Associative Memory Design for the Fast TracKer Processor (FTK) at ATLAS

for the AMchip Collaboration

FTK Architecture (final system)

Complex system, many units:

- **48** Data Formatters (DF)
  - Clustering Mezzanine
- **128** Processing Units
  - AUX Board (FPGA):
    - Data Organizer (DO)
    - Track Fitter (TF - 8 layers)
    - Hit Warrior (HW)
  - AM Board with ~10M patterns on AMchip04 custom CAMs
- **32** Final Boards (FPGA)
  - Final Fit (11 layers)
  - Final Hit Warrior

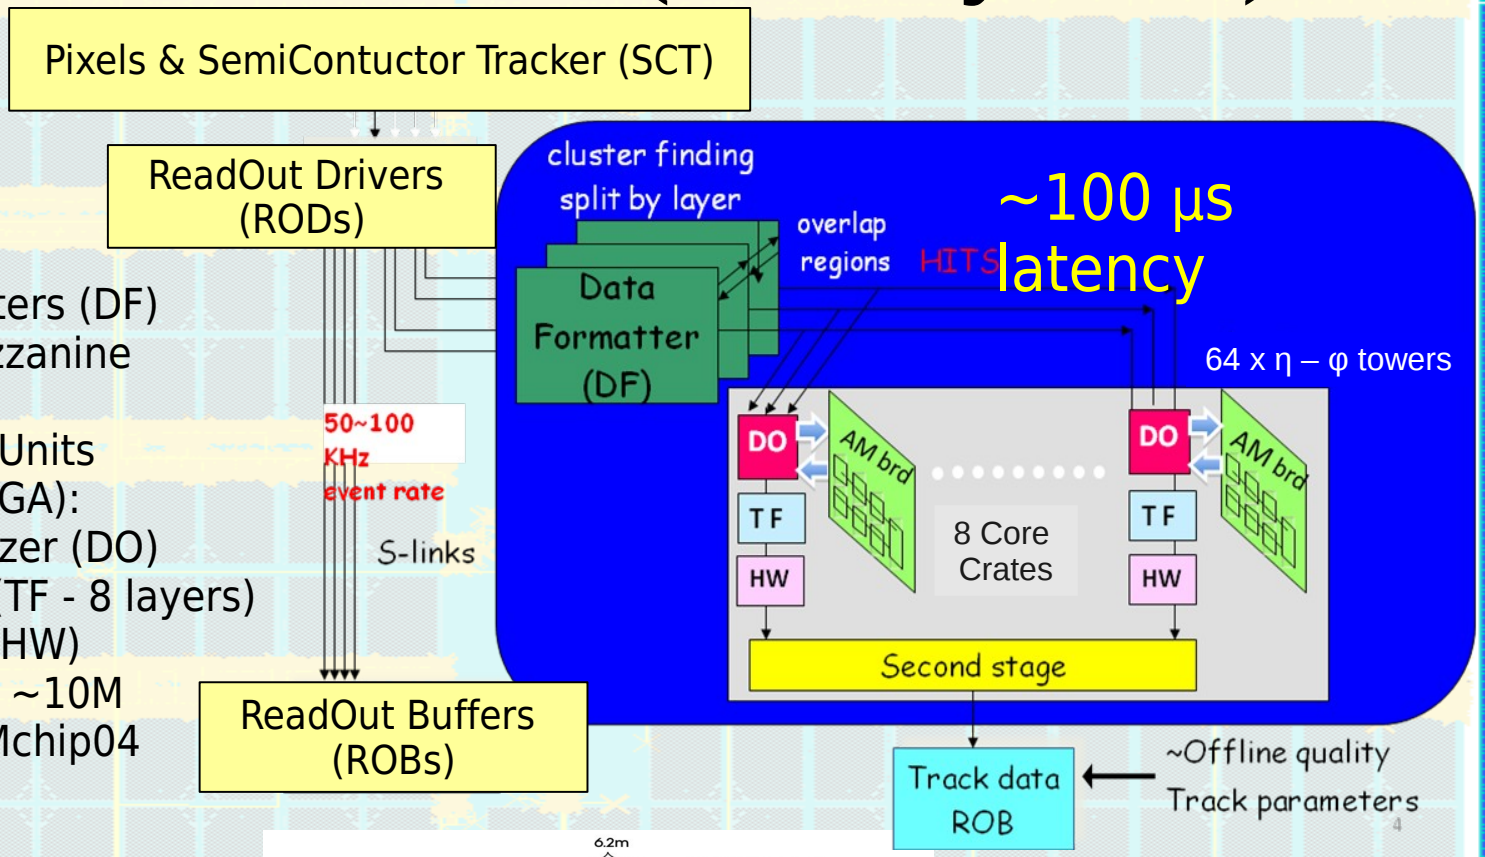

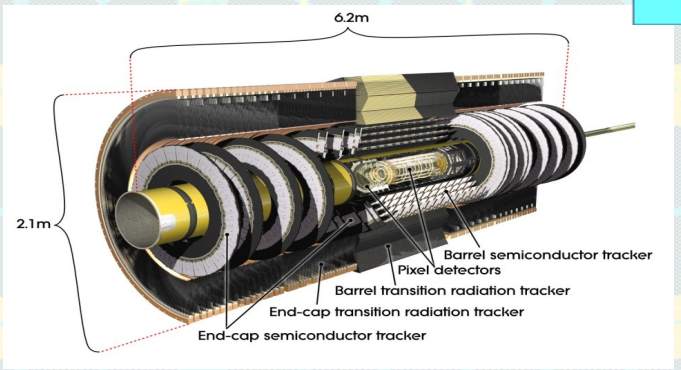

FTK will reconstruct tracks in all Inner Detector regions

The *Associative Memory*

- Dedicated device - maximum parallelism:
- Each pattern with private comparator
- Track search during detector readout

