

Fermilab Preprint: FERMILAB-TM-2493-CMS-E-PPD-TD  
DOE proposal ID#: 124552

**Title of Proposal: Development of 3D Vertically Integrated Pattern Recognition  
Associative Memory (VIPRAM)**

**Collider Detector Research and Development Program (LAB 11-438)**

**Fermi National Accelerator Laboratory**

**Principal Investigators:**

**Fermilab: Tiehui Ted Liu  
Scientist II  
Fermi National Accelerator Laboratory  
PO Box 500  
Batavia, IL 60123  
Telephone: 630-840-6675  
Fax: 630-840-4610  
Email: thliu@fnal.gov**

**Univ. of Chicago: Mel Shochet  
Kersten Distinguished Service Professor of Physics  
HEP-211  
Enrico Fermi Institute  
University of Chicago  
5640 S. Ellis Ave.  
Chicago, IL 60637 USA  
Telephone: 773-702-7440  
Fax: 773-702-1914  
email: shochet@hep.uchicago.edu**

**Laboratory Official  
Pier Oddone  
Laboratory Director  
Telephone: 630-840-3211  
Fax: 630-840-2900  
Email: pjoddone@fnal.gov**

**Requested Funding:**

**Year 1: 182K  
Year 2: 150K  
Year 3: 150K**

**Use of human subjects in proposed project: No  
Use of vertebrate animals in proposed project: No**

**Collaborators:**

Fermilab:

Gregory Deptuch, Jim Hoff, Ron Lipton, Ted Liu, Jamieson Olsen,  
Erik Ramberg, Jin-Yuan Wu, Ray Yarema

University of Chicago:

Mel Shochet, Fukun Tang

ANL:

Marcel Demarteau, Gary Drake, Jinlong Zhang

INFN Padova

Silvia Amerio

Tezzaron Semiconductor Corporation

Bob Patti, Gretchen Patti

**(Version V 2.4, March 12th, 2011)**

## Table of Contents

<b>Table of Contents</b> .....	<b>3</b>
<b>Abstract</b> .....	<b>4</b>
<b>Project Narrative</b> .....	<b>5</b>
<b>Introduction: Future Challenges of Pattern Recognition</b> .....	<b>5</b>
<b>Fast Pattern Recognition and Track Reconstruction</b> .....	<b>6</b>
<b>The unique requirement for AM for HEP application</b> .....	<b>8</b>
<b>Implementation of PRAM in 2D</b> .....	<b>9</b>
<b>Proposed Solution Based on a Novel Technology: from 2D to 3D</b> .....	<b>10</b>
<b>VIPRAM: Vertically Integrated Pattern Recognition Associative Memory</b> .....	<b>11</b>
<b>The collaboration</b> .....	<b>13</b>
<b>The VIPRAM 3D Stacking Process: MPW Prototyping</b> .....	<b>15</b>
<b>Prototyping Costs</b> .....	<b>17</b>
<b>Dividing the Labor among Collaborators</b> .....	<b>17</b>
<b>Costs and Timeline</b> .....	<b>19</b>
<b>Future Plan: Phase II</b> .....	<b>20</b>
<b>VIPRAM – Frequently Asked Questions</b> .....	<b>22</b>
<b>1. Why not simply use modern FPGAs to implement a Pattern Recognition Associative Memory?</b> .....	<b>22</b>
<b>2. Why not use the commercially available CAMs, such as in a network search engine?</b> .....	<b>22</b>
<b>3. How can the power consumption of AM be reduced significantly?</b> .....	<b>22</b>
<b>4. Why not simply increase the pattern density by going to 65 nm (or beyond) in 2D?</b> .....	<b>23</b>
<b>5. Is the 3D process for the VIPRAM really similar to the Tezzaron's 3D DRAM process?</b> .....	<b>23</b>
<b>6. What about the yield issues of the 3D process and how to address them?</b> .....	<b>23</b>
<b>7. What are the main architecture advantages of the 3D VIPRAM approach?</b> .....	<b>23</b>
<b>Appendices</b> .....	<b>25</b>
<b>“Diagonal Vias” approach for inter-tier communications</b> .....	<b>25</b>
<b>Diagonal Vias in Greater Detail</b> .....	<b>27</b>
<b>Majority Logic – New versus Old</b> .....	<b>27</b>
<b>The Fischer Tree for Readout</b> .....	<b>30</b>
<b>How CAM Works and 3D Advantages for PRAM</b> .....	<b>32</b>
<b>Tezzaron’s FaStack® Technology (see Tezzaron’s web site)</b> .....	<b>34</b>
<b>Expected Areas of Improvement in the True 3D Architecture</b> .....	<b>36</b>
<b>References</b> .....	<b>37</b>
<b>Principle Investigators: Curriculum Vitae</b> .....	<b>39</b>

## Abstract

Many next-generation physics experiments will be characterized by the collection of large quantities of data, taken in rapid succession, from which scientists will have to unravel the underlying physical processes. In most cases, large backgrounds will overwhelm the physics signal. Since the quantity of data that can be stored for later analysis is limited, real-time event selection is imperative to retain the interesting events while rejecting the background. Scaling of current technologies is unlikely to satisfy the scientific needs of future projects, so investments in transformational new technologies need to be made. For example, future particle physics experiments looking for rare processes will have to address the demanding challenges of fast pattern recognition in triggering as detector hit density becomes significantly higher due to the high luminosity required to produce the rare processes. In this proposal, we intend to develop hardware-based technology that significantly advances the state-of-the-art for fast pattern recognition within and outside HEP using the 3D vertical integration technology that has emerged recently in industry.

The ultimate physics reach of the LHC experiments will crucially depend on the tracking trigger's ability to help discriminate between interesting rare events and the background. Hardware-based pattern recognition for fast triggering on particle tracks has been successfully used in high-energy physics experiments for some time. The CDF Silicon Vertex Trigger (SVT) at the Fermilab Tevatron is an excellent example. The method used there, developed in the 1990's, is based on algorithms that use a massively parallel associative memory architecture to identify patterns efficiently at high speed. However, due to much higher occupancy and event rates at the LHC, and the fact that the LHC detectors have a much larger number of channels in their tracking detectors, there is an enormous challenge in implementing pattern recognition for a track trigger, requiring about three orders of magnitude more associative memory patterns than what was used in the original CDF SVT. Significant improvement in the architecture of associative memory structures is needed to run fast pattern recognition algorithms of this scale.

We are proposing the development of 3D integrated circuit technology as a way to implement new associative memory structures for fast pattern recognition applications. Adding a "third" dimension to the signal processing chain, as compared to the two-dimensional nature of printed circuit boards, Field Programmable Gate Arrays (FPGAs), etc., opens up the possibility for new architectures that could dramatically enhance pattern recognition capability. We are currently performing preliminary design work to demonstrate the feasibility of this approach. In this proposal, we seek to develop the design and perform the ASIC engineering necessary to realize a prototype device.

While our focus here is on the Energy Frontier (e.g. the LHC), the approach may have applications in experiments in the Intensity Frontier and the Cosmic Frontier as well as other scientific and medical projects. In fact, the technique that we are proposing is very generic and could have wide applications far beyond track trigger, both within and outside HEP.

## Project Narrative

### Introduction: Future Challenges of Pattern Recognition

Many next generation science experiments will be characterized by the collection of large amounts of data, taken in rapid succession, from which the scientists will have to unravel the underlying physics processes. More often than not, large backgrounds will overwhelm the physics signal and real-time data analysis will be indispensable to immediately separate interesting events from background, select them for further analysis and reduce the data size to manageable proportions. Scaling of current technologies does not seem to meet the scientific goals of future projects and investments in transformational new technologies need to be made to enable new scientific projects.

Many areas in science can be identified that currently face these challenges. One is the capability to perform fast pattern recognition and track reconstruction of particle trajectories in modern High-Energy Physics (HEP) hadron collider experiments. The Large Hadron Collider (LHC) at CERN has proposed a luminosity increase of a factor of five to ten over the original design as the goal for the upgrade, which will result in a corresponding increase in particle interactions and track densities in the detector. Most of these interactions contain events that are of no significance and should not be recorded. The ultimate physics reach of the LHC will crucially depend on the tracking trigger capabilities of its experiments to handle these high luminosities and discriminate between the interesting events and the background. The overall goal is to identify particle tracks at the trigger level, a capability that is crucial for many important searches for new physics. The CMS muon trigger, for example, will reach an unacceptably large rate at high luminosity due to the number of hits in the muon detectors. The first-level trigger can be reduced to an acceptable level if tracks are found in the inner detector and matched to the muon candidates. There are other important reasons for having tracking trigger capabilities at early stages of the trigger system. For example, the online identification of heavy fermions such as  $b$  quarks and tau leptons are important, since many interesting channels of new phenomena produce heavier elementary particles. Tracks coming from a secondary vertex not in the direction of the beam line identify a  $b$  quark. Tau jets can be separated from background using the number of tracks within a narrow “signal cone” and the number in a larger “isolation region”.

Another example of tracking needs at the high-energy frontier is pattern recognition at a Muon Collider (MuC). The hit densities in a vertex and tracking detector at a MuC are dominated by backgrounds from the decays in-flight of the upstream muons and upstream muons entering the detector. These upstream muons will not originate from the interaction point, but rather travel along the beam axis. Fast, efficient pattern recognition could identify and eliminate these tracks online. Besides the fast tracking trigger applications for the energy frontier experiments, the proposed R&D might be useful for other experiments in the future as well. Possible examples include intensity frontier experiments such as  $\mu 2e$ , and cosmic frontier experiments such as ground-based telescope arrays where fast triggering on the correlation of images from multiple telescopes into the sky is needed. Another possible example is to use the associative memory to correlate images in time, to detect any changes in time for the same detector

that rapidly scans the sky (such as the detection of Supernova in real time with large CCD cameras, such as LSST, the Large Synoptic Survey Telescope).

Instrumentation at photon science facilities could also benefit from the development of technologies that provide fast online pattern recognition of large sets of data. In Photon Correlation Spectroscopy (PCS), for example, the dynamics of a material are probed by analyzing the temporal correlations among photons scattered by the material. X-ray PCS (XPCS) offers the unprecedented opportunity to extend the range of length scales over which a material's low frequency dynamics can be probed down to inter-atomic spacing. With the advent of new coherent, brilliant X-ray sources, technologies such as the one proposed here, enabling online correlation spectroscopy could enhance the facility's scientific reach.

There are also potential medical applications. For example, in proton computed tomography (pCT), data taking rates and background contamination are serious limitations to reducing patient exposure time. Associative memory could provide the rapid pattern recognition needed to address both of these problems.

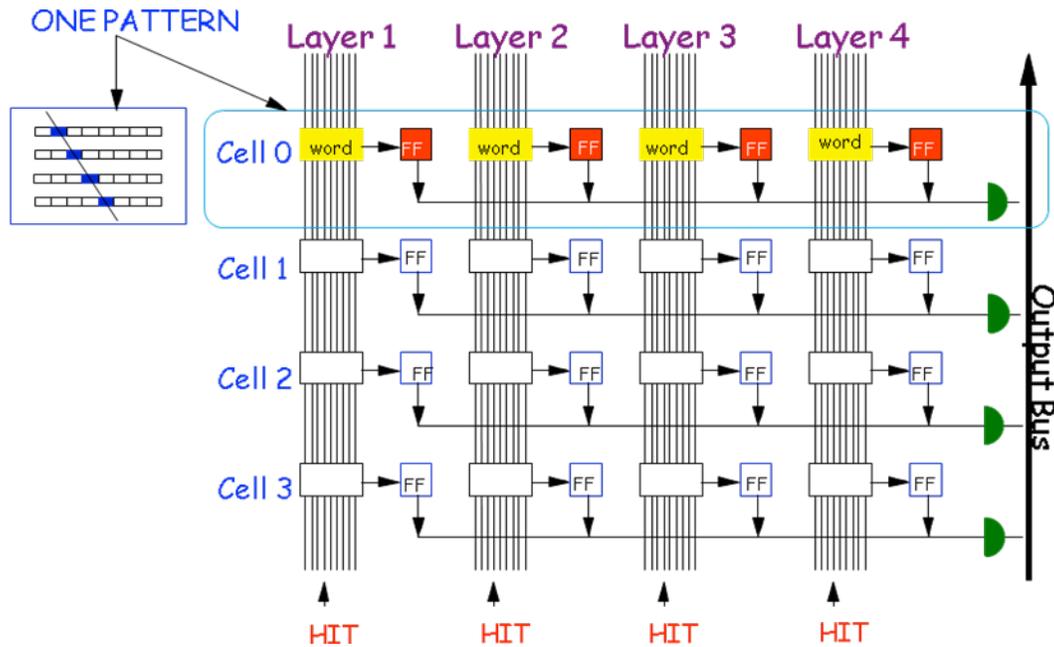


Figure 1. CDF SVT Associative Memory Architecture [1].

