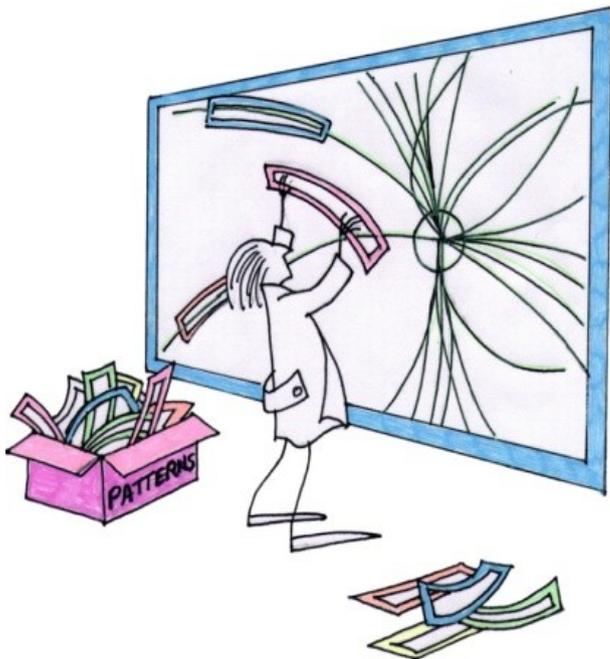


AMchip04: a new generation associative memory chip for HEP applications

F. Crescioli



OUTLINE

- Online tracking at Hadron Colliders
- SVT/FTK algorithm
- Fast pattern recognition with dedicated hw
- AMchip history
- AMchip04: the new FTK variable resolution associative memory chip
- Beyond AMchip04
- Conclusions