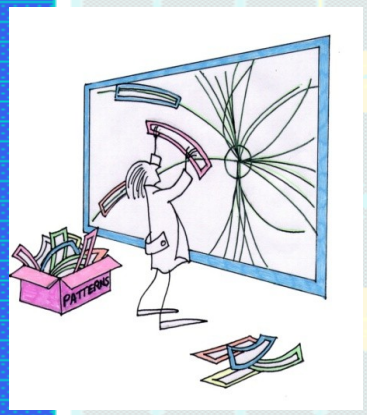

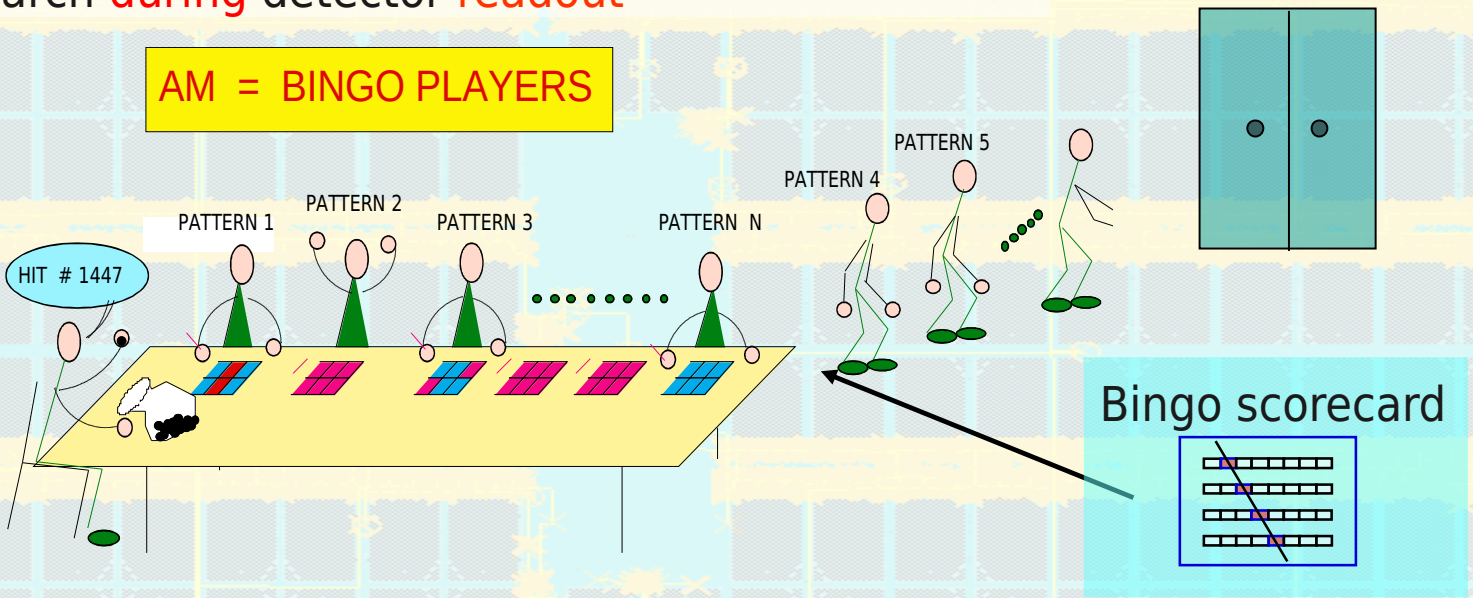

• **Full custom**	**700**	**nm:**	**0,128**	**6L**	**kpat/chip**
• **FPGA**	**350**	**nm:**	**0,128**	**6L**	**kpat/chip**
• **standard cell**	**180**	**nm:**	**5,0**	**6L**	**kpat/chip**
• **new for FTK**	**65**	**nm:**	**~80**	**8L**	**kpat/chip**
• **3D: 2 Tiers**	**65**	**nm:**	**~160**	**8L**	**kpat/chip**

AM working principle

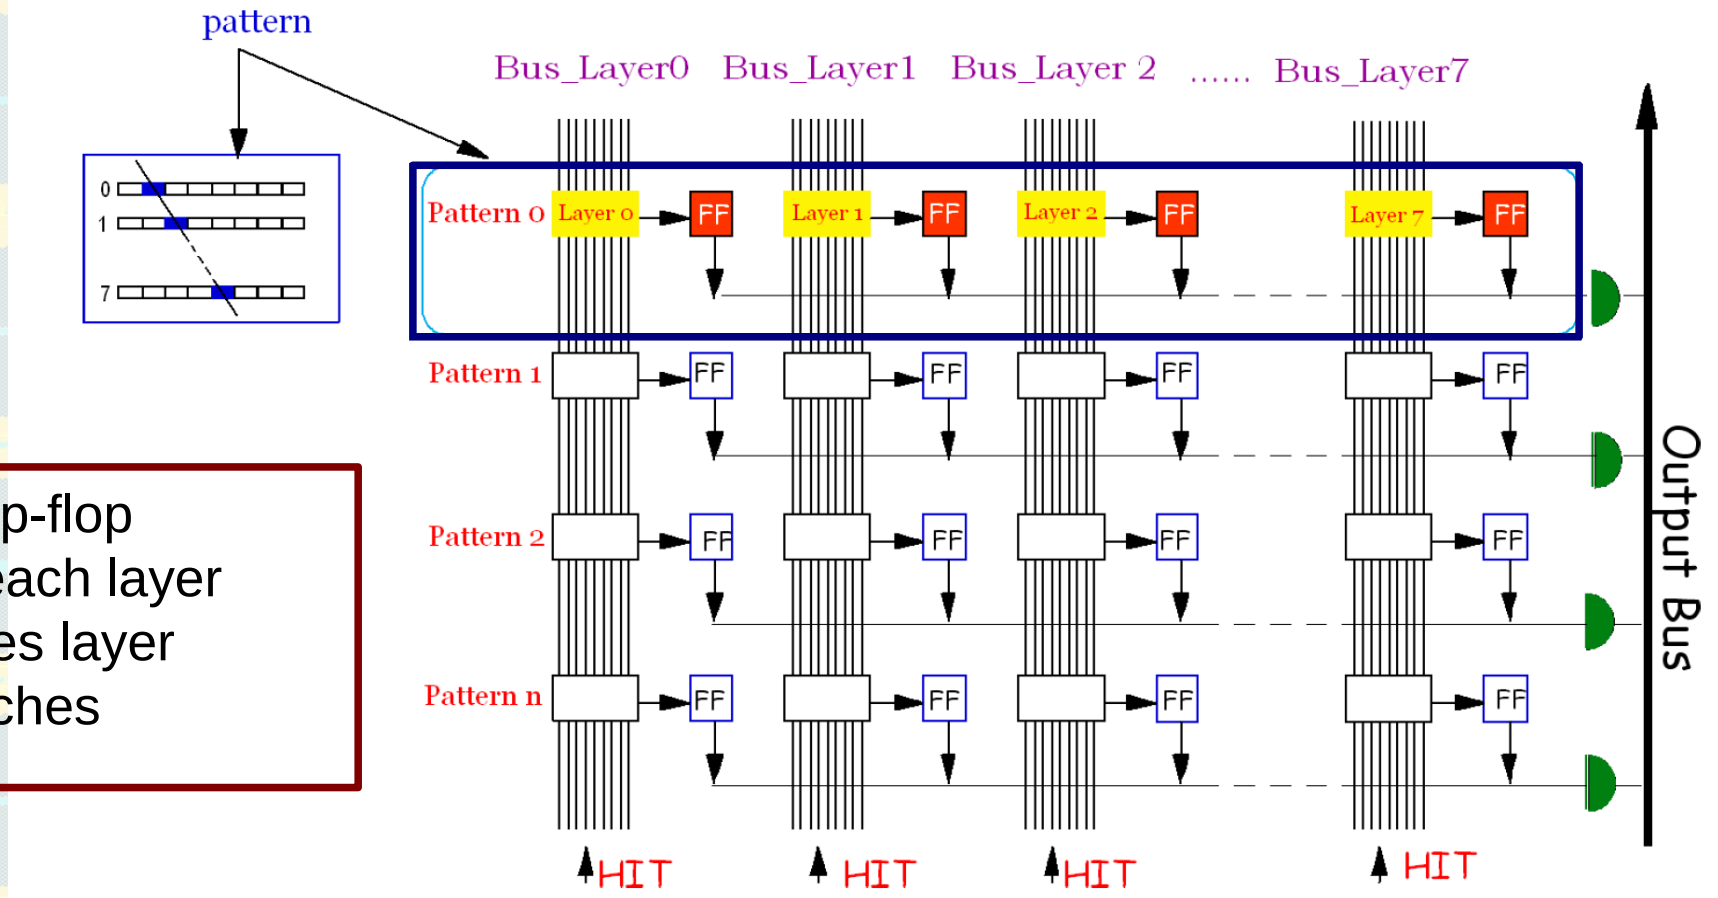

1 Flip-flop for each layer stores layer matches

All patterns compared in parallel with incoming data. Look for correlation of data received at different times. (Feature unique to AMchip)

Fast pattern matching. Flexible input.

Requirements for FTK

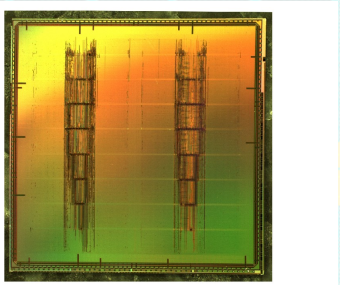

What we have now: **Standard Cell 180** µm
5000 pattern/chip for 6-layer patterns,
2500 pattern/chip for 12-layer patterns

"A VLSI Processor for Fast Track Finding Based on Content Addressable Memories",
IEEE Transactions on Nuclear Science, Volume 53, Issue 4, Part 2, Aug. **2006** Page(s):2428 - 2433

65 nm technology provides a factor 8 → 20000 patterns/chip
Full custom cell provides at least a factor 2 → 40000 patterns/chip

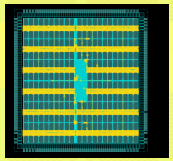

NEXT: NEW VERSION For both L1 & L2

8 layers instead of 12 provides a factor 1,5 → 60000 patterns/chip
1,2 x 1,2 cm^2 2D chip → 80000 patterns/chip
With a **2 D chip** we gain a factor 20 **(n pattern x m layer)** !