### Fast Pattern Recognition and Track Reconstruction

Traditionally, track triggers have been implemented using computational techniques to identify patterns and perform track fitting, often using processors running in the upper levels of a data acquisition system to perform the task. However, such algorithms are relatively slow, because of the serial nature of CPU processing time. It is also desirable to push this type of trigger into earlier levels of a trigger system. The CDF Silicon Vertex Trigger (SVT) at the Fermilab Tevatron is a great example. The method used there [1], developed in the 1990s, uses algorithms implemented in fast logic. The technique has two parts. The first part uses Associative Memory (AM) architectures

based on Content Addressable Memory (CAM) cells [2] to efficiently identify track patterns (roads) at high speed using coarse-resolution “hits” recorded in the tracking detector. Then, the patterns are processed using fast FPGAs to perform track fitting with full detector resolution hits. A block diagram of the Associative Memory architecture is shown in Figure 1. The method solves the combinatorial challenge inherent to the track finding by exploiting massive parallelism of associative memories that can compare tracking detector hits to a set of pre-calculated patterns simultaneously. Because each pattern is narrow, the usual helical fit can be replaced by a simple linear calculation. The track fitting stage for each matched pattern is much simplified and fast using tracking parameters with values for the center of the road, and applying corrections that are linear in the hit position in each layer.

The SVT approach was highly successful, and CDF was the first hadron collider experiment in HEP to incorporate a fast secondary vertex track trigger [3][4]. It finds all tracks emanating from each collision and precisely measures their properties within about 30 microseconds after the collision. This latency can be compared to the  $\sim 1$  second required when track reconstruction is done inside a modern computer. The SVT has been essential to many of the physics results to come out of the CDF experiment and it has significantly improved the CDF physics reach. For example, it is the critical device that enabled CDF to precisely measure the frequency of the long awaited Bs mixing [5], a process that is important for understanding the matter-antimatter asymmetry in nature. It also allowed the observation of the decay of the Z boson [6], a carrier of the weak nuclear force, into two energetic b-quark jets, a signature very similar to that of the Higgs boson decay.

In the era of the upgraded LHC (SLHC), it is desirable to implement this type of track finding capability in the early stages of the trigger because of the importance of identifying particle tracks at the trigger level in many searches for new physics. However, due to the much higher occupancy and event rates at the SLHC, and the fact that the LHC detectors are much more massive with orders of magnitude more channels in their tracking volumes, it is a difficult challenge to perform pattern recognition and track fitting at the trigger level. In addition there is the obvious challenge of data transfer from the detector to the trigger system. While processing power has increased steadily over time, the demands in fast track reconstruction have increased even faster because of the rapid rise in detector hit combinatorics as luminosity increases. Significant improvements in both pattern recognition (the associative memory) as well as track fitting performance are needed.

A critical figure of merit for an AM-based track reconstruction system is the number of predetermined track patterns or roads that can be stored in the Associative Memory bank. Generally speaking, wider roads using coarser resolution hits require less AM storage, but the number of AM roads satisfied by random hits and the number of fits at the track fitting stage downstream increases quickly due to the high detector occupancy. Also, the demand on the bandwidth would be higher because all the roads and hits have to be transferred from the AM stage to the track fitting stage. If the roads are very narrow, due to using finer resolution hits, the number of fake roads and fits are reduced, but the required total size of the AM increases dramatically. Therefore, the road width must be optimized. The required AM pattern bank size will be different for

different experiments and different for the same experiment at different luminosities. As an example, consider the implementation of a hardware-based track trigger like that used in the CDF SVT in the context of what is needed for the LHC. For this comparison, we will use the Atlas FastTrack (FTK) [7] project as an example, since the design requirements of the system are well known from extensive simulations, although this high-level extrapolation could apply to CMS as well. The original CDF SVT system, in operation during 2000 until 2005, had a total number of associative memory patterns of 384,000, while the proposed ATLAS FTK system for the Level 2 trigger will require  $\sim 1$  billion patterns in order to handle a luminosity of  $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$  [1]. This is three orders of magnitude more associative memory patterns. The Level 1 Track Trigger upgrade for both CMS and ATLAS would likely also require large number of AM patterns running at high speed for tracking trigger processing.

An upgraded AM chip (AMchip03) was developed in 2005 by the CDF Italian collaborators [8]. The AMchip03 was implemented in 180nm technology using the Standard Cell approach and the number of patterns per chip increased by a factor of 40 over the previous version used at CDF, from 128 to 5,000. This chip, which runs at 40MHz, was used to upgrade the SVT system, where the total number of patterns has increased to more than 6 Million. However, in order to meet the challenges for the LHC high luminosity running, another increase in pattern density by two orders of magnitude will be required. The AMchip pattern density can be improved by optimizing the design in single-layer chips (2D), using custom cell designs with smaller feature size technology. There is an R&D effort (Italy/FNAL/Germany) using 65 nm technology to improve the standard cell based AMchip03 design [9]. Some of this work is described in the Appendix and is very important for the near term performance improvement. However, due to the limitations in technology scaling, the gain is rather limited and may not be sufficient for applications at higher luminosity for LHC in the future.

The current technology using FPGAs and custom 2D chips cannot be scaled in a simple manner to satisfy future needs. Significant improvement in associative memory performance (pattern density and speed) is needed. The solution we are pursuing is to add a “third” dimension to the signal processing chain by developing multi-tier chips.

### **The unique requirement for AM for HEP application**

As mentioned earlier, the CDF SVT Associative Memory chip (from now on we will call it PRAM, Pattern Recognition Associative Memory, to emphasis its purpose for HEP) is a departure beyond conventional CAMs. Like conventional CAMs, PRAMs store address patterns and look for matches between incoming hits and those addresses for a given detector layer. At this level, the match is expected to be either exact (Binary CAM) or partial (Ternary CAM) and an array of Match Flags is the typical output. A PRAM has an array of Match Flag Latches which capture and hold the results of the match until reset for the next event. As the hits from the various layers of the detector for the same event arrive, the PRAM is looking for more than simple matches from one candidate address to one or more stored address patterns. The PRAM organizes stored address patterns into roads, which are linked arrays of several stored address patterns from different detector layers. Each stored address pattern in a road is from a different layer in the detector system and these linked arrays represent a path or road that a particle



Figure 2. A typical implementation of PRAM in 2D, shown for six detector layers.

might traverse through the layers of the detector (hence the name “road”). The ultimate goal of the PRAM is to match real particle trajectories to those roads. Like a conventional CAM, a PRAM flags a match when a candidate address matches a stored pattern address for a given detector layer. However, before the PRAM does anything with that match, it must find matches in all the elements (layers) that constitute a road.

It should be emphasized that compared to commercially available CAMs, such as Network Search Engine, the CDF AMchip has the unique ability to search for correlations among input words received on different clock cycles. This is essential for tracking trigger applications since the input words are the detector hits arriving from different layers at different times. They arrive at the chip without any specific timing correlation. Each pattern has to store each fired layer until the pattern is matched or the event is fully processed and thus all patterns can be reset. Even in the case of a level-1 trigger application, which is largely synchronous, this feature will still be important.

### Implementation of PRAM in 2D

In a 2D implementation, a PRAM can be thought of as an array of classic CAMs laid out side-by-side column-wise with an extra set of Road Glue Logic connecting each row. This is shown in Figure 2.

Each green column in Figure 2 is a set of classic CAM word cells dedicated to one particular detector layer; therefore, in this figure, six independent detector layers are serviced. Each row in Figure 2 is one complete set of related patterns. In a tracking detector in a high energy physics experiment, this would be one, pre-defined charged particle track. In other words, a charged particle traverses each layer of the detector and its passage is recorded on each detector layer. If those recorded addresses match the

addresses stored in each CAM word cell in a row in Figure 2, then the Glue Logic (in blue) will activate, indicating that a pattern match has been found. This way, each row has all the circuitry necessary for one complete road detection. The details of how CAM word cell and Majority Logic cell work are described in the Appendix.

One of the limits of the 2D implementation arises when a large number of CAM bits and large number of detector layers are required (AMchip03 has about 16 bits per layer and total of 6 layers, but the new 2D chip for FTK will require 8 layers): some of the match lines from CAM word cells to the Road Glue Logic or Majority Logic will be long and have large parasitic capacitance and resistance, and this would contribute to power consumption. In addition, the common control, interface and readout logic will need to be implemented in 2D. Thus, the routing of the 2D implementation is quite involved. Moreover, when more detector layers are required, the pattern density will be reduced accordingly due to the fact that more columns of CAM word cells will be required. As will be seen next, in 3D, the situation will be very different, making it possible to increase the pattern density significantly, independent of the number of detector layers needed.

### Proposed Solution Based on a Novel Technology: from 2D to 3D

In this proposal, we describe the use of 3-dimensional (3D) integrated circuit technology as a way to improve the fast pattern recognition and track reconstruction for HEP trigger applications.

3D technology is the integration of thinned and bonded silicon integrated circuits with vertical interconnects between IC layers using Through Silicon Vias (TSVs). The technology has wide applications in industry, ranging from memories to pixel arrays to microprocessors and FPGAs. Performance can be improved significantly by reducing interconnect R/L/C for higher speed and density. In addition, it provides the freedom to divide functionality among different layers or tiers to create new designs that are simply not possible in 2D. As Moore's law is approaching severe limitations, it is expected that 3D technology will be the next scaling engine. It is worth pointing out here that 3D technology is not just used as a mere extension of Moore's law, it actually offers novel design opportunities which are simply not possible in 2D. As will be described next, this is certainly the case for this proposal for the VIPRAM design.

Generally speaking, 3D technology becomes useful when a task can be partitioned into multiple sections that are physically and logically separable, and the interconnections among them are straightforward. Moreover, the use of 3D technology can have varied goals. For example, it can be used to increase transistor density, i.e. to increase the number of transistors per square micron. Such is a major goal of 3D DRAM design. Here, the DRAM task is first logically divided into a control/interface section and memory core. The control/interface section is physically separated onto its own tier, and the memory core is further divided into memory banks that are each implemented on their own tiers. A second, different example is the 3D integration of microprocessor systems. Here, different functions that have been traditionally separated can be brought together in a single monolithic structure and technological limitations can be eliminated. CPU and memory can be placed on separate tiers and the interconnection between them, i.e. the memory bus, can be reduced from on the order of tens of millimeters (a bus on a PC

board) to a few tens of microns (the length of a through silicon via). Also, the memory bus itself can be expanded from a few bits to hundreds of bits wide, dramatically improving the memory access bandwidth.

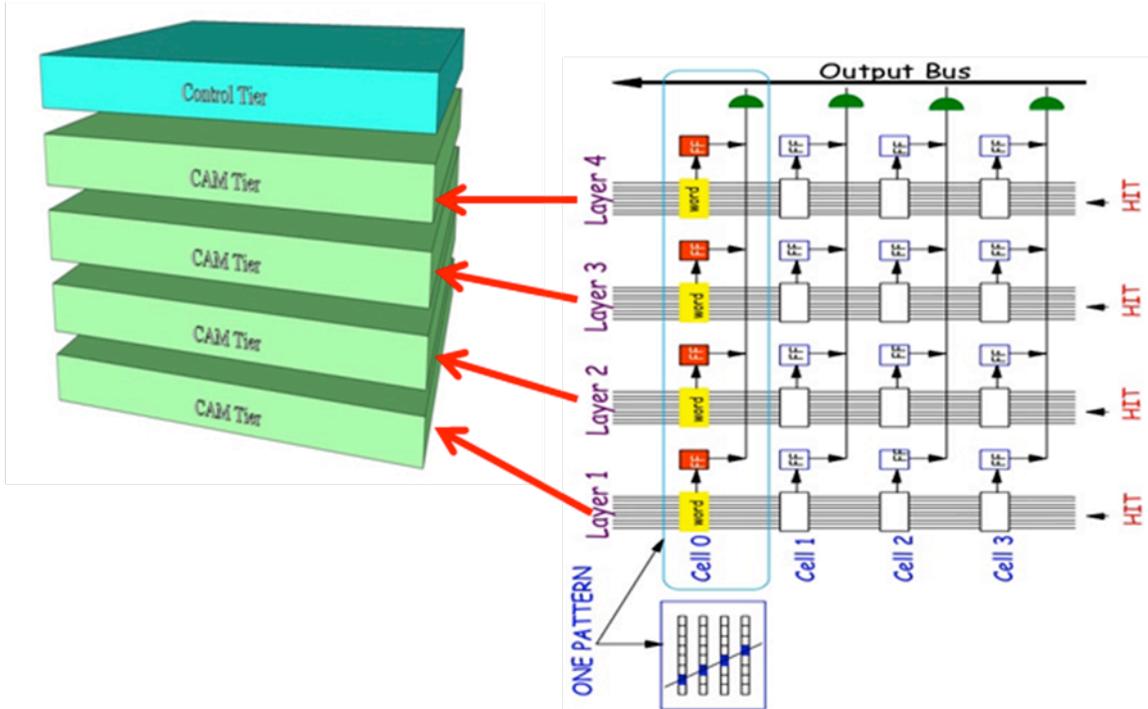


Figure 3. Associative memory architecture in 3D

**VIPRAM: Vertically Integrated  
Pattern Recognition  
Associative Memory**

Like CAMs, the PRAM's basic operation is to compare candidate addresses to stored address patterns. Unlike CAMs, a PRAM does not flag matches until it has matched candidate addresses from multiple sources (tiers representing layers) to an array or road of stored address patterns. This gives rise to a very natural 3D progression. Perhaps one way to see this more clearly is to view the Associative Memory architecture in Figure 1 in a different way: rotated by 90 degrees, as shown in