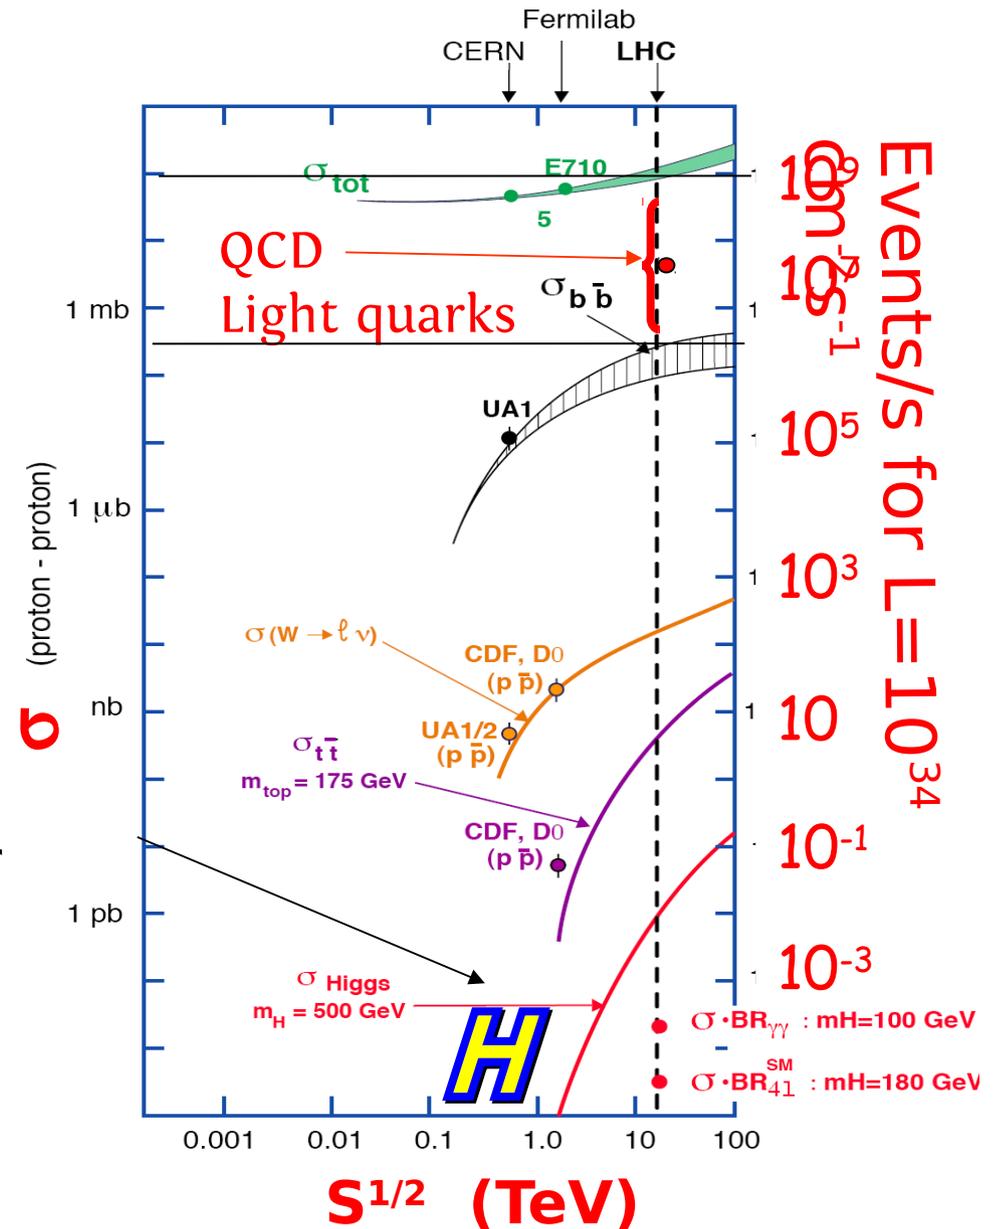
ONLINE TRACKING @ HADRON COLLIDERS

Search for rare SM or predicted BSM processes push the