1 AMboard: 128 chips → ~10 Mpatterns per board
1 Crate: 16 AMboard → ~160 Mpatterns per crate
100 MHz running clock

First prototype is 14 mm^2 with 8K pattern

AM chip internal logic

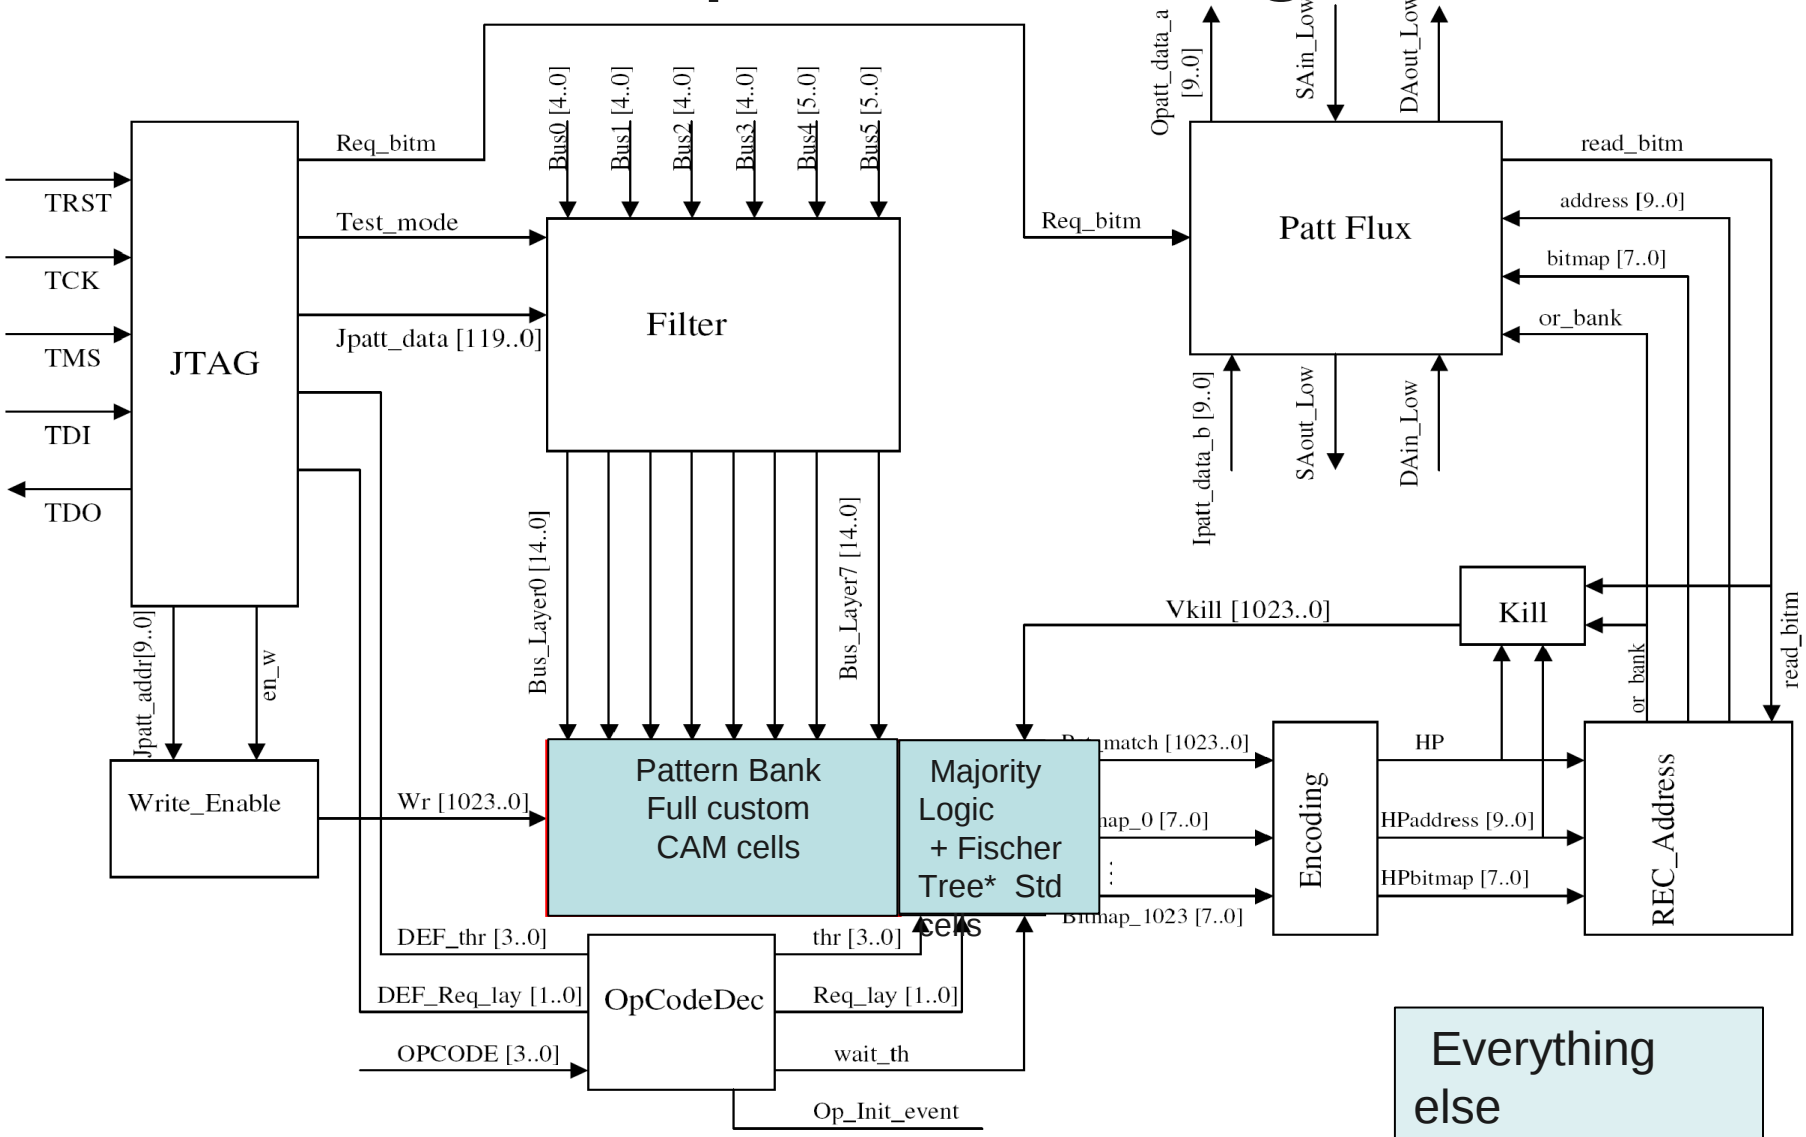

Everything else Is standard cells

* the AM readout is based on a modified Fischer Tree [1]

[1] P. Fischer NIM A461 (2001) 499-504

AM Chip Memory Layer

To save power we have used two different match line driving scheme:

- *Current race scheme*
- *Selective precharge scheme*

Current source:
3.7 x 1.8 um each

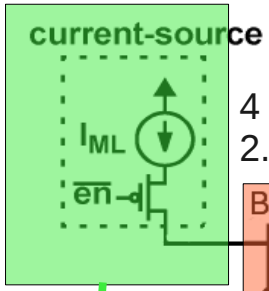

4 NAND Cells:
2.6 x 1.8 um each

14 NOR Cells:
2.6 x 1.8 um each

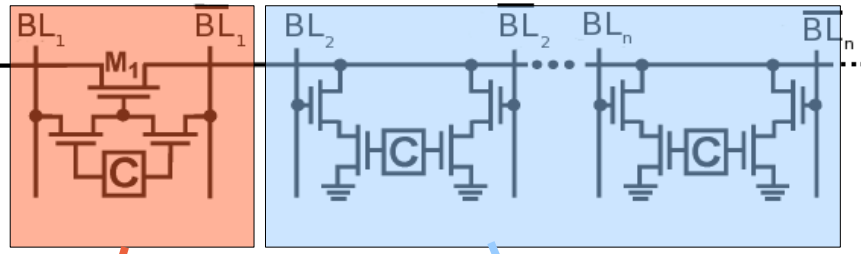

Latch + ML discharge:
4.7 x 1.8 um each

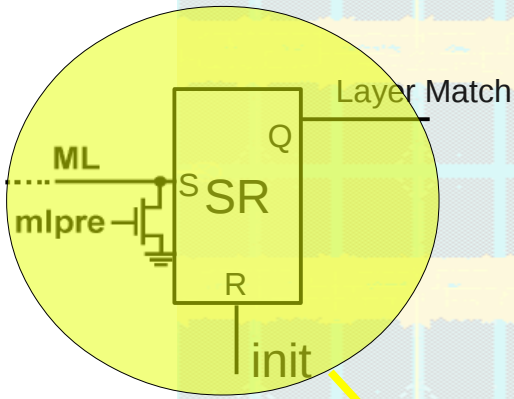

Full Layer layout : w=~53µm X h=1.8 µm

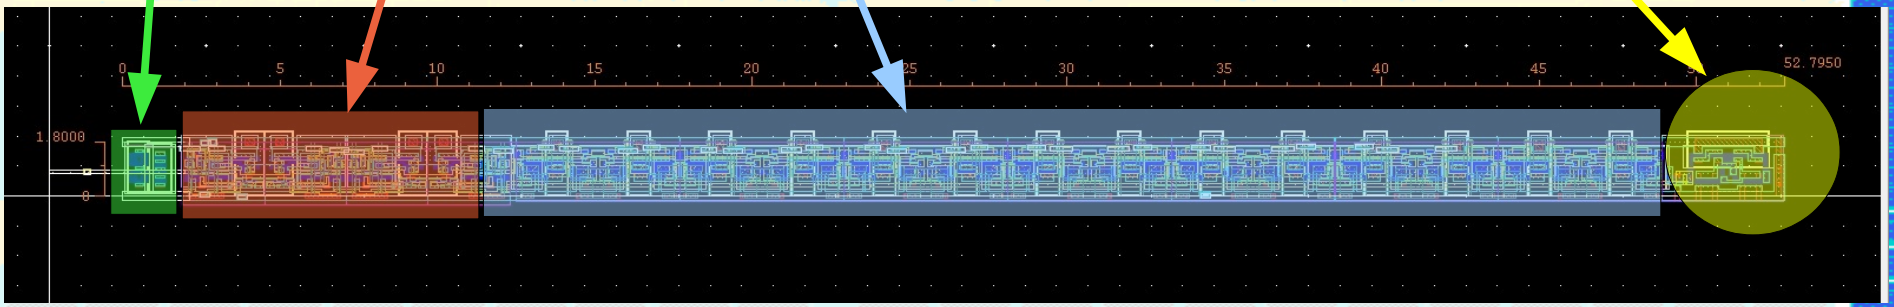

CAM layer timing diagram

Simulation done in nominal conditions:
Transistors models → TT
VDD → 1.2V
Temperature → 27 °C