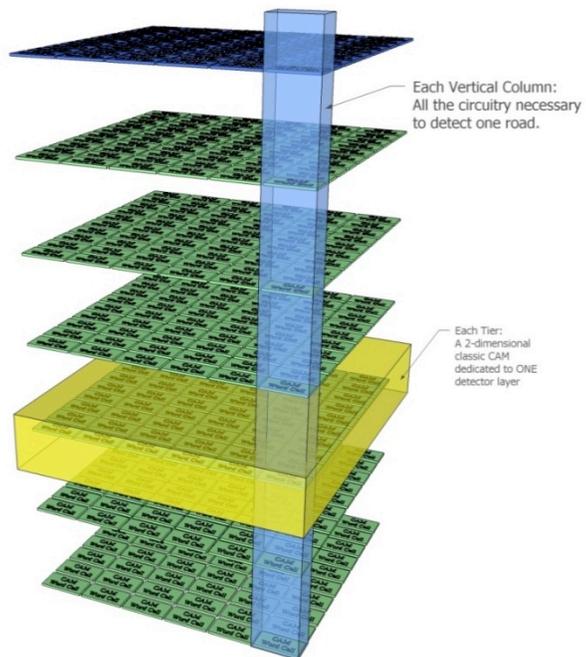


Figure 4. A PRAM in 3D where a vertical blue tube represents one independent road

Figure 3. The idea is to have a dedicated CAM tier for each detector layer, where the incoming hits are matched to the stored hit locations in an array of CAM word cells, and have one control/interface tier to collect and associate the hit matching information from each CAM word cell for a given CAM tier or detector layer. A significant improvement can be obtained when going in a third dimension. There is no need to spread the different patterns out horizontally; instead, stack them vertically with one layer per tier. There is no need to complicate and bloat the horizontal routing of signals; instead, route the individual match lines vertically using TSV to a Road Glue Logic cell on the top tier. The resulting architecture, i.e. the design of VIPRAM, is shown in Figure 4, where the vertical blue tube represents one independent pattern or road. The top tier is the control tier which has one Road Glue Logic or Majority Logic cell for each vertical blue tube or pattern, collecting match line signals from each CAM word cell on each CAM tier (or detector layer) and performing Majority Logic tasks to determine whether a road is fired for a given event.

Each tier now resembles a classic CAM word cell array in two dimensions. Each tier accommodates only one detector layer of data. Candidate addresses (i.e. hits) of a particular detector layer are driven along the bit lines of one tier and one tier only. Hit patterns are stored on the tier that corresponds to their detector layer. Each match line from a given layer for a given CAM word cell is driven vertically and directly into the Control tier. The details on how match lines are driven vertically from the identical CAM tiers to the control tier, as well as how input data bits are driven directly to their corresponding CAM tier is described in the Appendix, see the “Diagonal Vias”. In a 2D footprint the size of a single CAM word cell, all of the road detection circuitry can be implemented over a few tiers. This means that in an area that once contained only one CAM word cell, a 3D PRAM can process  $L$  layers of a road pattern, where  $L$  is the number of detector layers<sup>1</sup>. This is one of the main architectural advantages for HEP PRAM using 3D technology.

With Pattern Recognition Associative Memory for HEP tracking trigger applications, the task is logically and physically dividable and the interconnections are straightforward, making it almost an ideal candidate for 3D integration. Flagging an individual detector hit is not important. Rather, the path of a charged particle through many detector layers is what must be found. This effectively makes an HEP Associative Memory a “CAM of CAMs”, meaning that individual hits must be flagged and accumulated first for a given event, then related sets of detector hits – commonly called “roads” – must ultimately be flagged. Therefore, in its essence, an HEP Associative Memory (AM) bank is a CAM array that is a collection of independent roads, i.e. independent sets of hit addresses from different detector layers that represent a path or road that a charged particle might traverse through the detector. Like the 3D DRAM case, the AM can naturally be divided into a control/interface tier and a set of CAM tiers. What is unique for the 3D AM (VIPRAM) architecture is that when each CAM tier corresponds to a single detector layer, then the interconnections between the tiers become dramatically simplified. Logically, an AM road is an independent set of hit addresses

---

<sup>1</sup> Actually, this is conservative. “R” road patterns in a 2D PRAM took up the area of “ $R \times L$ ” CAM words plus the area of  $R$  Road Glue Logic cells plus extra area for routing. Therefore, we should fit *more* than  $R \times L$  roads in the same area occupied by  $R$  roads in a 2D PRAM.

from different detector layers; now, physically in the VIPRAM architecture, a road is a simple independent vertical tube in a 3D monolithic circuit that is a collection of CAM cells each programmed to detect the hit on a particular detector layer for that particular road and report the match directly to the control tier. Routing in 3D can be very efficient, especially if functional elements are arranged such that the interconnections among tiers are mostly vertical. This is the case for VIPRAM architecture: not only are the interconnects among tiers vertical, they are uniform (in fact identical) across the tiers as well.

In this proposal, we seek to first develop the 3D design of the VIPRAM, and then perform the ASIC engineering necessary to realize a prototype device of the 3D stacked VIPRAM as our near term R&D program over three years. This will be done initially in 130 nm as prototyping (with the initial goal of increasing the pattern density by a factor of 40 over that of the AMchip03). The final chip can be done in 65 nm (with > 100 x AMchip03 in pattern density expected).

One of the main challenges of the 3D stacking and integration approach will be power and thermal issues [10]. There has been a lot of work done in reducing the power consumption in the new 2D ongoing AMchip04, and the 3D design will benefit directly from that effort. Power dissipation in CAM is dominated by the dynamic power that is consumed by the match-line (ML) and search-line (SL) toggling during each clock cycle for search and match operations. The search-lines are switching to represent the new words to be compared and as a result match-lines are continuously switching based on the miss/match results. The AMchip04 design reduces power consumption in a significant way by using the pre-match power-saving technique [11].

In addition, since we plan to follow Tezzaron's 3D DRAM stacking approach [12], we will learn a great deal from Tezzaron's extensive experience in addressing power and thermal issues in 3D stacking.

### The collaboration

The proposed R&D would be carried out as a collaborative effort among Fermilab, Argonne, University of Chicago, INFN Padova in Italy, and Tezzaron. The proposed work is fully consistent with DOE's priorities for national labs: focus on transformational science, connect basic and applied sciences, re-energize the national labs as centers of great science and innovation, embrace a degree of risk-taking in research, and create an effective mechanism to integrate national laboratory, university and industry activities.

Some of the physicists in this collaboration have been involved in the design, building, commissioning, operation and upgrade of the CDF SVT system, as well as the current design work of the FTK system. Fermilab also collaborates closely with INFN Pisa and Frascati in Italy on the 2D development of AMchip04 [18] in 65 nm. Fermilab also contributes to the new Majority Logic design as well as the pattern readout algorithm using Fisher Tree approach in the control tier. The new Majority Logic and readout algorithm as well as the CAM word cell, currently being developed for the 2D AMchip04, will be directly useful for the 3D VIPRAM design. The extensive experience in associative memory and track fitting within the collaboration will be important for carrying out this R&D project. In addition, this proposal will leverage unique areas of

engineering expertise at Fermilab.

The 3D integrated circuit technology is actively being pursued by industry, since it enables heterogeneous integration of IC technologies, dense packing of transistors, and close integration of sensors and electronics. Partnering with an experienced industrial partner is key to the success of this project. Our partner in this R&D project is the company Tezzaron Semiconductor, located in Naperville next to Fermilab, Illinois. Tezzaron is one of the world-leaders in developing the 3D technology and specializes in cutting-edge memory products, 3D wafer stacking and TSV processes. Tezzaron's revolutionary FaStack® technology, which integrates several layers of DRAM with a powerful controller layer, will be used for the VIPRAM R&D work.

Fermilab was the first high-energy physics laboratory to recognize the potential of 3D integrated circuits for particle physics. It has started a focused R&D program to explore this technology, and is currently recognized as the world leader in exploring this technology for high-energy physics applications. In addition, Fermilab has already been developing a 3D chip (VICTR) to demonstrate the application of 3D technology to the formation of track-trigger primitives for the CMS level-1 tracking trigger upgrade. The proposed 3D fast pattern recognition and track fitting R&D would leverage this other work in 3D technology that has already begun. Moreover, Fermilab has built a successful relationship with Tezzaron over the course of the last few years. We plan to further develop our collaboration with them as part of this work.

### The VIPRAM 3D Stacking Process: MPW Prototyping

In ordinary 2D prototyping, a very large fraction of the cost is in the creation of the Mask Set, that set of images used in VLSI fabrication. The remainder of the cost is the actual fabrication. The magnitude of the expense has given rise to the so-called Multi-Project Wafers or MPWs, in which the cost is divided among several users, each of whom paying for the space they use on a reticle. The same principle can be applied in 3D fabrication and this is certainly a viable option for the VIPRAM R&D project.

The 3D MPW runs that have been available to Fermilab in the last few years will be available in the foreseeable future. For the discussion below, we assume wafer-to-wafer stacking. Other options are also available. These 3D MPW runs are two-tier single mask set processes. This means that the delivered chip will have two tiers, one on top of the other, and these tiers will be joined by face-to-face bonding. More importantly, “single mask set processes” means that both tiers of the

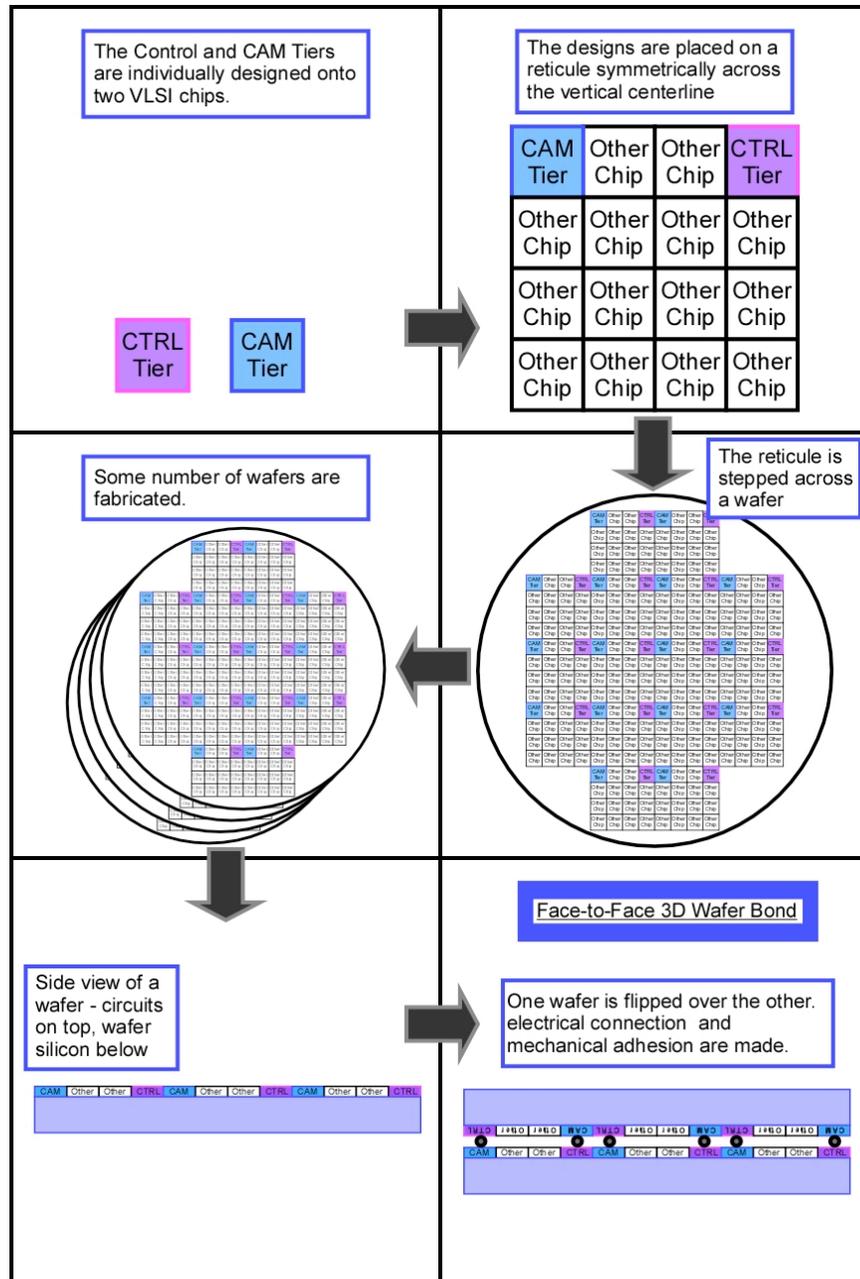


Figure 5 - A two-tier, Single Mask Set 3D MPW process

final delivered chip will be placed on the same reticle in the layout mask set and one of the two tiers will actually be flipped in the layout. This “pre-flipping” of one of the two tiers is done because when the two wafers are brought together face-to-face in 3D fabrication, one of the two must be flipped over and placed on top of the other. When this happens, the “pre-flipped” tier is “un-flipped” and the resulting 3D stack is correctly oriented.

