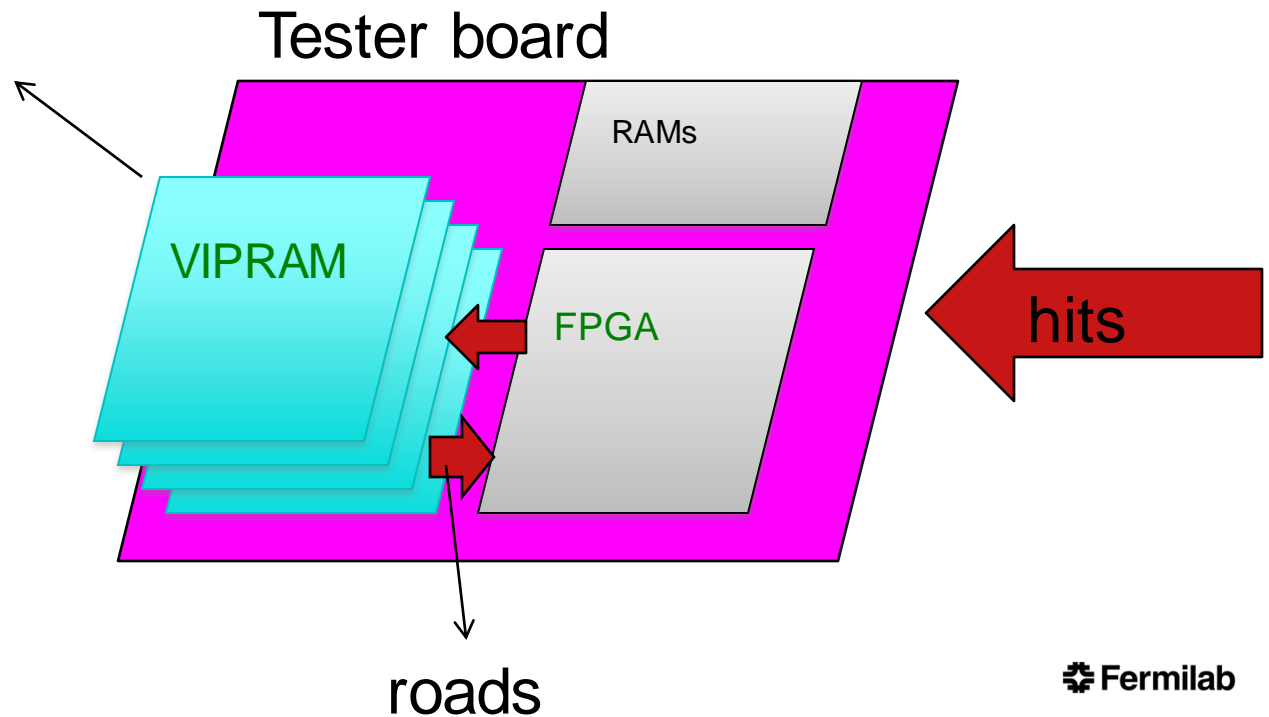

Submitted last April, FNAL/ANL/UC + SMU (EE)