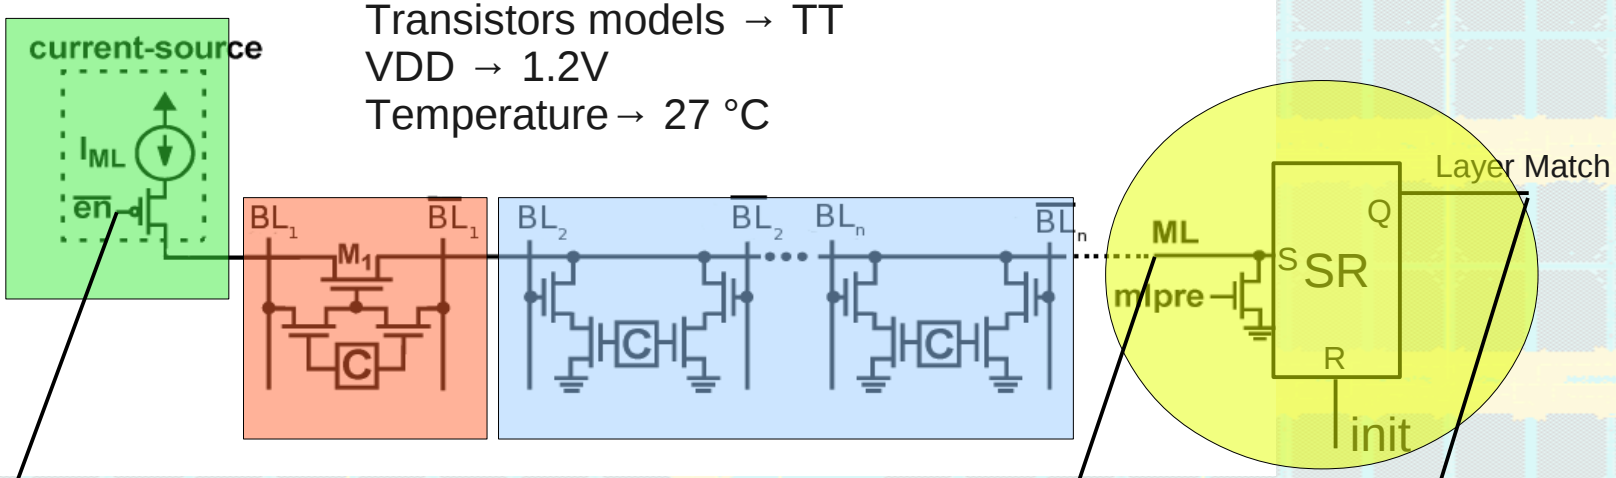

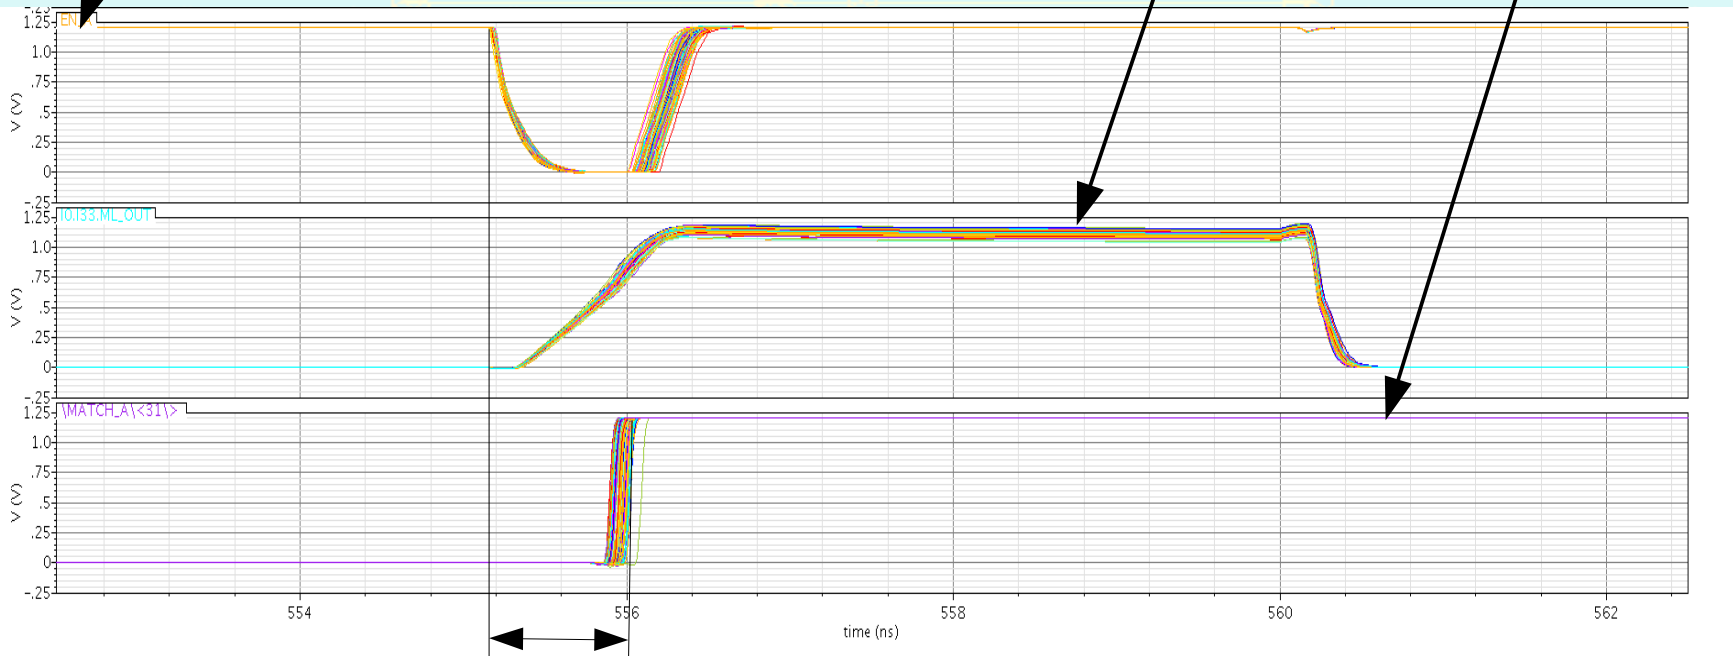

Current race and selective - precharge schemes

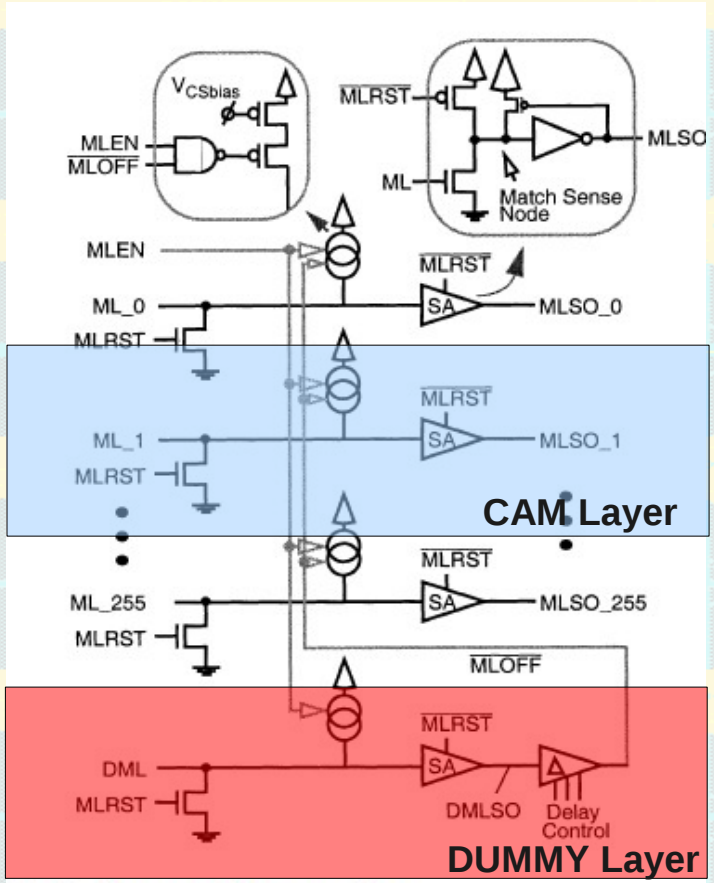

Fig. 5. Current-race ML sensing scheme.

Scheme from: "*A ternary content-addressable memory (TCAM) based on 4T static storage and including a Current-Race sensing scheme*", Ali Sheikholeslamiet Al. IEEE Journal of Solid-State Circuits, Vol. 38, NO. 1, January 2003

NAND type Cell

NOR type Cell

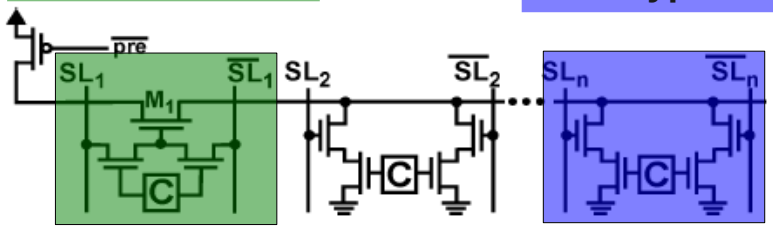

Fig. 16. Sample implementation of the selective-precharge matchline technique [43]. The first cell on the matchline is a NAND cell, while the other cells are NOR cells. Precharge occurs only in the case where there is a match in the first cell. If there is no match in the first cell, the precharge transistor is disconnected from the matchline, thus saving power.

Scheme from: "*Content-Addressable Memory (CAM) Circuits and Architectures: A Tutorial and Survey*", Kostas Pagiamtzis and Ali Sheikholeslami IEEE Journal of Solid-State Circuits, Vol. 41, NO. 3, March 2006