Looking more closely at the VIPRAM design, it has only two types of tiers, Control and CAM. In the final design, there will be one Control Tier and perhaps eight CAM Tiers. However, the design is such that it will function and can be tested with a Control Tier and only one CAM Tier; in other words, a two tier design, suitable for a Single Mask Set 3D MPW. This gives the collaboration considerable flexibility in prototyping.

Figure 5 shows the typical steps used in forming a two-tier, Single Mask Set 3D MPW chip. It illustrates the symmetric placement of the two tiers and, in the final frame in the lower-right, it is clear why one of the tiers must be “pre-flipped” in the reticle.

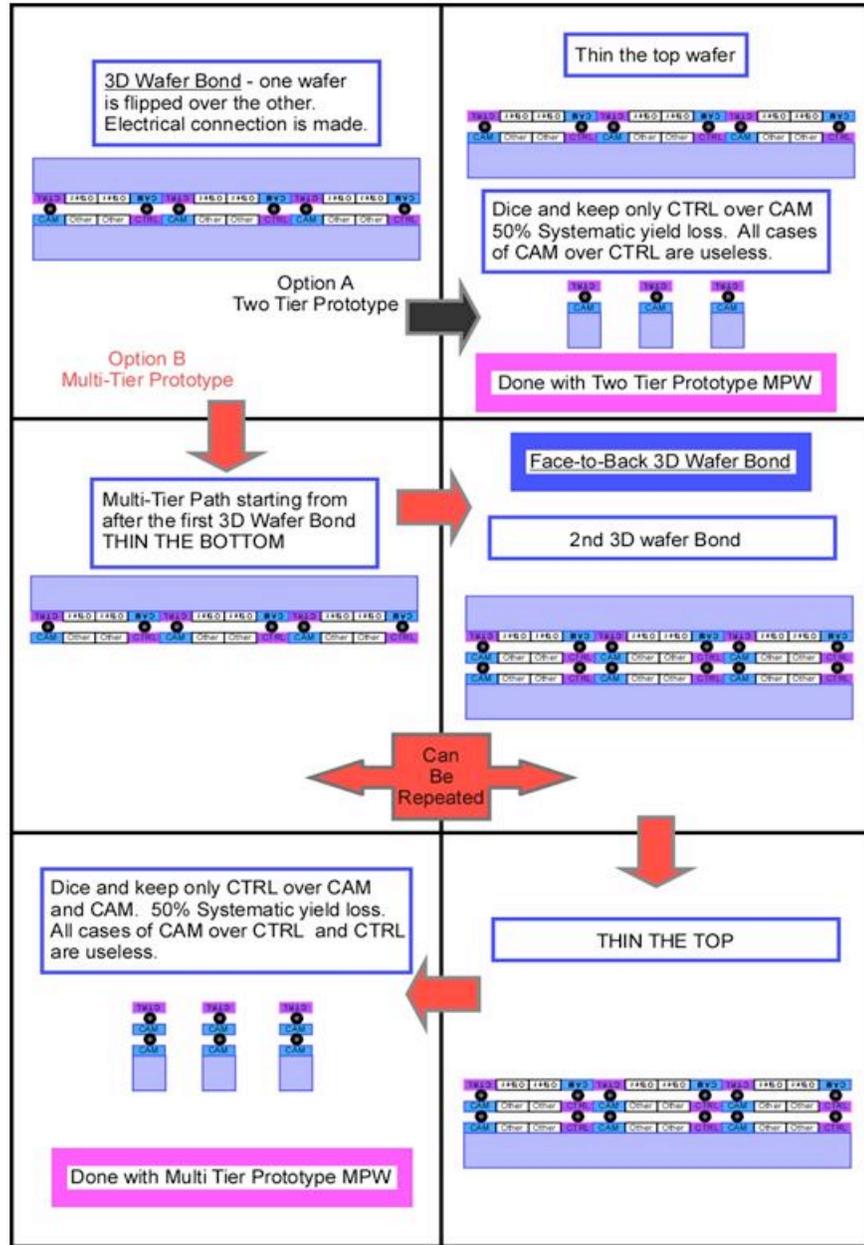


Figure 6 - The conclusion of a typical 3D MPW process OR an alternate process available to the VIPRAM.

The upper-left frame of Figure 6 is the continuation of lower-right frame of Figure 5. Figure 6 shows the two prototyping paths (shown as Option A and B) of VIPRAM using MPW.

In the first path (Option A), the MPW run can be concluded normally, just like everyone else in the MPW run. The non-supporting wafer can be thinned and then the wafer stack can be diced. The 3D chips that have CAM over Control are useless and are thrown away. The 3D chips that have Control over CAM are the correct combination and thus kept. (This is the systematic 50% yield loss implicit in all 3D Single Mask Set processes. In other words, this isn't unique to VIPRAM design). These two-tier VIPRAM chips can be then tested to verify most of the functionalities (except Majority Logic) of each tier and the success with which the tiers communicate.

In the second path (Option B), the 3D processing is continued once the first path is successful. The wafer of the top tier (in the figure) is used as a support wafer and the wafer of the bottom tier is thinned. A third tier can then be bonded to the two-tier stack in a face-to-back bond. These steps of wafer thinning/removal and additional tier stacking can be repeated if desired so that more CAM tiers can be added to the stack. The resulting 3D chips are then diced and again 50% of the chips will be CAM over Control, Control, etc. However, the other 50% will be Control over CAM, CAM, etc. and these multi-tier chips are the correct combinations and can be used for additional testing to fully demonstrate the “proof-of-principle” of the 3D design of the VIPRAM and the 3D stacking process. Incidentally, this two-path prototype technique is how Tezzaron prototyped its 3D DRAM stacking, thus VIPRAM prototyping is following a proven 3D process.

As for Phase I, the initial goal is to demonstrate the “proof-of-principle” of the 3D design of VIPRAM and the 3D stacking process by testing a Control + CAM + CAM combo stack, even though the design will be compatible with up to 8 CAM tiers (or detector layers).

## Prototyping Costs

### Dividing the Labor among Collaborators

At first glance it might seem that collaboration in 3D design is difficult if not impossible because of the rigidity of the geometric requirements. However, this is not the case. It is simply an extrapolation by another dimension of what is already done in ordinary 2D VLSI.

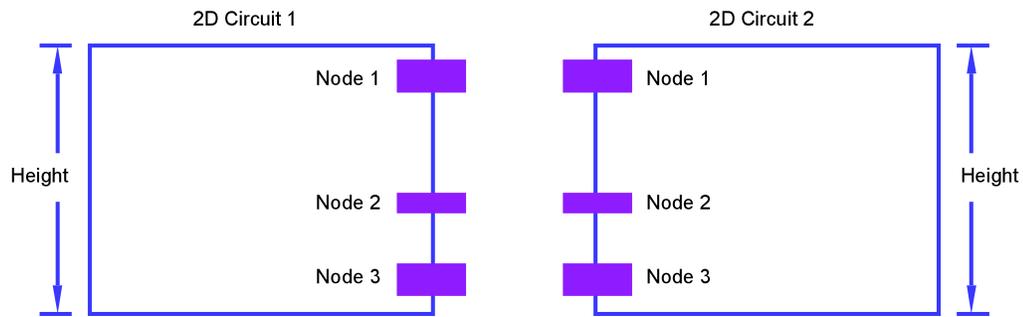


Figure 7 - Simple 2D collaboration.

In ordinary 2D design, collaboration depends on defining one dimension (here the height) and then defining the interface between those two circuits exactly and specifically. Here we see that Node 1, Node 2 and Node 3 cross from one circuit to the other. Their location and size must be defined exactly. In short, for a 2D design, collaboration depends on exact definitions along one dimension.

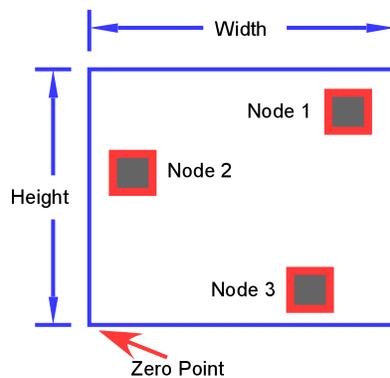


Figure 8 - An example 3D collaboration template

In a 3D design, collaboration depends on exact definitions in two dimensions. Interactions between tiers of a 3D circuit are accomplished with 3D vias. Their positions in different tiers must be known perfectly. Therefore, an exact *zero point* is defined. All 3D via locations are defined in size and 2-dimensional position relative to that zero point. By carefully placing both tiers in the reticle, the 3D via connectivity is maintained.

3D design collaboration, then, requires a definition of cell height and width, of a cell's zero point, of a cell's 3D vias and their exact position. Once these are defined, collaboration is essentially the same as in 2D VLSI design.

It is the intention of the VIPRAM collaboration to divide the VLSI design between Fermilab and the University of Chicago. Overall responsibility as well as responsibility for the cell definition and for the Control Tier will belong to Fermilab. The CAM cell will be the responsibility of the University of Chicago. This will include a translation into the expected 130nm Global Foundries process and the full custom layout of the structure. Extensive simulation will be required to ensure proper function.

## Costs and Timeline

A three year project is anticipated for the first phase. The first year will strictly be design. The second year will produce simple 2-Tier Control + CAM prototype stack for initial testing. If the testing of the first prototype reveals errors, then the third year will produce a second version of the 2-Tier VIPRAM prototype. If, on the other hand, the first version of the prototype is successful, then the third year will produce a multi-tiered VIPRAM. The minimal number of tiers needed to demonstrate the “proof-of-principle” would be a 3-tier stack: Control + CAM + CAM. To achieve this goal, the number of wafers needed is expected to be 16 assuming current, worst case yield forecasting of the 3D stacking process. This is conservative because the 3D stacking yield will be improved over time by the industry.

Most of the conceptual design work of VIPRAM has been done in 2010 [19]. The concept of VIPRAM was inspired by the initial work supported by the University of Chicago and Fermilab Strategic Collaborative Initiative Award in July 2010, for the proposal of “Rapid Identification of Heavy Quarks and Leptons at Large Hadron Collider” [14]. One unique feature of the VIPRAM architecture is its simplicity and, consequently, much of the actual design work has been already done this year, too (see Appendix for some of the details). Moreover, much can be based on the design work already done for the 2D AMchip R&D [18]. The CAM tier contains a large array of identical CAM word cells. Each cell stores the expected hit address for a given pattern and compares against the incoming hits from a given detector layer. The CAM word cell contains 15 CAM bits. The implementation of the CAM word cell is straightforward and will be largely based on the AMchip04 design [18] in which Fermilab has been deeply involved. Much of the design work for the CAM word cell has concerned power reduction. The design and layout of the CAM word cell will be done by Chicago engineer Fukun Tang. The Control tier contains a large array of identical Majority Logic cells. Each cell represents one pattern. As such the Control tier behaves very much like a pixel detector chip. The readout of the fired roads will be done using the Fischer Tree approach. Much of the design of the Majority Logic and the Fischer Tree readout has been done at the schematic level and simulated both digitally and electrically (SPICE) by Fermilab engineer Jim Hoff. The main work to be done is the layout of the Control tier. Note that the design and layout of Fischer Tree has been done by Fermilab ASIC group for a different project in the past. Much of simulation and verification work will be done by physicists. Silvia Amerio will spend 50% of her time on the project, and leads the effort on the simulation and preparation of test stand (based on AMchip test stand for CDF). She is from INFN Padova and she has the Marie Curie Fellowship to stay at Fermilab for two years. Her Fermilab supervisor is Ted Liu. She is an expert for the CDF SVT system, and she was the key person for the CDF SVT Gigafitter upgrade before she became the Marie Curie Fellow.

Due to the simplicity of the VIPRAM architecture and the fact that much of the design work has been already done, and that we will follow Tezzaron’s proven 3D process for its 3D DRAM design, the demand on Fermilab engineer time for the first phase of R&D will be relatively moderate. For this proposal, we would like to request funding to cover total of 9 months of Fermilab engineer cost, stretched over 3 years (5 + 2 + 2 months).