*Potential applications outside HEP:
Pattern recognition in space and time*

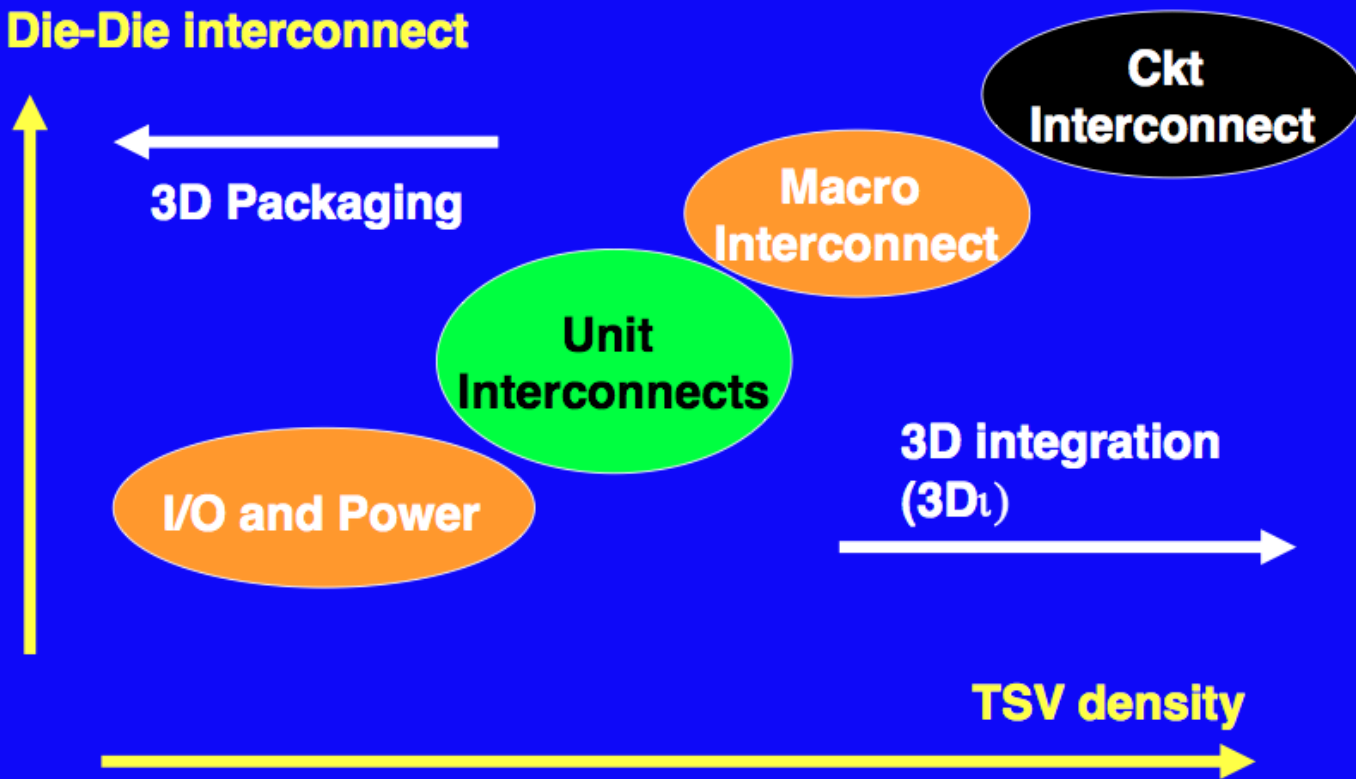
Initial Goal of R&D

- *proof-of-principle* demonstration
- Control tier + ~ 2 CAM tiers
- density: ~ 200K patterns/cm**2
- study performance vs speed/cost



3D Packaging Vs Integration

Die-Die Interconnect



3D ASIP conference December 2010 (S.S. Iyer)

As starting point for R&D

- 130nm Global Foundries CMOS
- Tezzaron' s 3D process
- ~18 bits in the CAM word (like AMchip04)
- initial design for stacking up to 4 CAM tiers
 - Only stack 1 Control + ~2 CAM tiers for proof-of-principle
- 4 μm center-to-center TSV spacing for compatibility with current Tezzaron' s 3D process.
- Simple to estimate the pattern density:
 - If a PRAM cell size is $\sim 20\mu\text{m} \times \sim 20\mu\text{m}$
 - ➔ this means $\sim 250\text{K}$ patterns per cm^{**2}

Design Work involved

Control/interface/readout design


Majority Logic cell design

CAM cell design

Initial design work by Fermilab ASIC group

Initial R&D goal:

Proof-of-principle demonstration



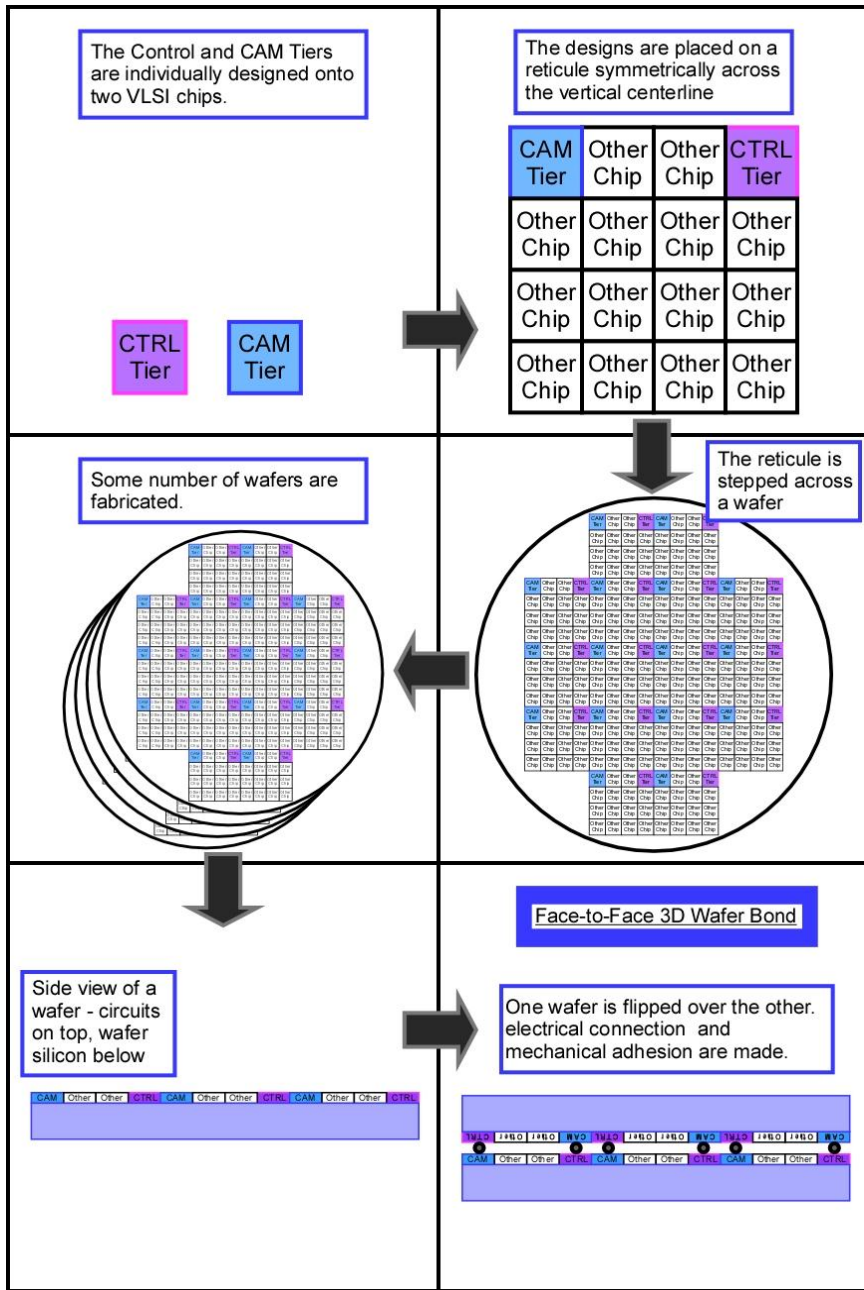


Figure 5 - A two-tier, Single Mask Set 3D MPW process

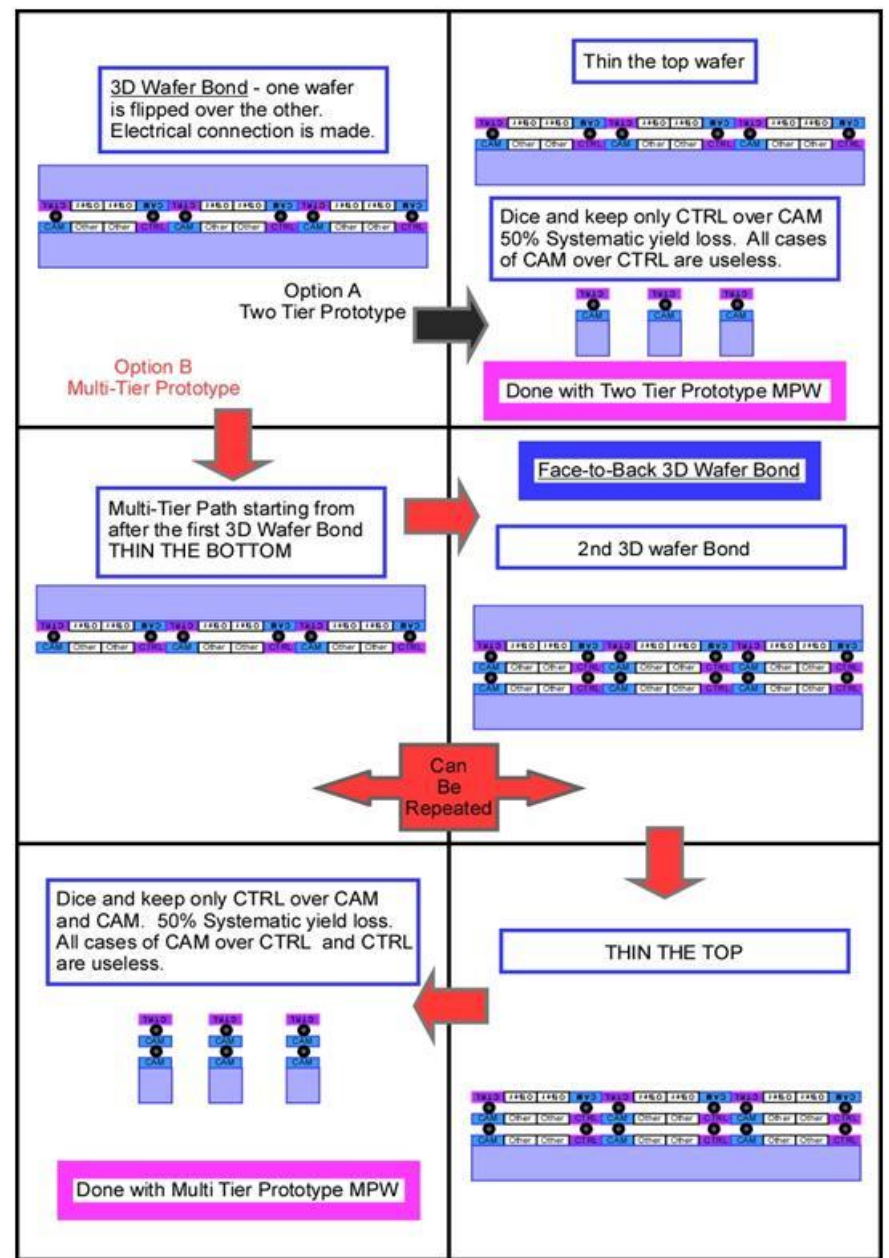

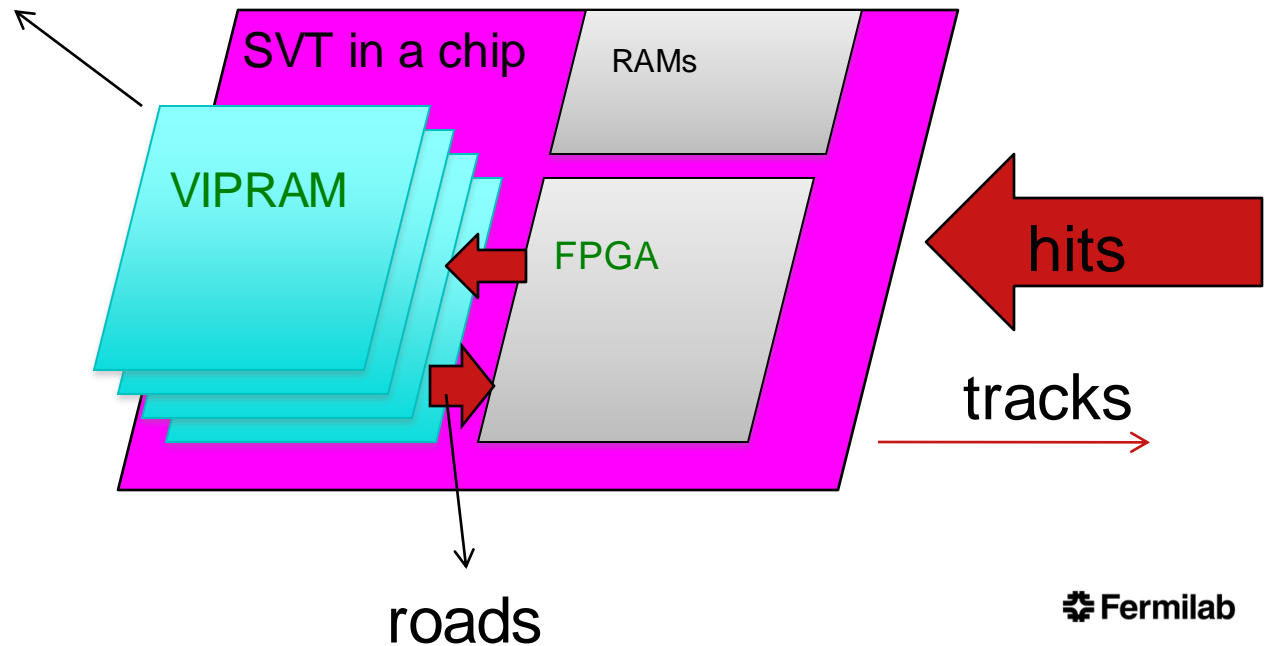


Figure 6 - The conclusion of a typical 3D MPW process OR an alternate process available to the VIPRAM.

Long term goal of R&D




- ~ 500K patterns/cm **2
- Running with > 100 MHz input rate
- N CAM tiers + Control tier
- integrated with FPGA/RAM
(*general purpose pattern recognition*)



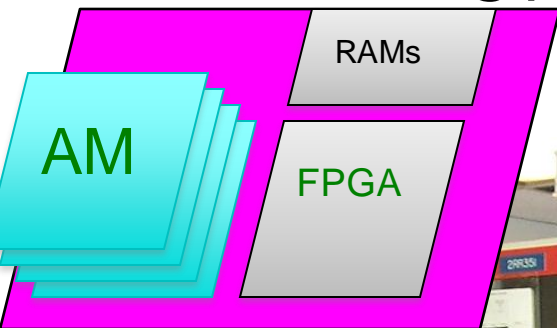
Integrate AM and TF/FPGA stages into one chip

- 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

SVT in one chip?: 2nd phase of VIPRAM project



2 meters





Original SVT system had 384K patterns total
Aim to reach ~500K per cm**2 for VIPRAM ...


Comments on future tracking trigger applications of VIPRAM

- VIPRAM project is “Generic R&D” at this stage, and current focus is the “proof-of-principle”: do our homework first
- It was motivated by FTK simulation studies and is based on existing AMchip general concept, as such,
 - Should be useful for future L2-like applications
 - “SVT in one chip” approach could simplify board and system design