collider's intensity to new frontiers

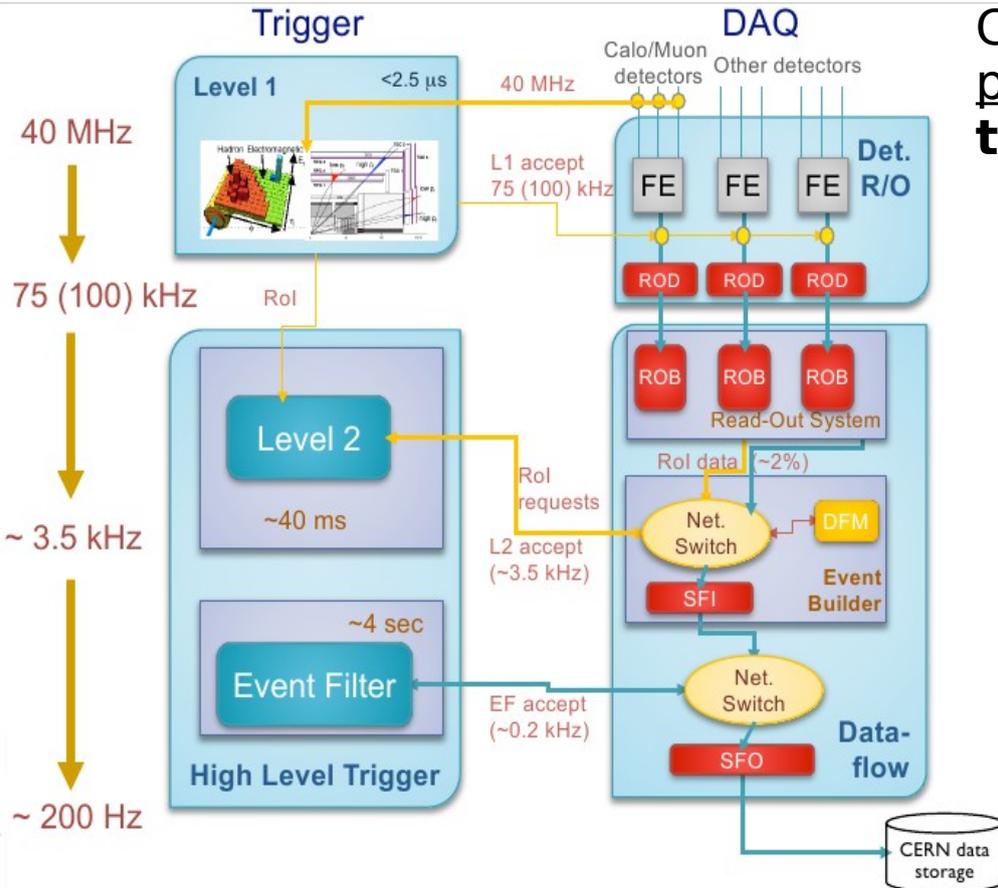
Rare processes are overwhelmed by well-known processes