1. Year 1 – Design (total requested: \$182K)
  - a. University of Chicago design of CAM cell and simulation of the VIPRAM design, 6 months engineering time; \$50,000.
  - b. Fermilab design of the Control Tier and the Peripheral Logic, as well as CAM tier design and interface specifications. Up to five months of Fermilab engineering time to finish the current design and layout work (\$61K direct, \$46K indirect)
  - c. Preparing test setup (10K hardware, 15K student);
2. Year 2 – First Prototype (total requested: \$150K)
  - a. Prototype fabrication by Tezzaron Semiconductor via MOSIS Semiconductor 3D multi-project wafer run. \$65,000.
  - b. Testing of first prototype (10K hardware, 15K student, 5K travel), two months of Fermilab engineer time (\$25K direct, \$30K indirect).
3. Year 3 – Multi-Tier Prototype and testing (total requested: \$150K)
  - a. If there are errors on the first Prototype, a second prototype will be made. \$65,000.
  - b. Two months Fermilab engineer time needed for revision and testing (\$25K direct, \$30K indirect)
  - c. If there are no errors for the first prototype, then additional wafers of the first multi project wafer run will be purchased from MOSIS.
    - i. Cost: 16 wafers x \$3000/wafer = \$48000.
    - ii. Tezzaron will fabricate the multi-tier fabrication (Control + CAM +CAM combo) at a cost of approximately \$17000.
    - iii. Total cost of multi-tier run: \$65000.
    - iv. Testing (25K students, 5K travel)

## Future Plan: Phase II

While this proposal is focused on the first phase of VIPRAM development, it would be useful to also briefly describe our plan for Phase II. This is an R&D program rather than a single effort to design a 3D chip.

Increasing the AMchip pattern bank size is one important way to increase performance. As the AM bank density and size increases, the number of fired roads will increase at high luminosity due to higher occupancy. This has two consequences. First, more full resolution detector hits associated with roads have to be retrieved and transferred from the AM stage to the track fitting stage, which demands higher bandwidth between the two. Secondly, for the track fitting, the fitting speed has to be high enough to keep up with the larger number of patterns found upstream. Motivated by existing system needs, CDF recently upgraded the track fitting stage of the SVT system, the “GigaFitter Upgrade”[13]. It significantly improves the track fitting speed and performance by taking full advantage of some of the advanced features (imbedded DSPs) of modern FPGAs [3]. The speed performance of the Gigafitter is approximately one fit

per nano-second, hence the name. This is the average time to fit the hits for a track candidate and extract a goodness of fit and track parameters, and it is several hundred times faster than in the original SVT track fitter. For the LHC trigger application at the phase-1 accelerator upgrade luminosity using ATLAS FTK as an example, even with almost 1 billion AM patterns, the total number of fits will be on the order of a million per event. Maintaining such a rate in a large system will be difficult even with the Gigafitter speed performance. In particular, there is a need to transfer large numbers of found roads and the associated full resolution hits from the AM stage into the track fitting stage. Since the track fitting stage typically has to be done on a separate module from the pattern matching stage, this would pose a significant design challenge at both the board and system level (for details, see FTK proposal [7]). The problem is much more severe for the phase-2 luminosity upgrade. It is therefore highly desirable for the AM stage and track fitting stage to be implemented in such a way that the two are very close to each other, preferably within the same chip. The track fitting stage can be viewed as the second stage of pattern recognition using full resolution information. The resulting integrated chip would be much more powerful for fast pattern recognition. In addition, both the board and system level design would be significantly simplified if this level of integration can be achieved.

For Phase II, we plan to integrate the VIPRAM design with the FPGA-based track fitting stage (VIPRAM +FPGA+DRAMs+SRAMs) into a single chip, possibly using a system-on-package approach (such as the silicon interposer approach recently used for Xilinx Virtex 7 FPGA). The goal of Phase I is to solve the pattern density limitation in 2D design by vertical integration, while the goal of Phase II is to solve the problem of very large data flow between the AM stage and the track fitting stage by integrating the two stages into one chip. This second part is ultimately what must be done to address the fast pattern recognition and track fitting challenges/issues for the LHC at very high luminosity. Note that since modern FPGAs can be used, the data input bandwidth will be significantly improved as well. In addition, large memories can be integrated into the same package this way. The large memory array could be used as a hit buffer to store the full resolution input hits in a database organized for rapid retrieval, as well as lookup tables for large sets of constants for track fitting purpose.

## VIPRAM – Frequently Asked Questions

### 1. Why not simply use modern FPGAs to implement a Pattern Recognition Associative Memory?

Pattern Density. The strength of an FPGA is its ability to be reconfigured as needed without the expense and labor of a VLSI submission. However, it cannot compare to custom VLSI in the sheer ability to maximize transistor density per unit area. Earlier AMchip collaborators used FPGAs as a means of testing their logic, but the VLSI chips – even though they were not full custom – stored orders of magnitude more patterns.

### 2. Why not use the commercially available CAMs, such as in a network search engine?

A PRAM is a Pattern Recognition Associative Memory. It uses CAM structures, but it searches for related sets of hit patterns – e.g. a set of hit address patterns created by a charged particle traversing a tracking detector. The PRAM has the unique ability to search for correlations among input hits received at different clock cycles. This is essential for tracking trigger applications since the input words are the detector hits arriving from different detector layers at different times without any specific timing correlation. Each pattern has to store each fired layer until the pattern is matched or the event is fully processed. A commercial CAM looks for a single pattern match in one clock cycle. Some might argue that one could wait for the full detector readout to finish and then take all combinations to form long words and then let CAM do the matching of each word. This doesn't quite work, due to the fact that the required width of word is usually too large for the CAM chip to handle, and the need to take all combinations would require many more patterns stored.

### 3. How can the power consumption of AM be reduced significantly?

There are known techniques on how to reduce power consumption for CAMs in a significant way. For example, one of the techniques used for AMchip04 R&D is the selective pre-charge scheme. There are other ways to reduce power as well. First, custom cells need to replace semi-custom cells for each CAM cell and Majority Logic block in the design. This will reduce extraneous logic changing levels. Second, match line length needs to be reduced as much as possible because longer lines mean larger capacitances and larger capacitances mean greater power consumption. In 2D PRAMs, the custom cells will have a very positive effect on this as well because the custom cells will be smaller. In 3D, though, the improvements are more significant. The 3<sup>rd</sup> dimension of routing allows the match-lines to be as short as possible – much shorter than the 2D routing can possibly be. Tezzaron's FaStack process (see Appendix) addresses the thermal stress issues with 3D stacking by ultra-thinning. Ultra-thinning reduces the wafer thickness to as little as 8 microns, uniform to within +/- 0.5 micron. FaStack's aggressive wafer thinning prevents excess thermal buildup and allows the stack to behave as one thermal unit, and copper bonds facilitate heat dissipation.

#### **4. Why not simply increase the pattern density by going to 65 nm (or beyond) in 2D?**

Pattern density, speed and power. 3D routing represents a significant improvement in overall routing efficiency. Therefore, for the same feature size more patterns can be fit per unit area. The improved routing reduces trace length for increased speed and reduced power. Therefore, regardless of the chosen technology, 3D represents an immediate improvement in all three significant figures of merit.

Finally, the cost of fabrication should be considered. As VLSI feature sizes get smaller, the cost of fabrication is getting exponentially larger. 3D technology offers the possibility that even superior pattern density can be obtained at a lower fabrication cost.

#### **5. Is the 3D process for the VIPRAM really similar to the Tezzaron's 3D DRAM process?**

Yes. In fact, they are virtually identical both in the prototyping process and in any production level process we might choose to use. For details, see Appendix on Tezzaron's FaStack process.

#### **6. What about the yield issues of the 3D process and how to address them?**

3D processes have yield issues relative to 2D VLSI processes. This is obvious given the simple fact that for any 3D approach, first one or more 2D VLSI designs must be fabricated and then wafers need to be joined to make them 3D. At this time, 2D VLSI yields are routinely above 90%. The wafer bonding steps used in 3D processes are currently considered to have a 50% yield. This yield in wafer bonding is one of the important factors that the industry is trying hard to improve. History suggests that wafer bonding yield will improve a lot in the 3 years of the VIPRAM R&D project.

#### **7. What are the main architecture advantages of the 3D VIPRAM approach?**

PRAMs are almost tailor made for 3D VLSI design. They are logically divisible to Control and CAMs. Within the CAMs, PRAMs are further divisible detector layer by detector layer. Each of these divisions is largely independent of the other divisions. Communication between the divisions is simplicity itself – each CAM sends one bit of information to the Control and only to the Control. In a 3D VLSI design, the Control gets its own tier and each detector layer CAM gets its own tier and communication is strictly vertical up to the Control Tier. This literally should become a textbook example for a 3D VLSI design book.

It is not surprising that so many clever and simple ideas come together in the VIPRAM design with the simple “blue vertical tube” concept. The first of these is the Diagonal Via, Tezzaron's inter-tier communication patent. This allows the VIPRAM to stack identical CAM tiers one above the other with no mask alterations whatsoever. This dramatically reduces development and production costs. The second is pass transistor logic which is used in the Majority Logic (Glue Cells) to determine if a road has been found. The pass transistor logic approach significantly reduces the transistor counts and allows LEGO-style layout, and the majority function can be even subdivided detector layer by detector layer, which in 3D VLSI terms means that it can be subdivided tier by

tier. The third is the Fischer Tree or Mephisto readout logic which is a self-selecting readout logic system.

Finally, then, the advantages of a 3D VIPRAM are increased pattern density at an increased speed with decreased power density. The improved pattern density comes from a reduction in the area required to build a PRAM through transistor reduction and vertical integration. The decreased power density comes from a reduction of transistors and a minimization of parasitic capacitance. The increased speed comes from architectural modifications such as the Fischer Tree for readout.

## Appendices

### “Diagonal Vias” approach for inter-tier communications

The “diagonal vias” idea by Bob Patti has been used extensively for Tezzaron’s 3D DRAM stacking with the Control + DRAM tiers design, and we plan to follow the same approach for the 3D VIPRAM design for vertical communications between the Control tier and the CAM tiers. This solution is called the “Diagonal Via” and was patented about 10 years ago (Patti, Robert, Connection Arrangement for Enabling the Use of Identical Chips in 3-dimensional Stacks of Chips Requiring Address Specific to Each Chip, U.S. Patent 6,271,587, filed September 15, 1999 and issued August 7, 2001).

Figure 8 shows a mock-up of a diagonal via showing pads to a tier above and pads to a tier below. In face-to-back bonding, the pads to a tier above would be the upper metal layer bonding interface and the pads to a tier below would be through silicon vias (TSVs). Of course, this is not a real layout, rather a conceptual diagram. The red and yellow lines, in reality, would be made up of vertical metal-metal vias and horizontal metal traces. In short, the diagonal via is a compact method for routing signals from a tier above to a tier below or from a tier below to a tier above.

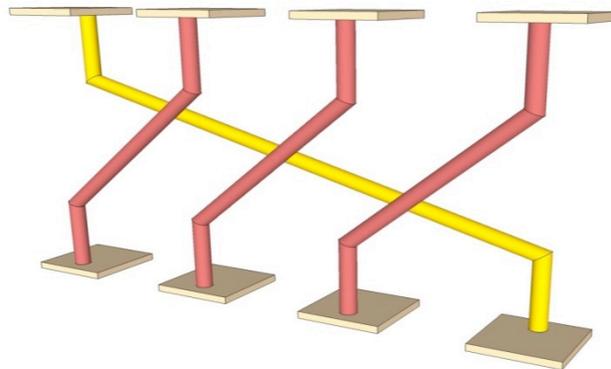
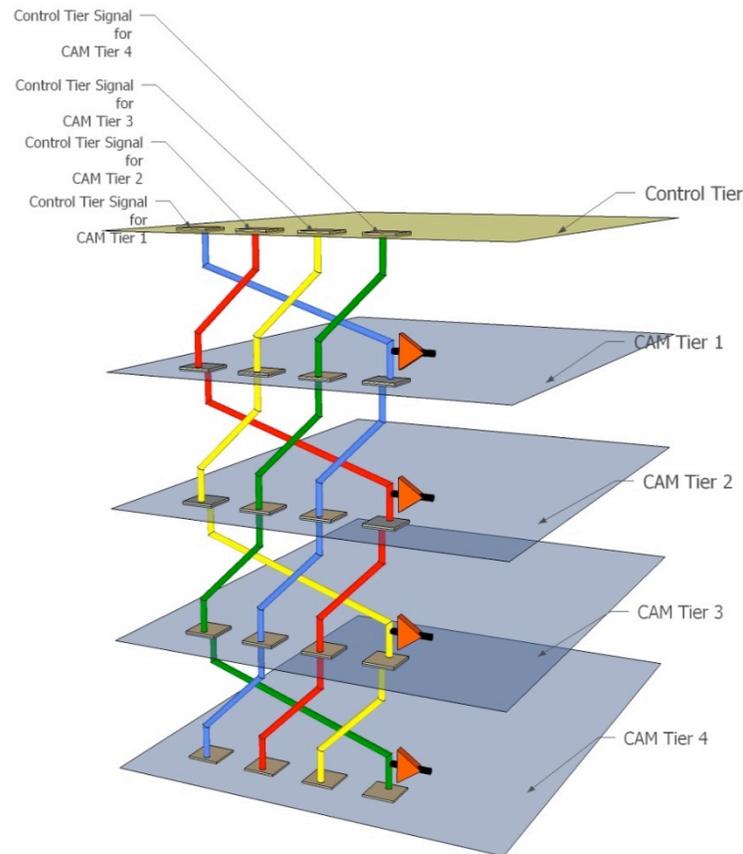