- For L1 tracking trigger
 - The *ultimate goal* is to design it for L1 application
 - *VIPRAM architecture is inherently flexible & open (highly desirable)*
 - need to work out system level design vs chip level
 - need extensive simulation studies with physics cases for guidance
 - *Inputs and collaboration are welcome (CMS/ATLAS/others) ...*

Detector design for triggering



Tracking Trigger Issues (L1&L2)



Pattern Recognition

Associative Memory approach Others ...

Track Fitting

GPUs



HLT

Associative Memory approach is a proven technical for tracking trigger at hadron collider. The potential limitation is the (scaling of) performance of Associative Memory technology.

3D technology offers new design opportunities and VIPRAM is a promising way to go ...

Backup slides

- 3D basics
- VIPRAM 3D stacking requirements
- Two recent successful 3D R&D projects using the same technology:
 - <http://www.gtcad.gatech.edu/3d-maps/>
 - http://web.eecs.umich.edu/~dfick/files/fick_isscc2012_slides.pdf
- Diagonal Via technique

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

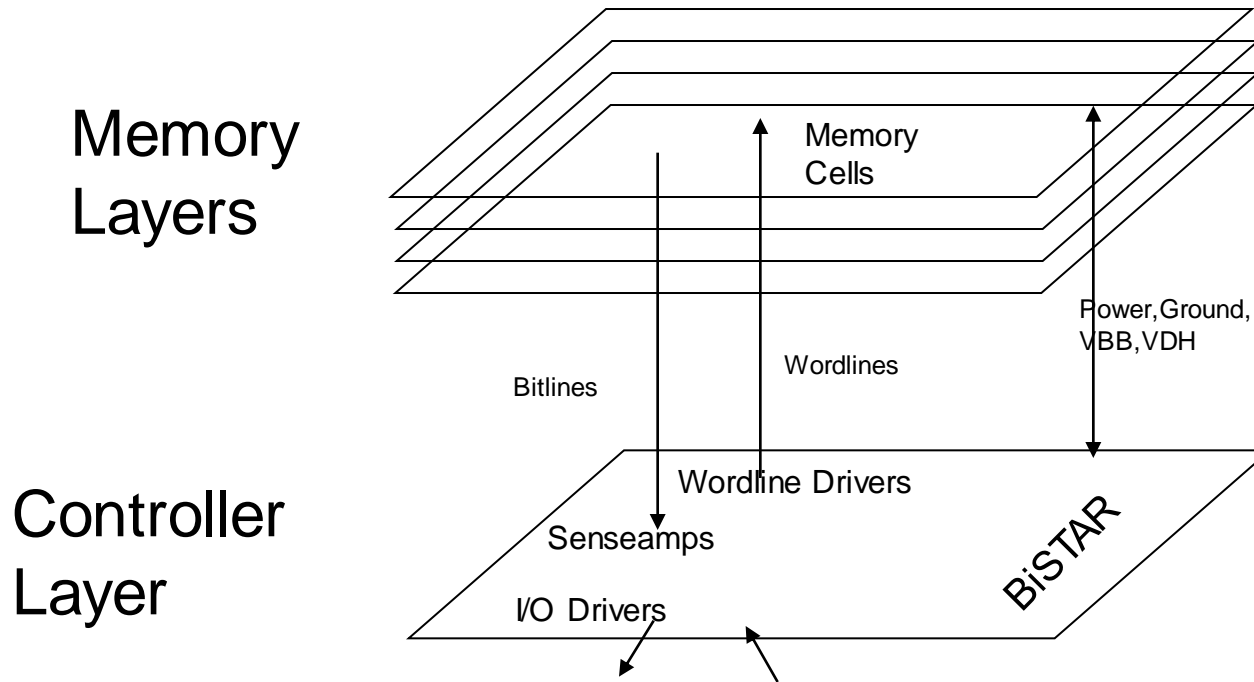
Examples of commercial applications of 3D Technology