The Full Custom Cell

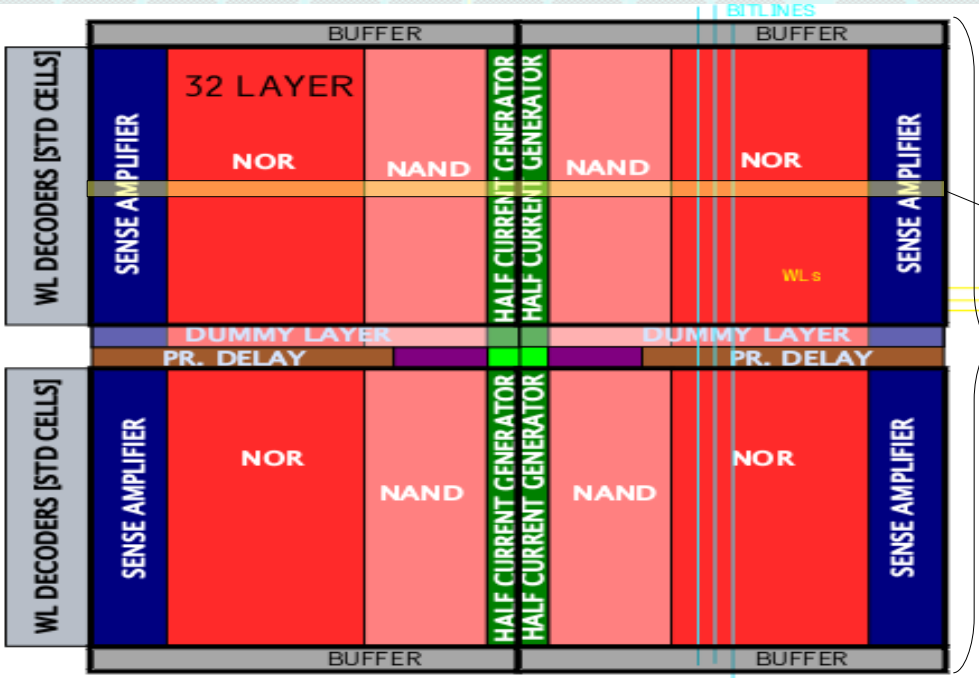

VBIAS REFERENCE
RESISTANCE

Two Layers = ¼ pattern

64 Layers + 1 Dummy layer in the middle

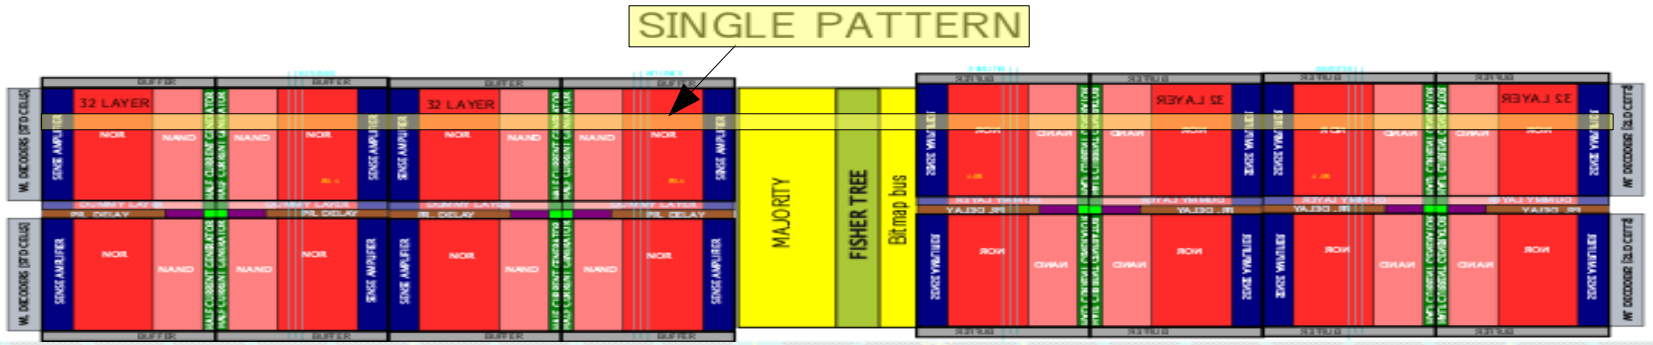

64 patterns (vertically)

Memory Block Layout

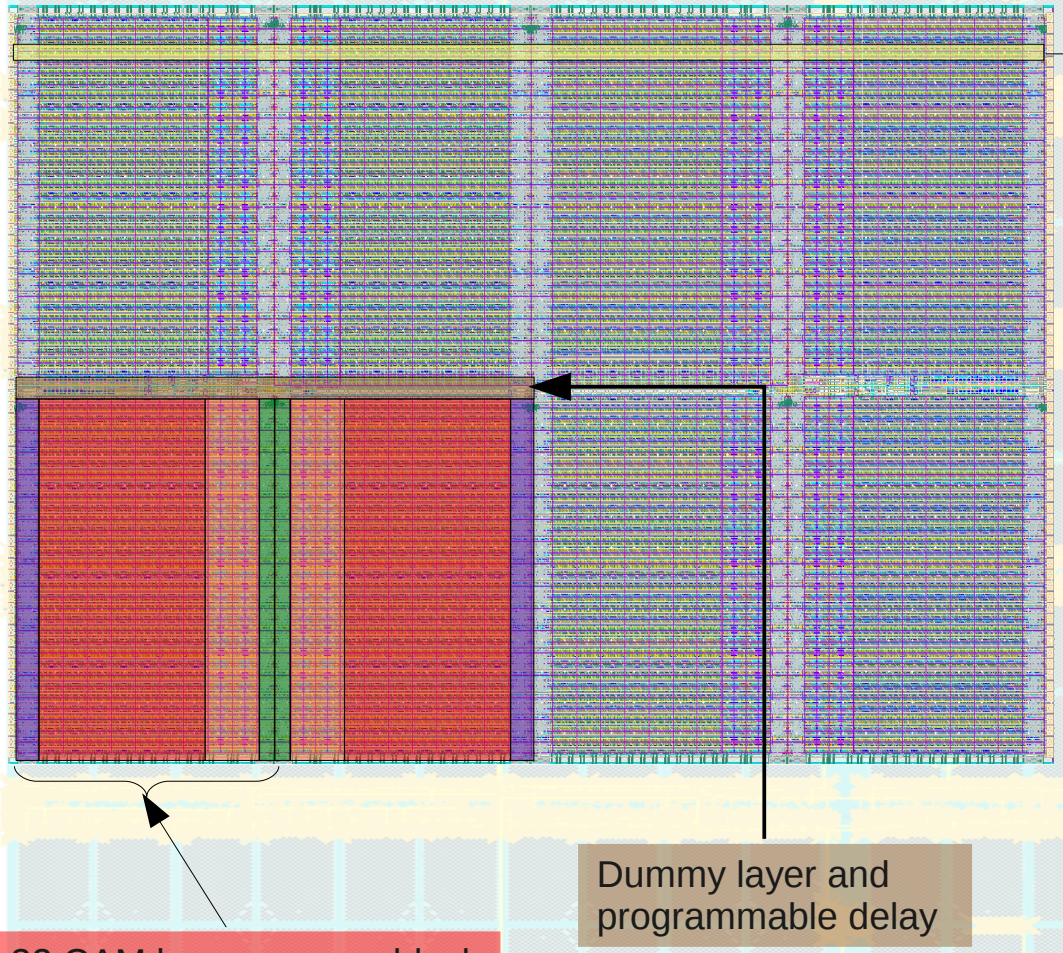

Full custom Layout of 64 x 4 CAM layers (half pattern):w=~226 µm X h=~123 µm

without including
- major logic
- readout logic
- control logic

Six metal layers are used to route signals, power supply and ground.

Bit lines are routed vertically while control lines and memory output are routed horizontally

Chip Layout prototype

The AMchip has an area of 14 mm²

CAM is organized as 22 column x 12 row matrix of full custom memory blocks

Each block is 64 x 8 layers

Between two row of blocks there is the majority logic and the fisher tree made using std. Cells.

In the center there is the control logic made using std. Cells.