Figure 9. Diagonal Vias Concept



**Figure 10. Diagonal Vias in multiple tiers.**

The diagonal via structure allows inter-tier communication with automatic tier-self ID, without the need of any extra transistors. To see this more clearly, Figure 9 shows how this is done for signals driven from Control Tier to each of the CAM Tier. In each case, the signals are shuffled one pad to the right and the rightmost pad is routed back to the leftmost pad. In this case, the Control Tier is sending layer/tier specific data to each tier (such as input data bus from each detector layer). This same structure works with drivers on each CAM Tier and with each CAM tier sending layer/tier specific data to the Control Tier (such as match line signal from each CAM word cell). In a structure with one Control Tier and four CAM Tiers, the Control tier sees four vias, one for each CAM tier. All CAM Tiers all have exactly the same layout physically, as is required. The leftmost via on the Control Tier is for CAM Tier 1. It is obvious that the blue diagonal via takes the Control Tier information and passes it down to the receiver on CAM Tier 1. Note that the blue route continues down through CAM tiers 2-4, but it does not ever arrive again at a receiver. Only CAM tier 1 receives the information dedicated to CAM tier 1. Similarly, the rightmost via on the Control Tier is dedicated to CAM Tier 4. This is the green route in the figure. The signal begins by passing to the left on CAM tiers 1-3, but then it passes to the right on CAM Tier 4 and arrives at CAM Tier 4's receiver. Again, the only receiver to get this data is the receiver on CAM Tier 4.

## Diagonal Vias in Greater Detail

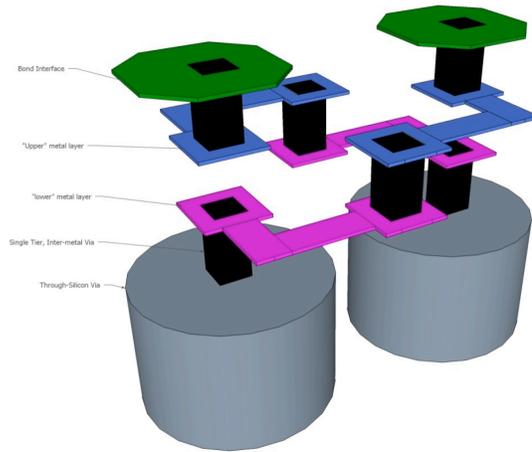


Figure 11: A more detailed representation of a Diagonal Via.

Strictly for the curious, Diagonal Vias are not a new technology. They represent no increased fabrication risk. It is simply a clever idea. The following diagram illustrates one possible routing scheme for a two-via Diagonal Via. All of the geometry shown is straightforward, routine VLSI and all of it is drawn on one tier.

On the bottom are two cylindrical through-silicon vias. They are connected to lower metal layers by simple, old-fashioned inter-metal vias. In this picture, the lower metal layer is, in fact, metal1, the lowest metal layer. This is obvious because the through-silicon via to metal via is one simple cube (the black cube that connects the grey through-silicon via to the magenta metal). In truth, this lower metal layer could be metal2 or metal3 or almost any metal layer. The lower metal layer routes the signal away from the TSV and brings it up to a higher metal layer via another standard VLSI metal-

to-metal via. This second metal layer routes the signal around and uses a third standard VLSI metal-to-metal via to connect the signal to the bond interface.

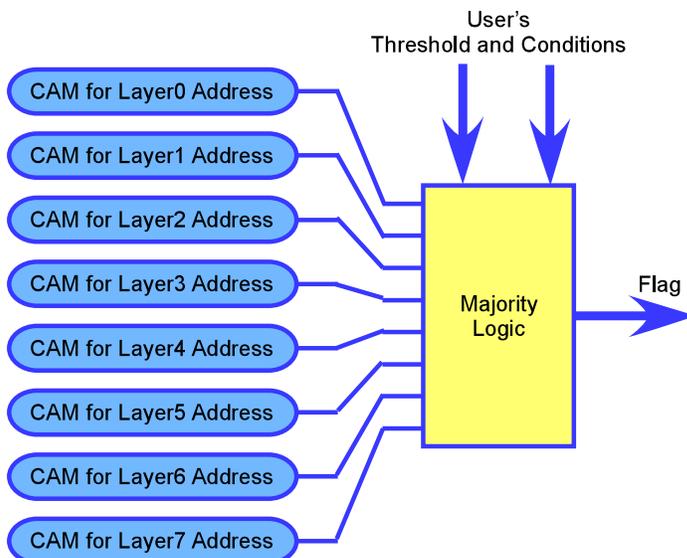


Figure 12 - The Majority Logic

Following the signals through the diagram, it is obvious that the signals move diagonally in the vertical direction even though the geometries are standard, run-of-the-mill two dimensional VLSI.

### Majority Logic – New versus Old

A Pattern Recognition Associative Memory is really a

CAM of CAMs. A simple CAM sifts through candidates and selects those that have the right combinations of bits to match its stored value. In a PRAM, many CAMs sift through candidate addresses and select those that match their internal addresses. When there is a match, it means that a charged particle passed through a particular location on a particular layer in the detector. However, tracking detectors are not looking for a hit at a location. They are looking for the track of a particle as it passes from one end of the detector to another. Therefore, after the CAMs, some other logic must sift through the **CAM matches** to determine if the right **combinations of addresses** have been matched. The block that sifts through the CAM matches is called the Majority Logic<sup>2</sup>.

Several factors complicate the task of the Majority Logic. First, while it would be simpler to insist upon a perfectly complete set of matches, inefficient detectors and the various limitations of electronic equipment make it necessary to consider slightly imperfect roads. Therefore, a user-controlled threshold is necessary to allow a user to select only perfect matches, matches with one missing CAM address or matches with two missing CAM addresses. Second, certain conditions might be necessary. For example, a user might want to disable road flagging for some period of time such as during start-up or a user might want to insist on road flagging regardless of the state of the CAM cells such as during bench testing.

In the original AMchips, the Majority Logic solved this problem with adders and digital comparators. The number of matches was summed and if it exceeded a user-defined number and if the user-defined conditions were favorable, the flag was fired. This method has its drawbacks. First, it requires a large number of transistors. Second, it requires time and, given the speeds the AMchip operates at, it therefore requires multiple pipeline stages. Third, this method is particularly ill-suited to 3D implementation. The algorithm cannot be sub-divided into tasks that can be easily placed across several tiers.

For the new AMchips, the Majority Logic has been re-designed. Rather than using adders and comparators, the new Majority Logic uses pass-transistor multiplexors, arrayed in stages by detector layer. In other words, there is one stage of the Majority Logic per CAM cell. The idea is not to count each match and compare to an arbitrary threshold; rather, each stage accepts a pattern as an input. If that stage's CAM cell outputs a match then the input pattern is passed to the output unchanged. If the stage's CAM cell does not output a match, then the pattern is left shifted. This "Match Stage" logic is shown in Table 1.

Table 1, Match Stage Outputs

Stage Input	Stage Output: Match	Stage Output: Mismatch
111	111	011
011	011	001
001	001	000
000	000	000

<sup>2</sup> The Majority Logic is sometimes called the Glue Logic although technically speaking the Majority Logic is *part of* the Glue Logic. The Glue Logic is all that is necessary to create a PRAM array and therefore includes address bits and bus drivers, etc.

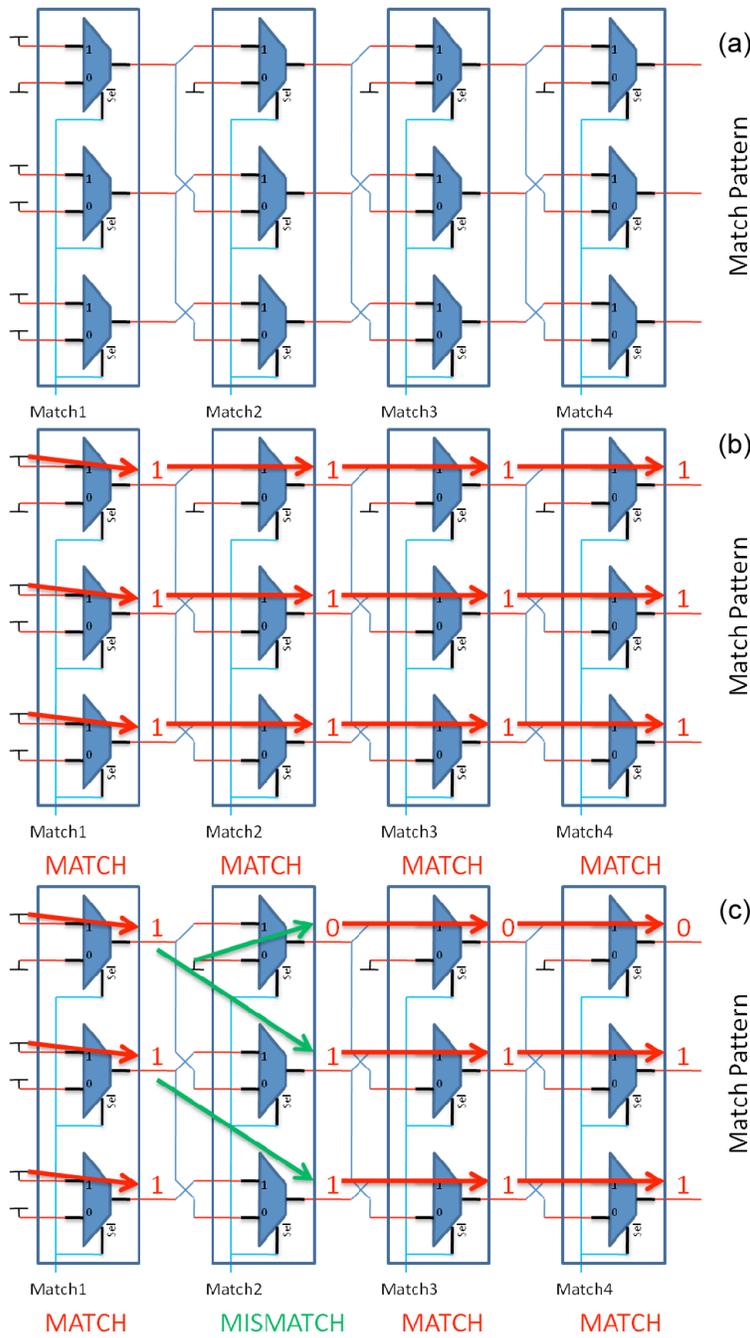


Figure 13 - Pass Transistor Multiplexers in the Majority Logic

Concatenating Match Stages as shown in Figure 13(a) grows the Match Stages into the Majority Logic. On the far left, the first input pattern is fixed at 111. If all the CAM cells output matches (Figure 13(b)), then each Match Stage will pass this 111 to its stage output meaning that the rightmost Match Stage will output a 111, indicating a perfect match. If, however, the rightmost Match Stage outputs a 011, then one and only one of the match lines was a zero (no match)(Figure 13(c)). If the rightmost Match Stage outputs a 001, then two of the match lines were zero. Finally, if the rightmost Match Stage outputs a 000, then three or more of the match lines were zero. This is the case regardless of which layers are matched and which are not. The output is only dependent on the number of missing layers. This “Majority

Pattern” is compared to the user defined threshold and affected by the user-defined conditions, resulting in the Flag. Because there are only four possible patterns, a complete digital comparator is unnecessary.