- 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

“Dis-Integrated” 3D Memory



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



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

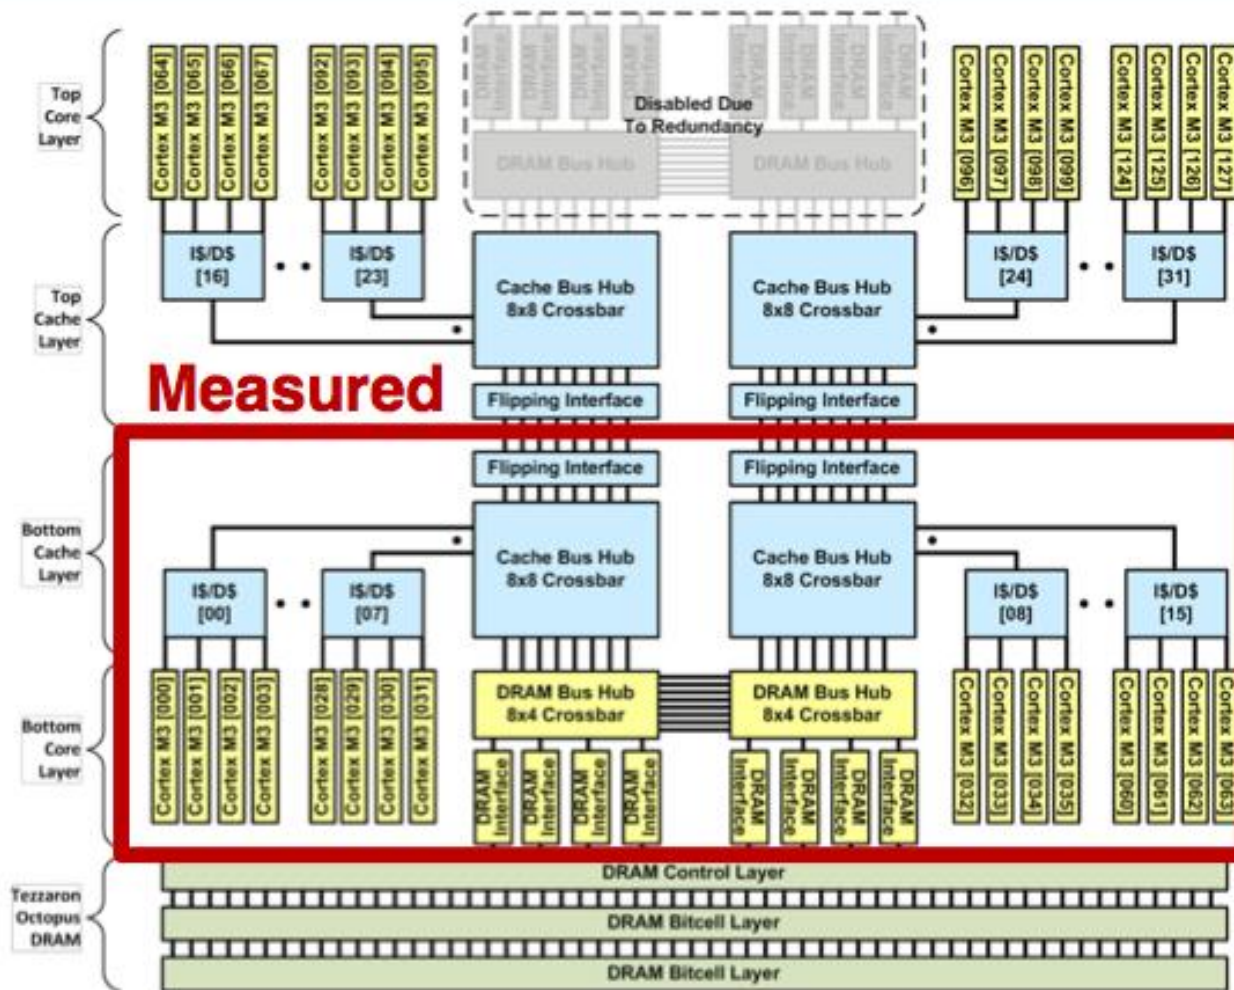