- 8-9 orders of magnitude between Higgs and total cross-section
- Need to prioritize the physics output and leave flexibility for the unexpected

Need **sophisticated trigger techniques** to maximize the bandwidth for interesting events



ONLINE TRACKING @ HADRON COLLIDERS

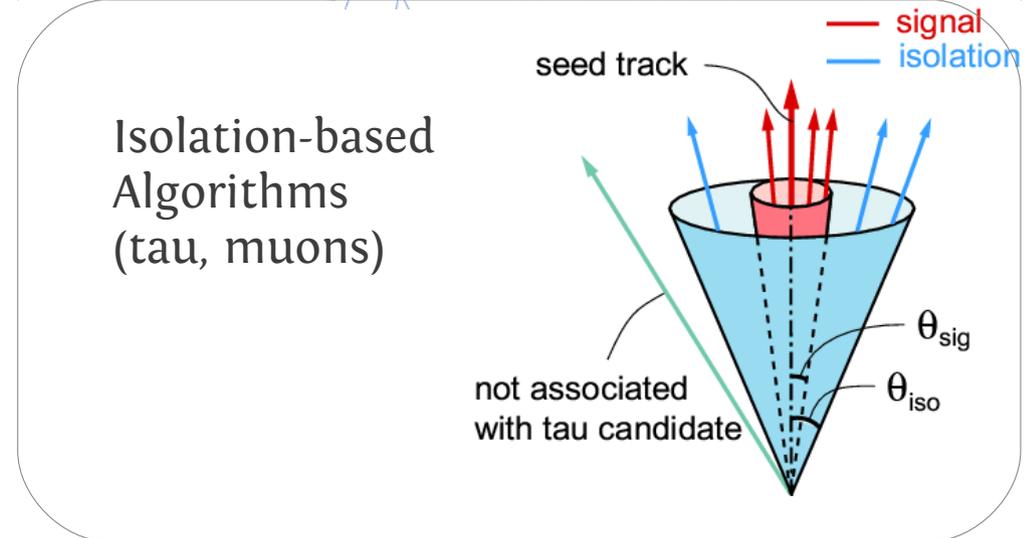
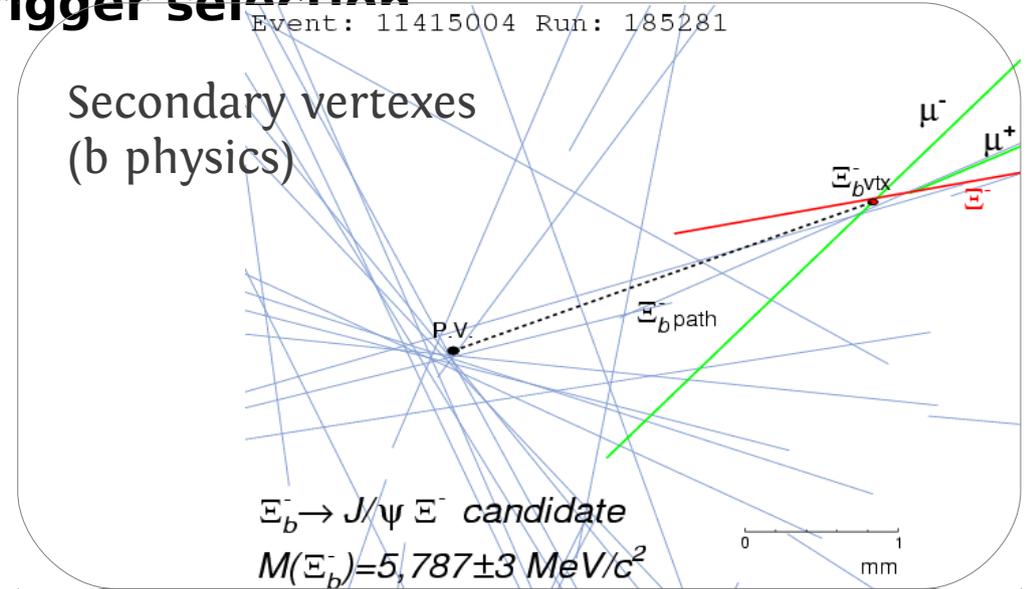


Track reconstruction is a complex and time-consuming computing task to perform in a short time budget.