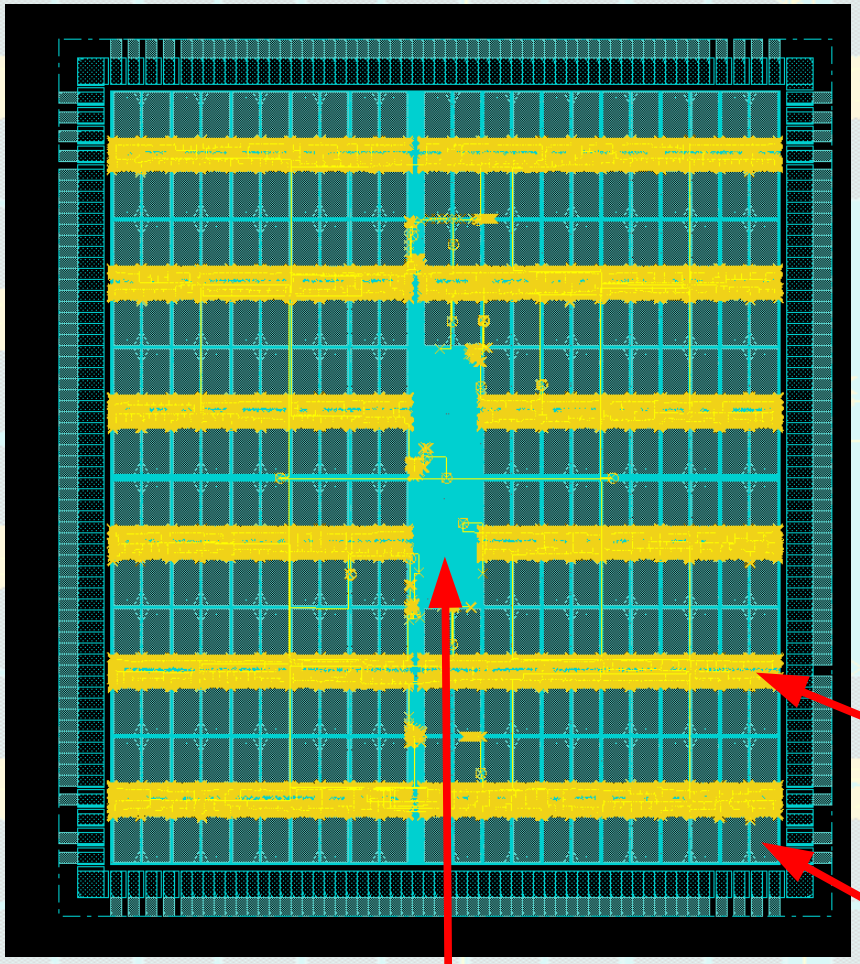

"Variable resolution" in the AMchip

Ternary CAM: feature: Don't Care Bits

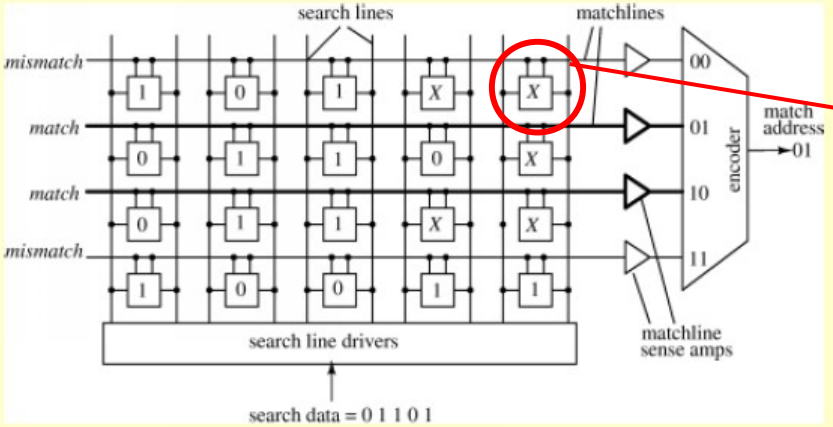

We can use **don't care** on the least significant bit when we want to match the **pattern layer @ Large resolution** or use all the bits to match it @ **Thin resolution**

Coincidence window is programmable layer by layer and pattern by pattern

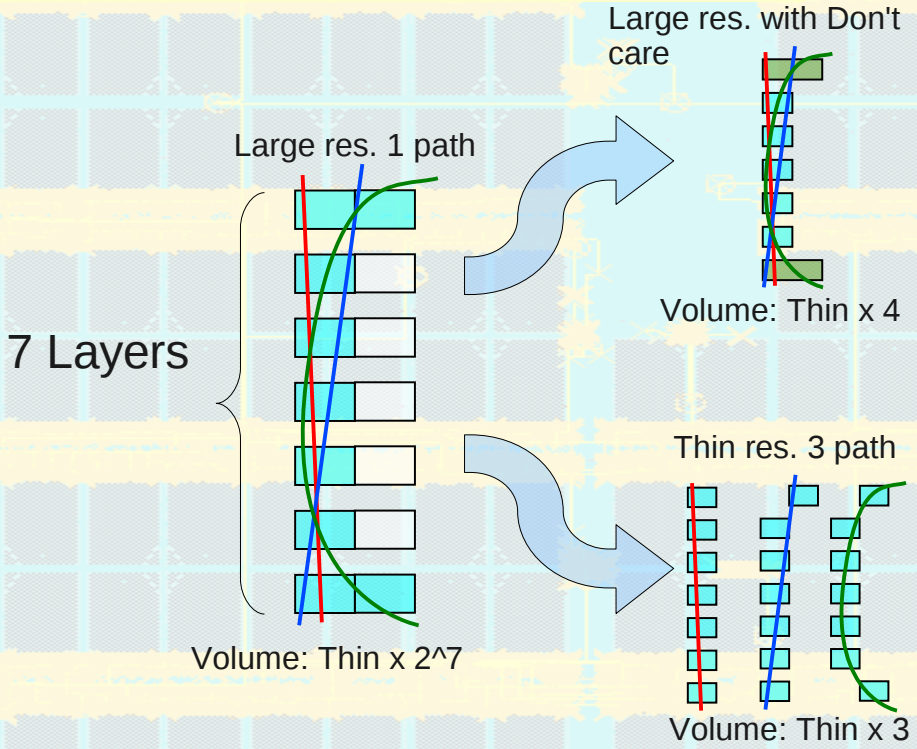

Good rejection but occupy only 1 bank location

Good rejection but occupy 3 locations in the bank

Ternary CAM Cell with two NOR type cells

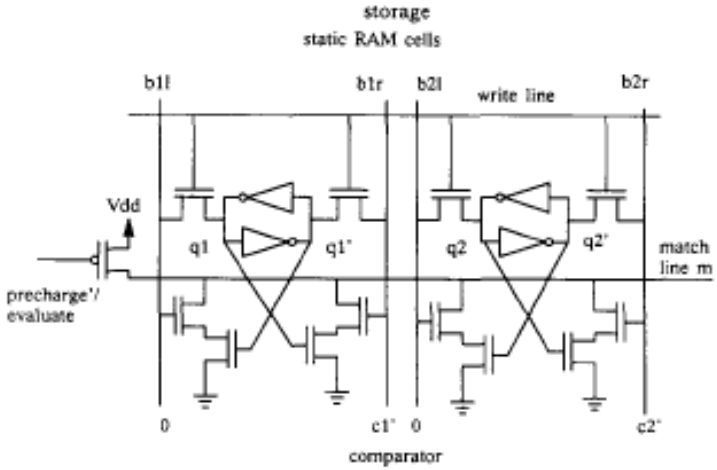

Fig. 8. Two adjacent static binary CAM cells.

storage scheme
stored values

	q1q2
0	01
1	10
*	00

(a)

retrieval scheme
presented values

presented ternary value	encoded value in the bit lines of two binary static CAMS. c1c2	b1l	b1r	b2l	b2r	binary CAM equivalent operation l r
0	01	0	1	0	0	0 M*
1	10	0	0	0	1	M 0
*	11	0	0	0	0	M M

*M is the masking of a bit operation common in commercial binary CAMS.

(b)

Fig. 9. Encoding and retrieval schemes for don't-care in two static binary CAM's cells with masking capability. (a) Encoding scheme. (b) Retrieval scheme.

Images from: "*Encoding Don't Cares in Static and Dynamic Content-Addressable Memories*", Sergio R. Ramirez-Chavez, IEEE Transactions on circuits and system-II: Analog and Digital Signal Processing, Vol. 39 NO. 8, August 1992

Low Power Approach

Technology

We have used *Tsmc 65 nm* low power technology for applications with 1.2V core design, and 3.3V capable I/Os

Power saving design techniques

TABLE IV
SUMMARY OF MATCHLINE SENSING TECHNIQUES

Scheme	ML Energy (fJ/bit/search)	Cycle Time (ns)	Scheme Simplicity	Noise Immunity	Reference(s)
Conventional	9.5	3.9	++	+	[25]
Low-swing[1]	4.2	3.1	−	−	[21], [34], [41]
Current Race	5.5	3.7	+	−	[42]
Selective Precharge[2]	5.6	3.5	+	+	[43]
Pipelining[3]	5.8	3.8	−	+	[49], [50]
Current Saving	4.3	3.7	−−	−	[52], [53]

[1] ML swing of 0.45 V. [2] Single bit used for first segment.
[3] ML divided into three ML segments.

TABLE V
SUMMARY OF SEARCHLINE-DRIVING APPROACHES

Scheme	SL Energy (fJ/bit/search)	Noise Simplicity	Noise Immunity	Reference(s)
Conventional	4.6	++	+	[25]
Eliminating Precharge	2.3	++	+	[21], [42], [52], [53]
Hierarchical Searchlines[1]	1.4	−−	−	[49], [50]

[1] Simulated for a 1024 × 144 CAM.

Tables from: "*Content-Addressable Memory (CAM) Circuits and Architectures: A Tutorial and Survey*", Kostas Pagiamtzis and Ali Sheikholeslami IEEE Journal of Solid-State Circuits, Vol. 41, NO. 3, March 2006

Power consumption

Work is under progress

The best power estimation is done using a layout extracted model of 64 x 8 layer memory block.