University of Michigan



Centip3De System Overview

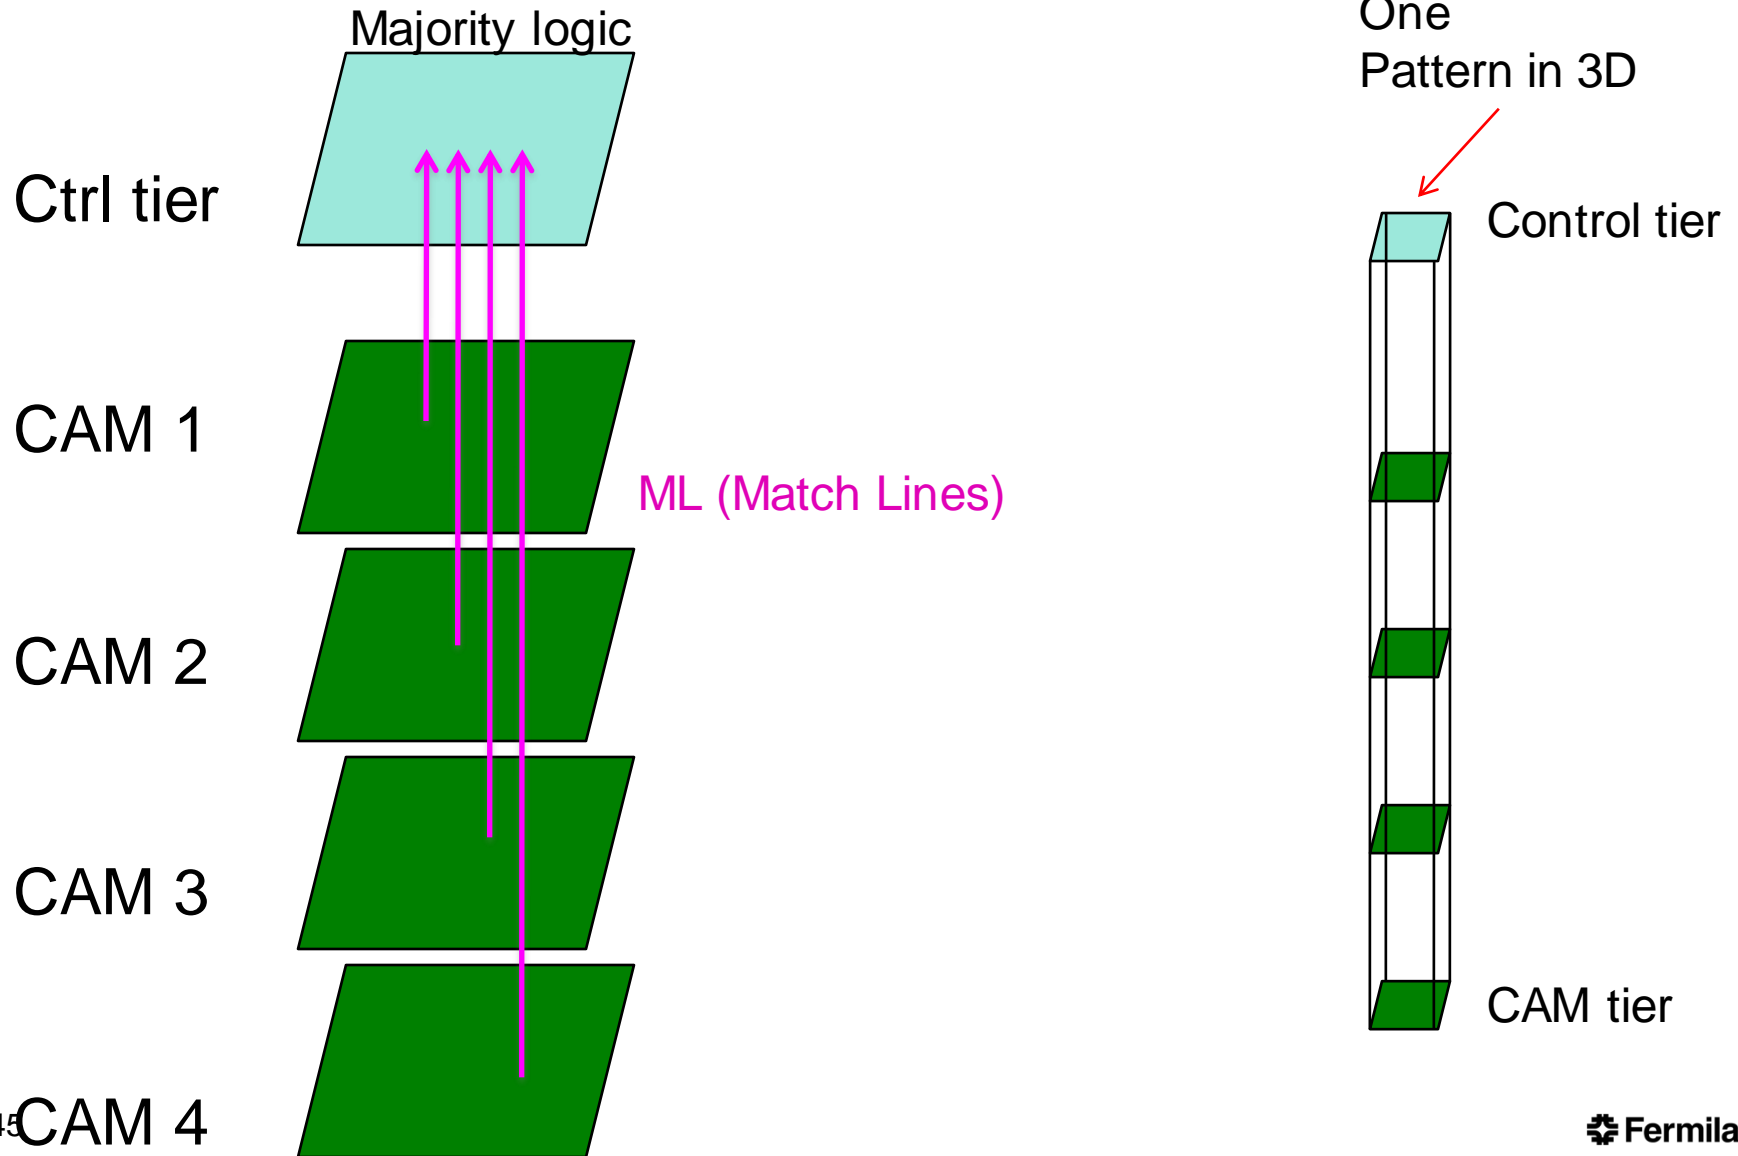
- 7-Layer NTC system
- 2-Layer system completed fabrication with measured results
- Full 7-layer system expected mid 2012