Hardware processors have been conceived for this task.

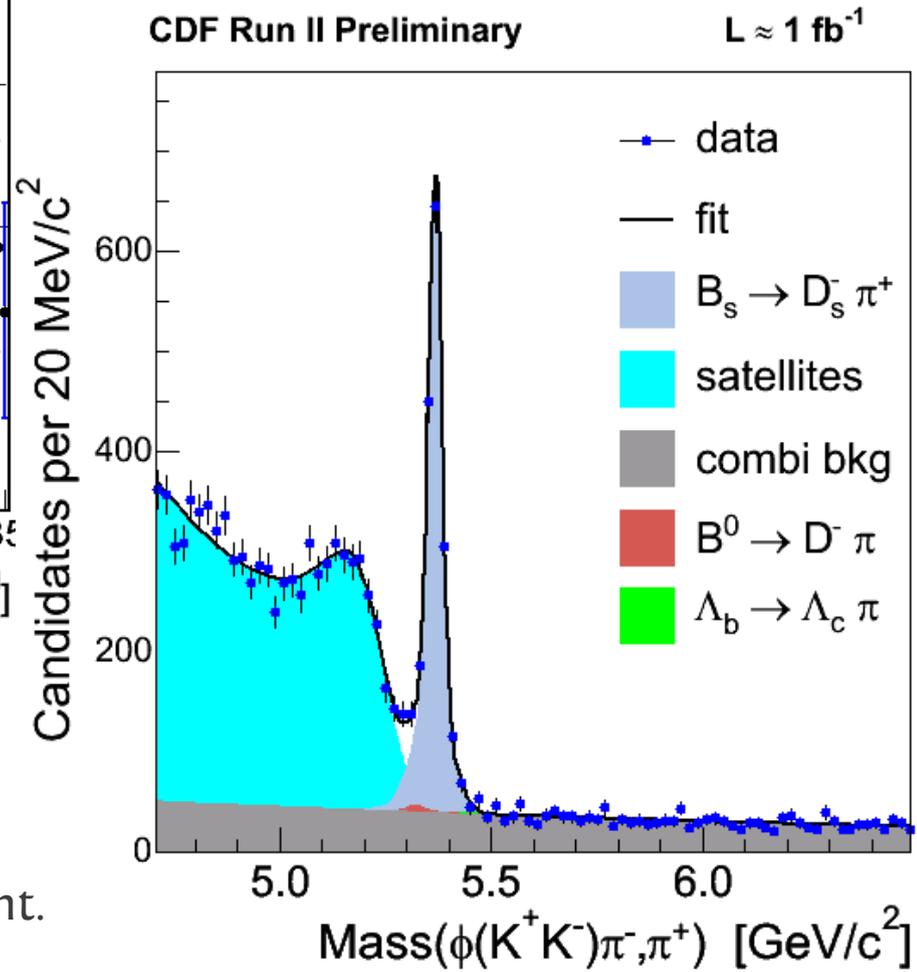
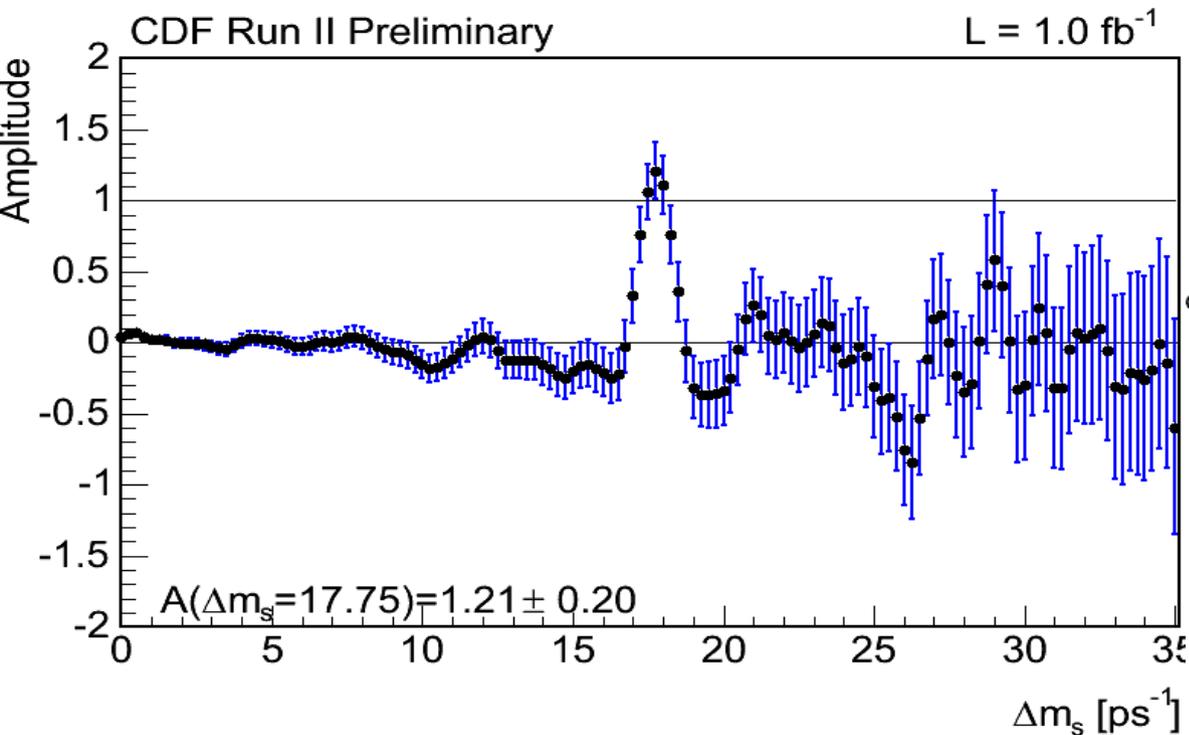
Examples: SVT (CDF), FTK (ATLAS)