The new Majority Logic shows a dramatic reduction in transistor count (down to approximately 270 from more than 1000 for the original logic) and a significant increase in speed (approximately 2ns propagation delay from match arrival to flag). Moreover,

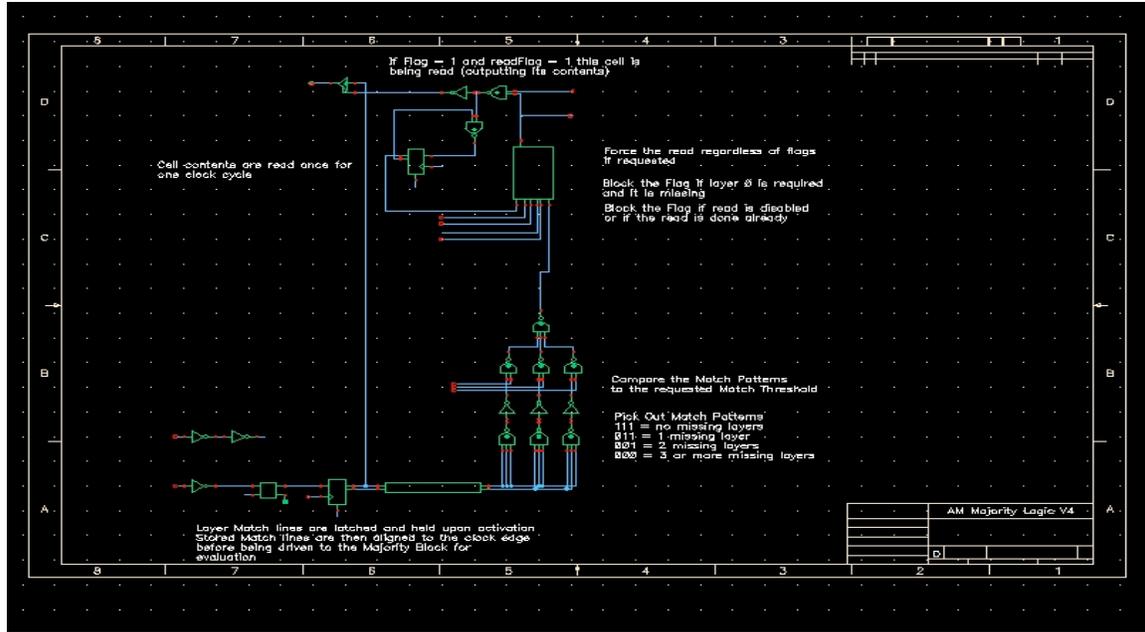


Figure 14 - The Majority Logic schematic

the Majority Logic stages are one-per-layer, making this design inherently sub-dividable by layer. Therefore, this new approach is well suited to 3D implementation.

The Majority Logic design is very advanced. It has been designed and simulated both logically and electronically (SPICE). All that remains to be done is for it to be laid out in a technology appropriate to 3D implementation. Its schematic is shown in Figure 14.

### The Fischer Tree for Readout

With the 3D VIPRAM design, the Control tier would greatly resemble a pixel detector or mini-strip readout in both geometry and behavior. Geometrically, both pixel detectors and the VIPRAM Control tier are rectangular arrays of devices whose positions are indicative of the device's addresses. In the case of a pixel detector, the address indicates the location of the individual pixel. In the case of the VIPRAM Control tier, the address indicates which road fired. Behaviorally, high speed readout is essential. In the case of a pixel detector, readout speed is determined by the likelihood that data might be lost if a second charged particle were to pass through the pixel before it could be read out. This is largely a function of luminosity and event rate. In the case of the VIPRAM Control tier, a particular road will fire once and only once per event so multiple firing is not an issue because of the nature of the PRAM's task; rather readout speed is determined by the number of fired roads present in an event. This, too, is a function of luminosity and event rate.

In both cases, there is a need for high-speed readout across a large silicon area and the same techniques that have been successful for pixel readout can be used for VIPRAM readout. In particular, the Fischer Tree [17] is particularly qualified for this purpose.

Fischer Trees were first introduced by Peter Fischer in 2001<sup>3</sup>. They are a simple binary tree each node of which is shown below in Figure 15. The logic is very simple. If either “Channel 1 Flag” or “Channel 2 Flag” is active, then “Flag on 1 or 2” is active. As long as “Pick 1 or 2” is inactive, nothing further happens. However, once “Pick 1 or 2” activates, then either “Choose Channel 1” or “Choose Channel 2” will activate, but not both of them. The configuration shown in Figure 15 is an “Up Dominant” Fischer Tree in that if both “Channel 1 Flag” and “Channel 2 Flag” are active, then “Choose Channel 1” is activated. It is just as easy to create a “Down Dominant” Fischer Tree.

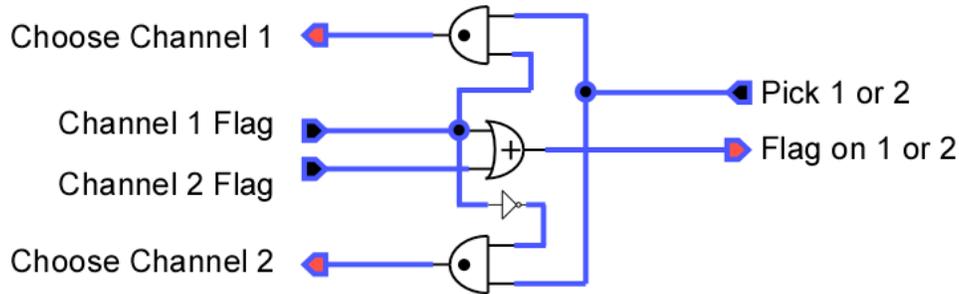


Figure 15 - A Fischer Tree node

Given the design of a single node, it is easy to extend the Fischer Tree to any number of channels that is a power of 2. A four channel Fischer Tree is shown in Figure 16.

There are several distinct advantages to the Fischer Tree. First and foremost is speed. The Fischer Tree is purely combinatorial, so no clocks are necessary for its use. The flags propagate forward and the choices propagate backward in logN time so larger and larger Fischer Trees do not grow dramatically slower and slower. Moreover, Fischer Trees are self-addressing as is shown in Figure 17 (ignoring the need for inversion to

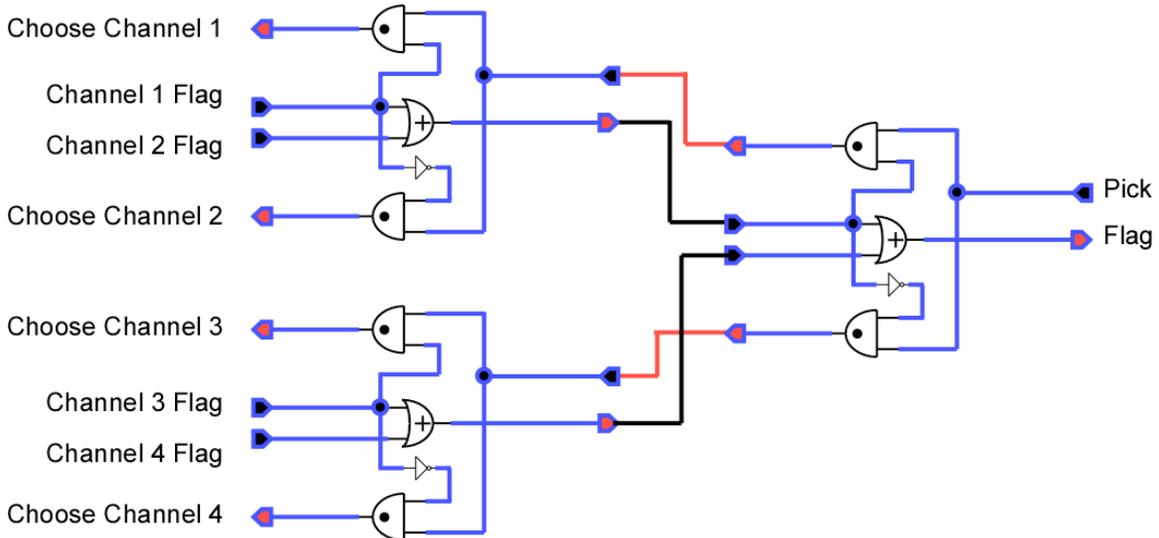


Figure 16 - A four channel Fischer Tree

<sup>3</sup> “First implementation of the MEPHISTO binary readout architecture for strip detectors” Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and

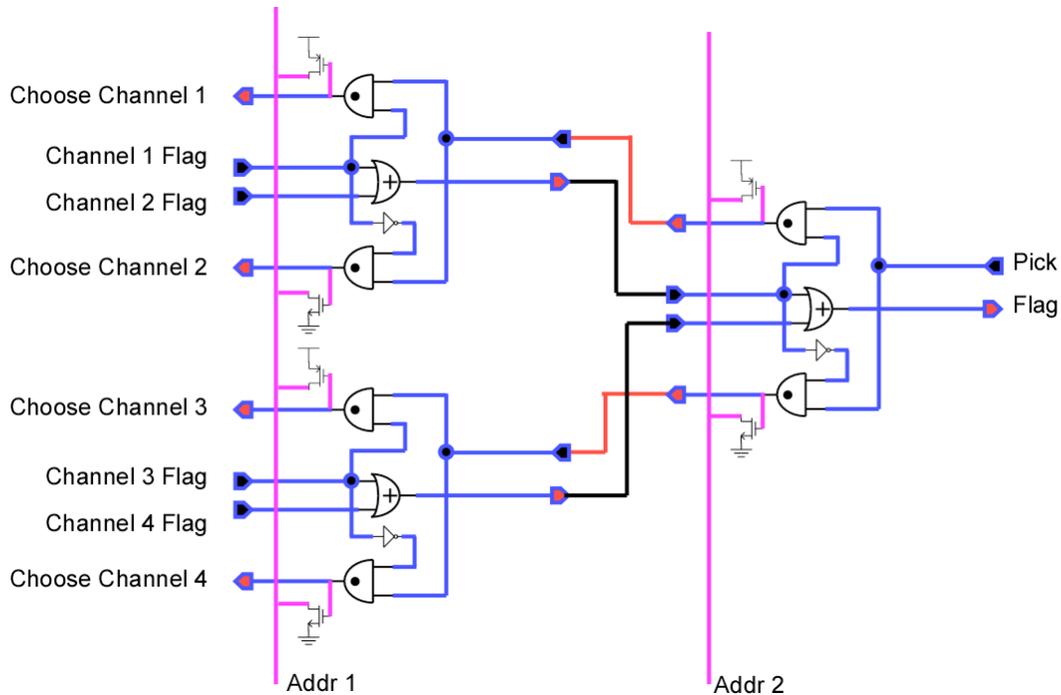


Figure 17 - A 4-Channel Fischer Tree with Addressing.

properly operate the pFets). Here it can be seen that if “Choose Channel 4” is active, then “Addr2” is pulled down to a zero and “Addr1” is also pulled to a zero. If “Choose Channel 2” is active, then “Addr 2” is pulled up to a one and “Addr 1” is pulled to a zero. This is a simple consequence of the binary nature of the Fischer Tree. It means that as soon as a channel is chosen, its address is already available at the periphery.

At the moment, schematics and simulations are complete for Fischer Trees of up to 512 channels. The propagation delay is approximately 3ns in the absence of parasitic capacitance, which will slow them down. Layouts are required for final determination of speed.

### How CAM Works and 3D Advantages for PRAM

To understand the advantages of the 3D technology for PRAM, it is useful to provide an overview of the fundamental architecture of both the CAM as well as the Associative Memory. We will first take a look at the basic architecture of conventional CAM, identify the uniqueness of the HEP AM, and describe how we might take advantage of the 3D technology to enhance the AM performance.

Conventional CAMs store an array of address patterns that a user wishes to compare to a stream of candidate addresses [2]. Each new candidate address is presented to the chip where it is compared simultaneously to each stored address pattern in the array. If there is a match, a flag is raised. This simple algorithm is somewhat complicated by the fact that more than one stored address pattern can flag a match. In such a case, a priority encoder must select one of the matches as the CAM’s chosen match.

A CAM has a regular architecture with a few basic components such as the CAM cell, search-lines (SL), match-lines (ML) and match-line sense amplifiers (MLSA) [2]. A CAM cell serves for two basic functions: bit storage and bit comparison, which can be a NOR or NAND-type cell. When multiple cells are connected in parallel to form a CAM word the match-line of each cell is shorted to the ML of an adjacent cell. In the case of AMchip03, the design is for 6 detector layers and each has about 16 bits, therefore the total number of CAM bits is about 100. For a large number of CAM bits like this, a long ML line can be created which has parasitic resistance and capacitance and contributes to the power consumption. Figure 11 shows a CAM model consisting of 4 words, with each word containing 5 bits arranged horizontally. A CAM search operation begins with loading the search-data word into the search-data registers followed by pre-charging all match-lines high, putting them all temporarily in match state. Next, the search-line drivers broadcast the search word onto the search-lines, and each CAM cell compares its stored bit against the bits on its search-lines. If there is a match, the match-lines remain high; in case of a miss, match-lines discharge to ground. Match-line sense amplifiers detect whether each ML has a match or a miss condition. Finally, the encoder maps the matching location to its matching address [15] [16].