Diagonal via concept

One key issue for VIPRAM design:

How to communication between the control and each CAM tier, given that the CAM tiers are physically identical?



Offset stacking idea- is this feasible?



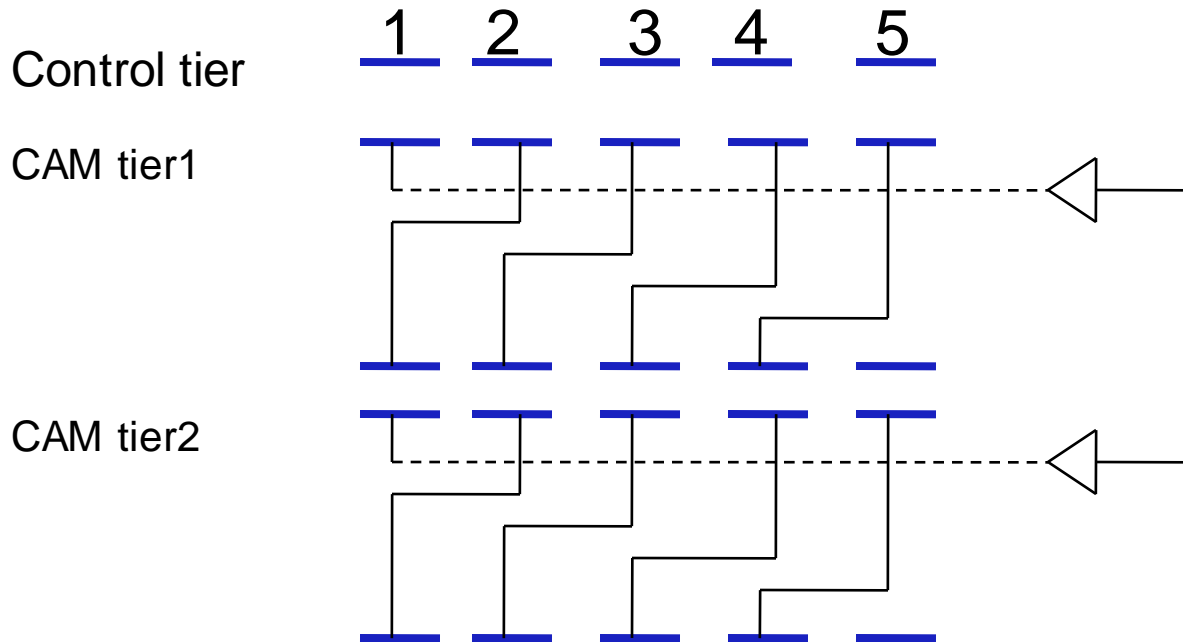
- The example for 4 identical CAM tiers, offset in one direction
- Every vertical connection has 3 extra connections on CAM tier
- Point to point communication done by offset (to/from Control)
- Power, clock etc lines have all 4 connected together

No extra transistor needed, pure geometrical solution.

⁴⁶But requires offset at wafer stacking stage...



Turns out Bob Patti at Tezzaron had a simpler idea to solve this problem long ago -- patented in 1999

- The idea was used for 3D DRAM stacking, to solve the same problem we are having, using “diagonal via”





- One example

- Diagonal via structure for 4 CAM tier case



The same can be done for all input and output signals,
No extra transistor is needed.
This trick solves the tier communication problem in a simple and clean way.