Charged particle trajectories provide very powerful information for **sophisticated trigger selection**



ONLINE TRACKING @ HADRON COLLIDERS

SVT @ CDF



Online tracking (XFT @ L1 and SVT @ L2) has been a key element of many CDF analyses.

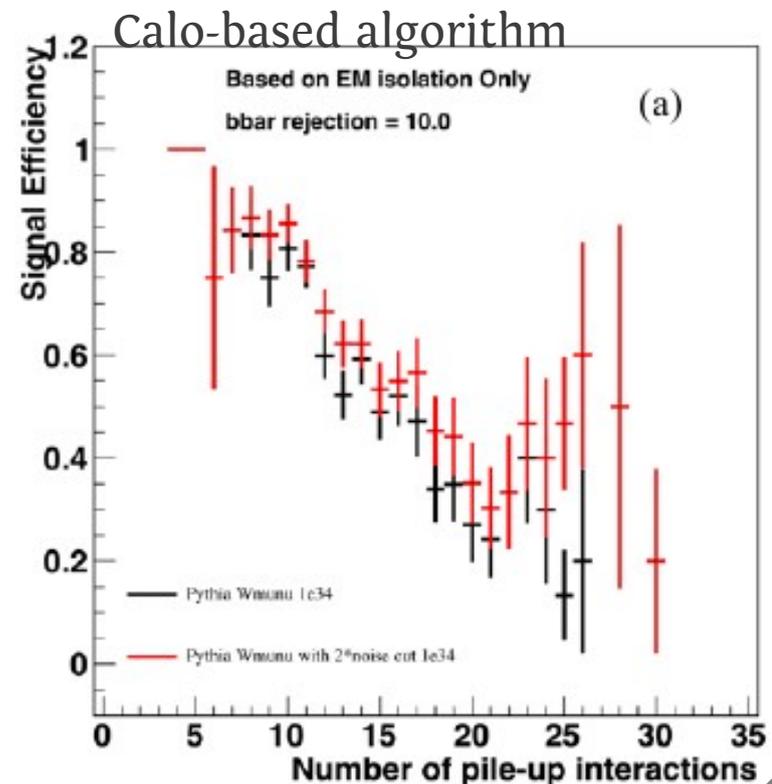
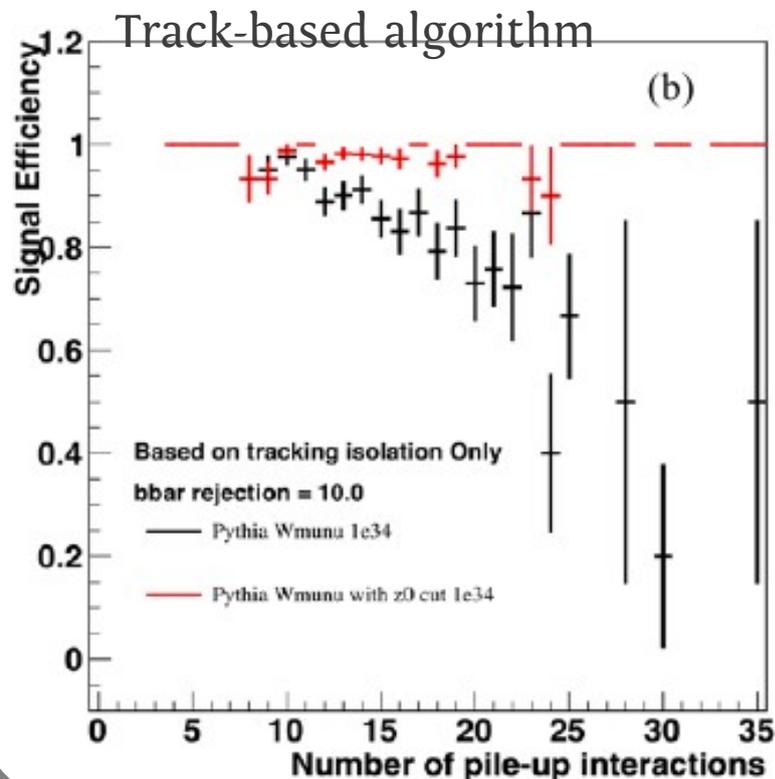
Example: the B_s oscillation frequency measurement. Many hadronic B_s collected thanks to SVT-based two-track trigger.