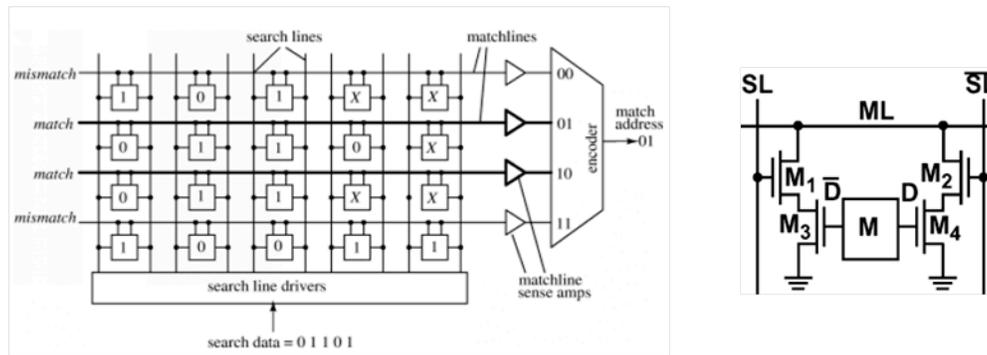


Figure 18. CAM model for 4 words with each word containing 5 bits [16]

A CAM compares input search data against a table of stored data and returns the address of the matching data. CAMs have a single clock cycle throughput making them much faster than other hardware or software based search systems. However, the speed of a CAM comes at the cost of increased silicon area and power consumption. As CAM size increases, so does the power consumption. Thus, power reduction is the main challenge in CAM design without sacrificing speed or area.

In a PRAM it is logical to divide the Pattern Address Array into banks of address pattern by layer number. This effectively divides the PRAM into  $N$  parallel conventional CAMs, one for each detector layer. In standard 2D integration, such a division increases the design size due to the routing necessary to link each Match Flag in a road. In 3D, that routing area can be virtually eliminated. In this way, there will be two types of tiers: one is the top tier or Control tier, which houses the IO and Road Glue Logic; the other is the CAM tiers – one per detector layer – which house the individual CAM arrays. The resulting architecture is the basic design concept of the VIPRAM, as shown in Figure 4.

It turns out that the VIPRAM approach is remarkably similar to the approach Tezzaron uses for their 3D DRAM stacking. Tezzaron Semiconductor develops multi-tiered 3D memory arrays. While these are not CAM arrays, CAM arrays and memory

arrays share much in common. Using a true 3D approach, Tezzaron divides the functionality between the tiers. Being cost conscious, they limit to two types of tiers. The top tier is the control tier that contains the IO logic, the sense arrays, the decode logic and the address line drivers. The remaining tiers are DRAM tiers, connected to the control tier by through silicon vias. Tezzaron takes the further step of fabricating the two different types of tiers in two different CMOS processes. The control tier is optimized by using CMOS high-speed processes that create high-performance transistors. The DRAM tiers use a high-density NMOS process that creates high-quality capacitors. The end result is a faster, denser memory without any changes to the design<sup>4</sup>.

### Tezzaron's FaStack® Technology (see also Tezzaron's web site)

Tezzaron's FaStack technology creates fast, dense, highly integrated 3D chips. The heart of the process is copper thermal diffusion stacking with very dense arrays of vertical interconnects. FaStack can bond either die-to-wafer or wafer-to-wafer, use either of two different types of vertical interconnects, and built the interconnect into the wafers with any of three different process flows. For this proposal, we choose to implement the VIPRAM by following the same FaStack method that Tezzaron uses for its own 3D DRAM products.

The FaStack method that we propose to use begins with hundreds of thousands of tungsten "Super-Contacts" built into the circuitry of each wafer during normal wafer processing. The wafers are then metalized by coating them with a 0.5 micron SiO<sub>2</sub> insulating glass layer and then a 1.0 micron Cu metal bond point layer with a proprietary layout design. The first two wafers are aligned face-to-face and bonded using a copper thermal diffusion process at less than 400 °C. The structural base (back side) of the upper wafer is then thinned to less than 10 microns using a combination of conventional wafer grinding, spin-etching, and chemical-mechanical polishing (CMP). The thinning exposes the Super-Contacts that were built into the top wafer. The back side of the thinned wafer, with its exposed Super-Contacts, can be metalized with bond points and bonded to the front side of a third metalized wafer. Thinning, metalizing, and bonding are repeated as desired. Once the wafer stacking process is completed, one side of the stack is thinned to the Super-Contacts and padded out for I/O; the other side is back-lapped to remove excess silicon.

A semiconductor wafer is usually about 750 microns thick, but its electrical activity is confined to a surface layer from 4-10 microns thick. The functional part of a wafer is thus a tiny proportion of its thickness; the rest of the wafer provides only structural support. The Tezzaron FaStack process uses most of the structural base of the first silicon wafer, but keeps less than 15 microns of each additional wafer in the stack. This produces multi-layer chips that fit easily into standard packaging. Unlike many other stacking methods, FaStack bonds the wafers before thinning. This means that the structural base of the first

---

<sup>4</sup> See <http://www.tezzaron.com/technology/FaStack.htm>

wafer supports the additional wafers as they are thinned. FaStack does not require thin wafer handling, temporary bonds, or the use of “handle wafers.”

The Super-Contact density can reach 300,000 per square mm (typical designs use ~10,000 per square mm). The alignment precision for 200 mm wafers has a 3-sigma process tolerance of  $\pm 1$  micron, but precision with  $\pm 0.3$  micron is typical. Ultra-thinning reduces the wafer thickness to as little as 8 microns, uniform to within  $\pm 0.5$  micron. The industry has expressed concern about potential thermal stress in 3D stacked chips. The FaStack process addresses this issue by aggressive ultrathinning of the wafers to prevent thermal buildup, allowing the stack to behave as one thermal unit. The copper used in the bonding process provides additional relief by facilitating heat dissipation.

Note that Tezzaron’s first key breakthrough in 3D development was the “Super-Via,” a vertical copper structure that adapted standard process flow wafers to Tezzaron’s 3D stacking process. In addition to vertical interconnect, the Super-Via structure provided alignment marks, thinning control, and bonding surfaces in a single structure. Since the early development and success with the Super-Via, Tezzaron developed a second generation of interconnect, the tungsten Super-Contact. This second generation interconnect adds more design flexibility while drastically decreasing the 3D interconnect footprint. The size of the Super-Via was 4.0 x 4.0 micron in its first incarnation; the Super-Contact is 1.2 x 1.2 micron, while face-to-face bonding has a size of 1.7 x 1.7 micron. Minimal pitch is 6 micron, < 4 micron, and 2.4 micron respectively.

Tezzaron built the first working 3D IC prototypes (six different devices) in 2004. In 2008, Tezzaron began producing custom stacked components under contract and now provides stacking services for a number of customers. FaStack devices have many advantages over their single-layer counterparts: they are much denser and their short vertical interconnects allow them to operate at higher speeds with a lower power budget. As an example, Tezzaron’s prototype FaStack 8051 processor, built in 2004, runs at either 5 times the speed of a normal 8051 or 10% of the power. In addition, FaStack allows disparate elements to be processed on separate wafers for simpler production and greater optimization. For details, also see Bob Patti’s paper “3D Integration at Tezzaron Semiconductor Corporation”, Handbook of 3D Integration 2008.

FaStack offers benefits to a variety of applications. Sensor arrays, for example, achieve unprecedented density by moving the support circuitry to a different layer than the sensors themselves. “System-on-Chip” (SoC) devices built with FaStack reduce power consumption, footprint, and interconnect delays. Microprocessors built with FaStack incorporate a huge, fast memory cache on a separate layer. FaStack also enables enormous improvements in memory technology and allows seamless integration of differing substrates. As 3D processing moves into the mainstream, entirely new products will emerge to capitalize on this technology.

### Expected Areas of Improvement in the True 3D Architecture

As described previously, True 3D architecture offers an increase in pattern density over 2D designs. There are other expected improvements as well:

1. **Power/Thermal** – Only the Control Tier will operate at full speed – i.e. data from all detector layers pushing through the tier. CAM tiers will operate at  $1/L$  speed (where  $L$  is the number of tiers or detector layers) – i.e. only data from a particular detector layer will be pushed onto a particular CAM tier. Therefore, the sensitive internal layers should have lower power consumption. In the Tezzaron's FaStack process, ultra-thinning reduces the wafer thickness to as little as 8 microns, uniform to within  $\pm 0.5$  micron. The FaStack's aggressive wafer thinning prevents excess thermal buildup and allows the stack to behave as one thermal unit, and copper bonds facilitate heat dissipation.

2. **Speed** – The Control Tier of the VIPRAM contains a 2 dimensional array of Road Glue Logic cells. Moreover, there should be much less routing on both the Control Tier as well as the CAM tiers within this 2 dimensional array, comparing to the AMchip0x 2D design. The extra routing space and the regularity of the 2 dimensional arrays of Road Glue Logic cells can be exploited for speed. For example, the actual road addresses can be placed on the periphery. A detected road would activate one "row" address and one "column" address that, taken together, would constitute the unique address of the detected road. Previous work in pixel readout architectures can be used to maximize the VIPRAM's speed, such the Fisher tree approach[17].

A rough estimate shows that with the VIPRAM design in 3D, one could gain at least two orders of magnitude in pattern density over the AMchip03. One of the goals of the present R&D is to quantify this gain.

## References

- [1] M. Dell'Orso and L. Ristori, "VLSI Structures for Track Finding," *Proceedings in Nuclear Instruments and Methods*, vol. A278, pp. 436-440, 1989.
- [2] T. Kohonen, "Content-Addressable Memories," 2<sup>nd</sup> edition, New York, Springer-Verlag, 1987.
- [3] M. Dell'Orso (for the CDF Collaboration), "The CDF Silicon Vertex Trigger", *Nuclear Physics B - Proceedings Supplements* Volume 156, Issue 1, June 2006, Pages 139-142
- [4] J. Adelman et al., "The Silicon Vertex Trigger upgrade at CDF", *NIM A* 572 (2007) 361
- [5] The CDF Collaboration, "Observation of  $B_0(s)$ - $B_{\bar{0}}(s)$  Oscillations", *Phys. Rev. Lett.* 97, 242003 (2006).
- [6] J. Donini et al., "Energy calibration of b-quark jets with  $Z \rightarrow b\bar{b}$  decays at the Tevatron collider", *Nucl. Instr. and Meth. A* (2008)
- [7] "FTK: A Hardware Track Finder for the ATLAS Trigger," Technical Proposal, 2010, [https://edms.cern.ch/file/1064724/1.0/FTK\\_TP\\_USG.pdf](https://edms.cern.ch/file/1064724/1.0/FTK_TP_USG.pdf)
- [8] A. Annovi et al. "A VLSI Processor for Fast Track Finding Based on Content Addressable Memories," *IEEE Transactions on Nuclear Science*, vol. 53, no. 4, pp. 1-6, 2006.
- [9] A. Annovi et al., "Associative Memory Design for the Fast Track Processor (FTK) at Atlas", Proceedings of the "IEEE - 17th Real-Time Conference" RT10, Lisboa (2010), [cdsweb.cern.ch/record/1273162/files/ATL-DAQ-PROC-2010-013.pdf](http://cdsweb.cern.ch/record/1273162/files/ATL-DAQ-PROC-2010-013.pdf)
- [10] K. Banerjee, S. Souri, P. Kapur, and K. Saraswat, "3-D ICs: A Novel Chip Design for Improving Deep-Sub micrometer Interconnect Performance and Systems-on-Chip Integration," *Proceedings of the IEEE*, pp. 602-633, 2001.
- [11] C.A. Zukowski, and S. Y. Wang, "Use of Selective Precharge for Low-Power Content-Addressable Memories," Proceeding in the IEEE International Symposium of Circuits and Systems, vol.3, pp. 1788-1791, 1997.
- [12] Tezzaron Semiconductor, <http://www.tezzaron.com/>.
- [13] S. Amerio et al., "The GigaFitter: Performance at CDF and perspectives for future applications", *J.Phys.:Conf.Ser.* 219 022001, 2010

- [14] Ted Liu, Mel Shochet, Jinlong Zhang, “Rapid Identification of Heavy Quarks and Leptons at the Large Hadron Collider”, University of Chicago and Fermi National Accelerator Laboratory Strategic Collaborative Initiative Award. July 2010.
- [15] K. E. Grosspietsch, “Associative Processors and Memories: A Survey,” *IEEE Micro*, vol. 12, no. 3, pp. 12-19, 1992.
- [16] K. Pagiamtzis and A. Sheikholeslami,  
“Content-addressable memory (CAM) circuits and architectures: A tutorial and survey,”  
*IEEE Journal of Solid-State Circuits*, vol. 41, no. 3, pp. 712–727, March 2006.
- [17] P. Fischer, “First implementation of the MEPHISTO binary readout architecture for strip detectors” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, Volume 461, Issues 1-3, 1 April 2001, Pages 499-504 8th Pisa Meeting on Advanced Detectors.
- [18] A. Annovi, et al. “Variable Resolution Associative Memory” for High Energy Physics. Submitted to the Advancements in Nuclear Instrumentation, Measurement Methods and their Applications (ANIMMA), Ghent Belgium, 6-9 June, 2011.
- [19] T. Liu et al. “Proposal for the development of 3D Vertically Integrated Pattern Recognition Associative Memory (VIPRAM). Fermilab Technical Memo: FERMILAB-TM-2482-E-PPD.