Price to pay:
a set of vias per signal
Number of vias = number of tiers



Diagonal via structure

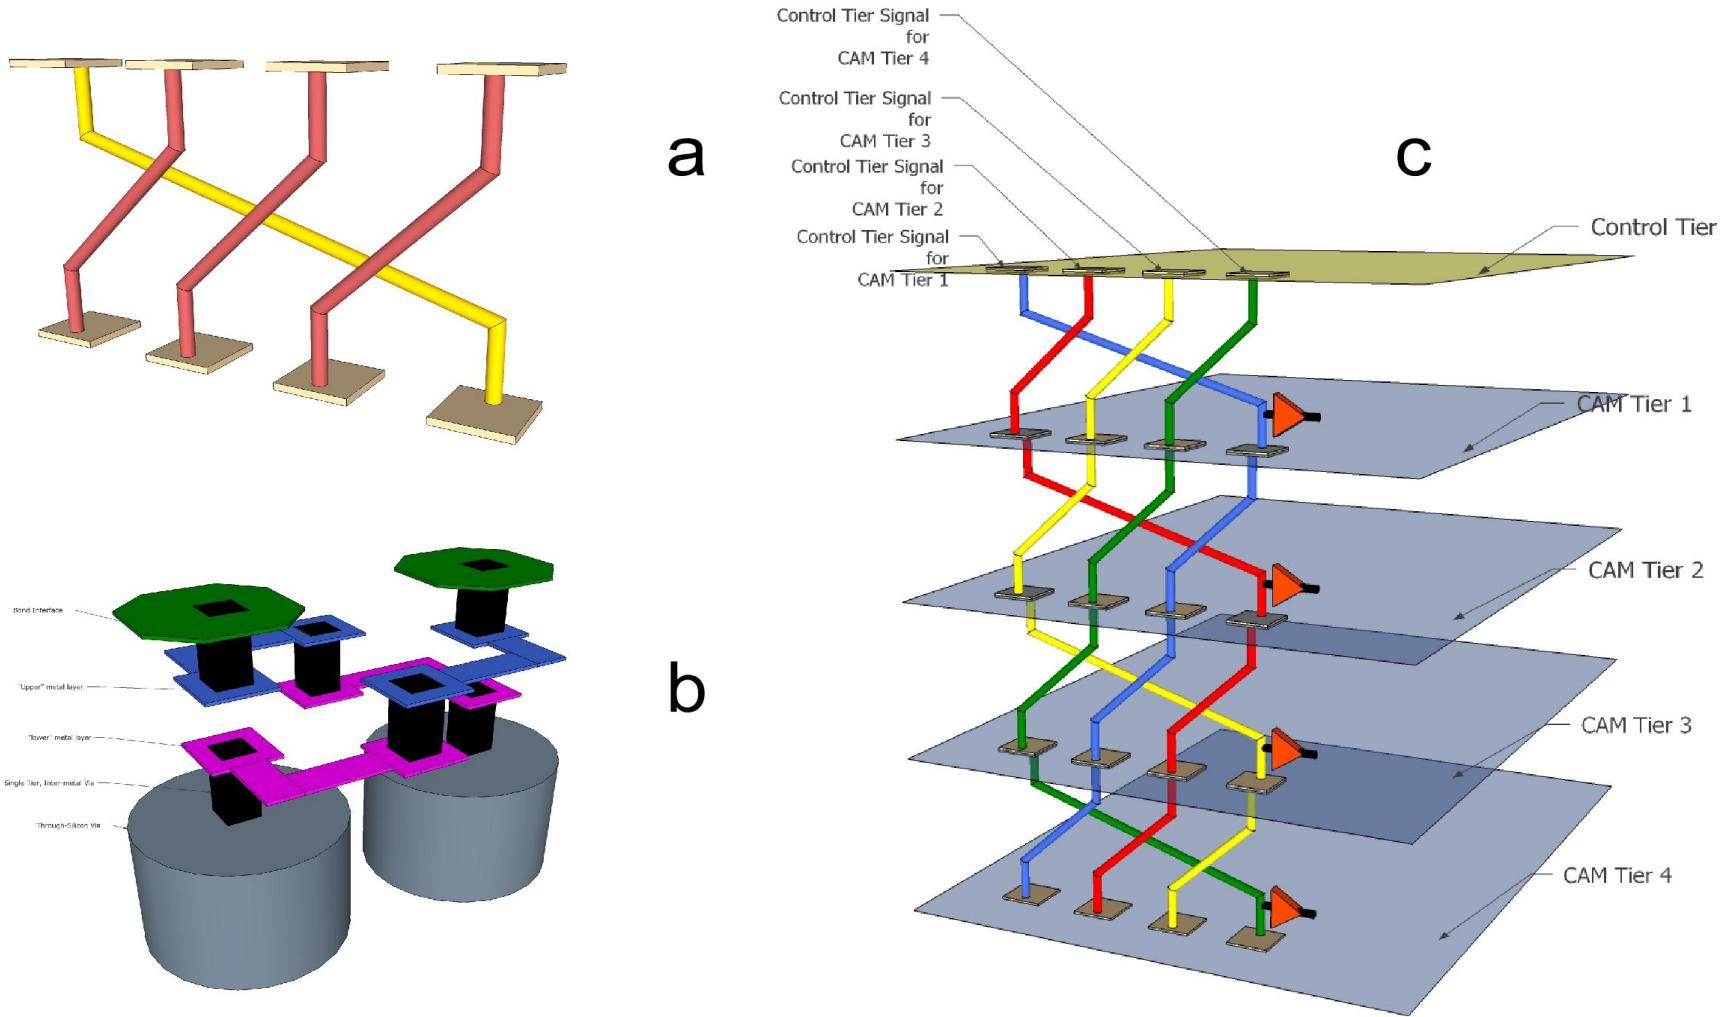


Fig. - Diagonal Vias: (a) shows a cartoon of the function of 4 diagonal vias; (b) shows a simple VLSI implementation of two diagonal vias from the Through-Silicon Vias on the bottom (in gray) up to one layer of metal (purple) to a second layer of metal (blue) and finally up to the bond interface (green) where it would connect to the next tier. (c) shows a cartoon of four 4-via diagonal vias working together to connect four different signals from a Control tier uniquely to four different CAM Tiers.