ONLINE TRACKING @ HADRON COLLIDERS

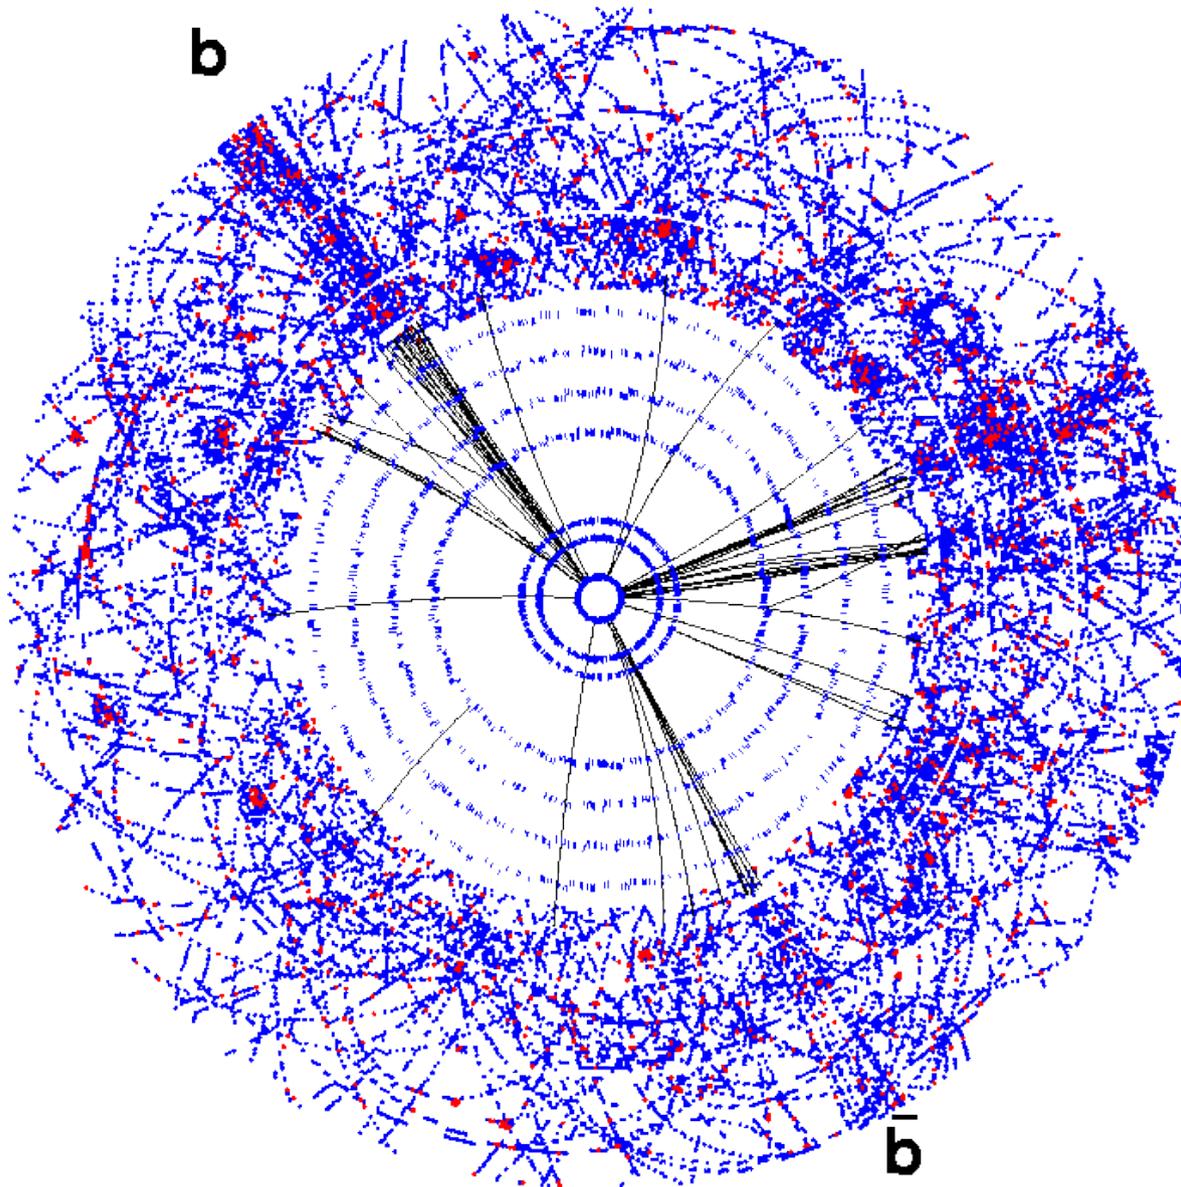
FastTrackr (FTK) has been recently approved by the ATLAS collaboration as an upgrade for its Level-2 trigger.

Online tracking is very important for high-luminosity, FTK is designed to handle track reconstruction up to $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$