The power consumption is evaluated during a read cycle of 800 ns in which progressively all layers match.

We use the nominal simulation condition:
Transistor models : Typical
Power supply : 1.2 V
Temperature : 27 °C
Frequency :100 MHz

In this condition the average power is: **7.63e-04 W** for *each memory block*
For the *total full custom memory blocks* is: **200 mW**

This value does not take into account the standard cells part of the chip.

AM chip status

Completed

- Full Custom memory block layout
- Floor plan of the complete chip
- Pin placement
- Simulation of 2 complete memory blocks with extracted layout model
- Placing and routing of the std. Logic part
- Creation of a memory block verilog model for full chip simulation

Work in progress

- Refinement of the standard cells placing and routing to increase the number of double vias and improve timing
- Improvement of the verilog memory core model to obtain a more reliable model

Future plans

2.5D??? What is that?

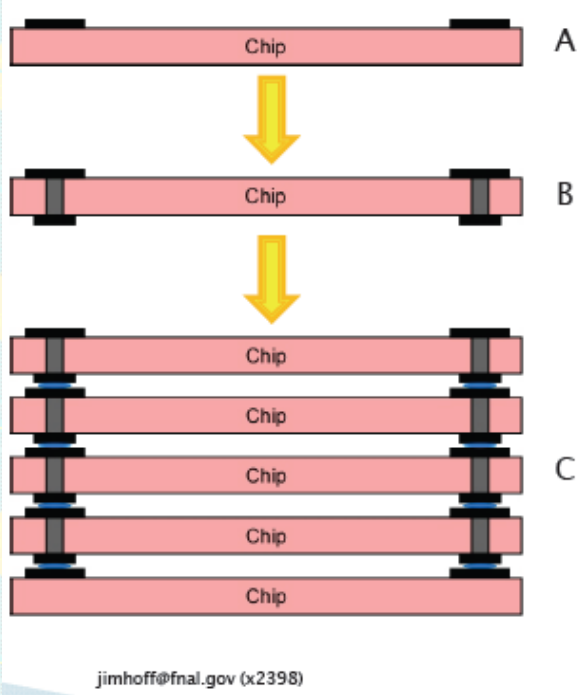

- Less expensive.
- Simpler technology.
- (Possibly?) Greater reuse of existing technology because you are not specifically designing a new 3D chip (maybe).

Future plans (Fermilab)

The True 3D: 1 tier/ Layer + 1 control tier

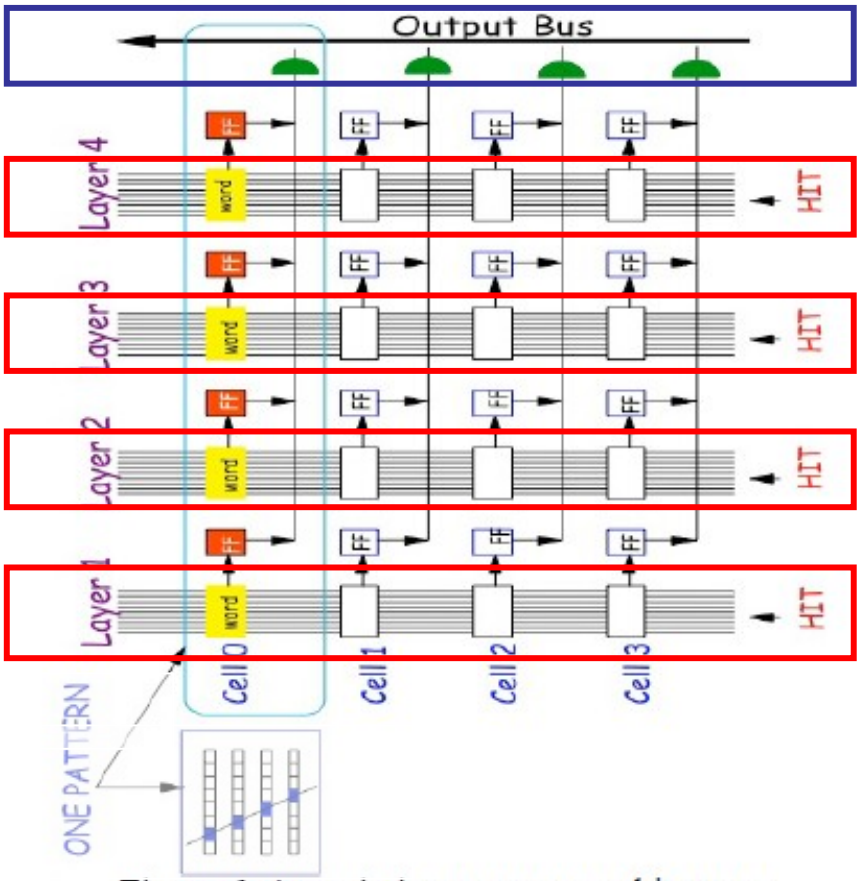

Figure 6. Associative memory architecture

Control Tier

Tier 4

Tier 3

Tier 2

Tier 1

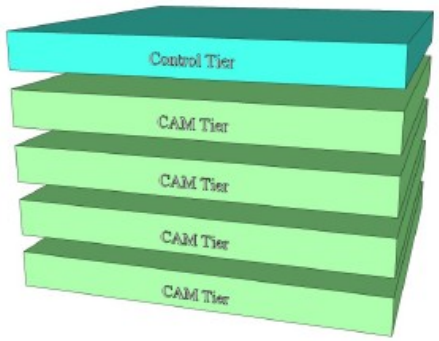

Figure 8. Control and CAM tiers for 3D implementation of VIPRAM.

Summary

In this talk we have shown:

- How the associative memory works and how it inserts in the ATLAS acquisition.
- How to implement power saving architecture and full custom design to gain in memory density
- The layout of the full custom CAM block and the full memory architecture and floorplan.
- The current project status and future development.

AMchip Summary
8k patters 8 layers (each layer 12bits+3ternary bits)
14mm^2 at TSMC 65nm
Working frequency >100MHz

SPARE SLIDES

Expected number of roads

We simulate WH events with different numbers of pile-up events

Average number of roads/AMBoard

N pileup (event complexity)	**N patterns**	**N roads Large resolution**	**N roads Large res. With 2 don't care per layer**	**N roads half resolution**
17.6 evts	5 Mpatterns	5040	**959**	1060
40 evts	80 Mpatterns	37000	**6500**	5720
75 evts	380 Mpatterns	53500	**8250**	5950
			(+ 20% hardware)	**(5 to 10X hardware)**

We gain a lot in resolution adding only few space in hardware

Matchline sensing approach

Conventional (Precharge-High) Matchline Sensing

the overall matchline power consumption of a CAM block with *w* matchlines

$$P_{ML} = w C_{ML} V_{DD}^2 f$$

Where CML is the matchline capacitance, VDD is the power supply voltage and f is the frequency.

Current-Race Scheme

This scheme precharges the matchline low and evaluates the matchline state by charging the matchline with a current *IML* supplied by a current source. the power consumption of a match and a miss are identical, the overall power consumption for all *w* matchlines is

$$P_{ML} = w C_{ML} V_{DD} V_{tn} f$$

Where Vtn is the latch treshold.

Selective-Precharge Scheme

performs a match operation on the first few bits of a word before activating the search of the remaining bits

Bit lines driving approach

Conventional Approach

The equation for the dynamic power consumption of the bitlines is:

$$P_{BL} = n C_{BL} V_{DD}^2 f$$

Where C_{BL} is the total capacitance of a single bit line, n is the total number of bit lines pairs, and V_{DD} is the power supply voltage.

Eliminating bit line precharge

We can save power eliminating the BL precharge phase.
The matchline sensing schemes that precharge the matchline low eliminate the need for BL precharge.
The equation for the reduced power is

$$P_{BL} = 1/2\, n C_{BL} V_{DD}^2 f$$

We have obtained a 50% reduction in bit line power