Isolated muon trigger efficiency wrt pile-up

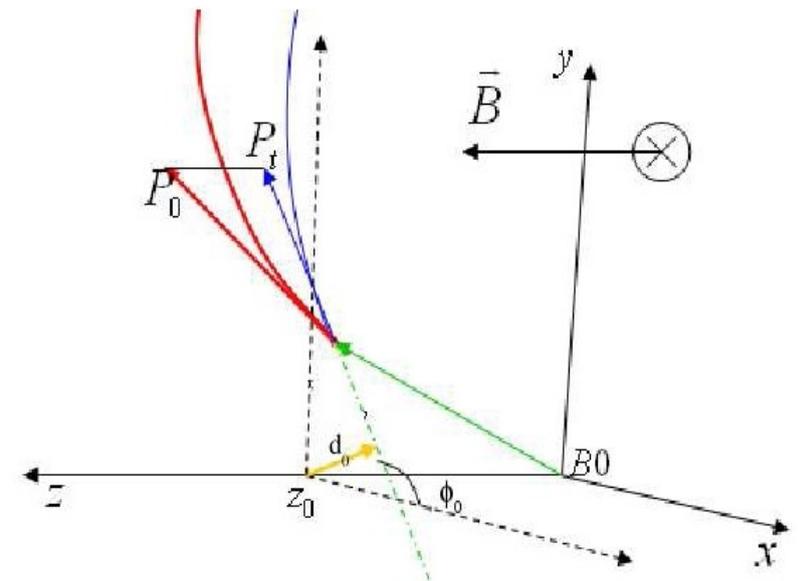


ONLINE TRACKING @ HADRON COLLIDERS

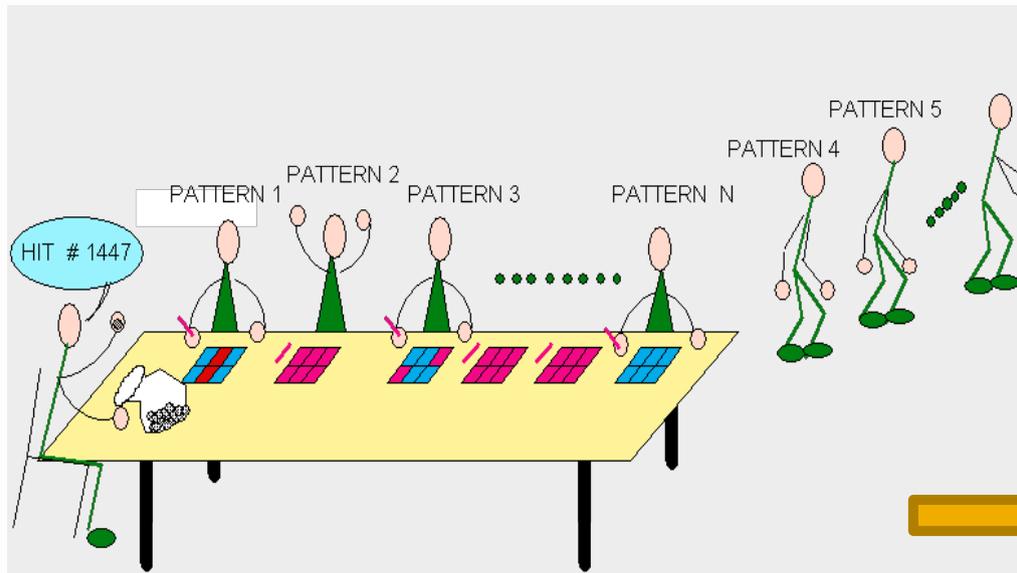
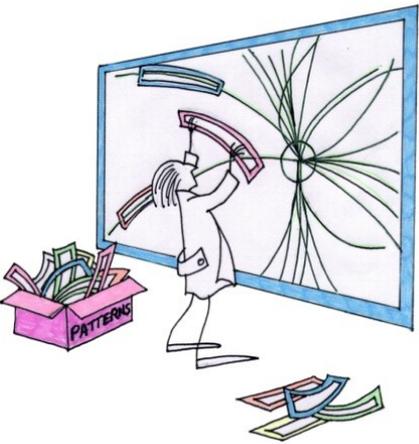


A Tracking processor must:

- recognize collection of hits belonging to a particle trajectory (pattern recognition)
- reconstruct particle trajectory parameters (origin, momentum) (track fit)



FTK/SVT ALGORITHM



An associative memory (AM) chip allows a full parallel search in the pattern bank.

Similar to commercial CAMs + extended functionalities.

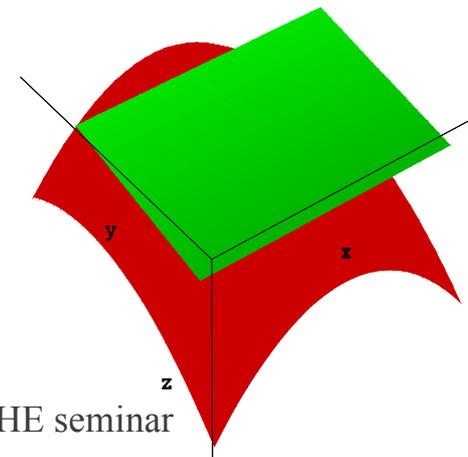


In FTK, tracking is divided into two separate steps

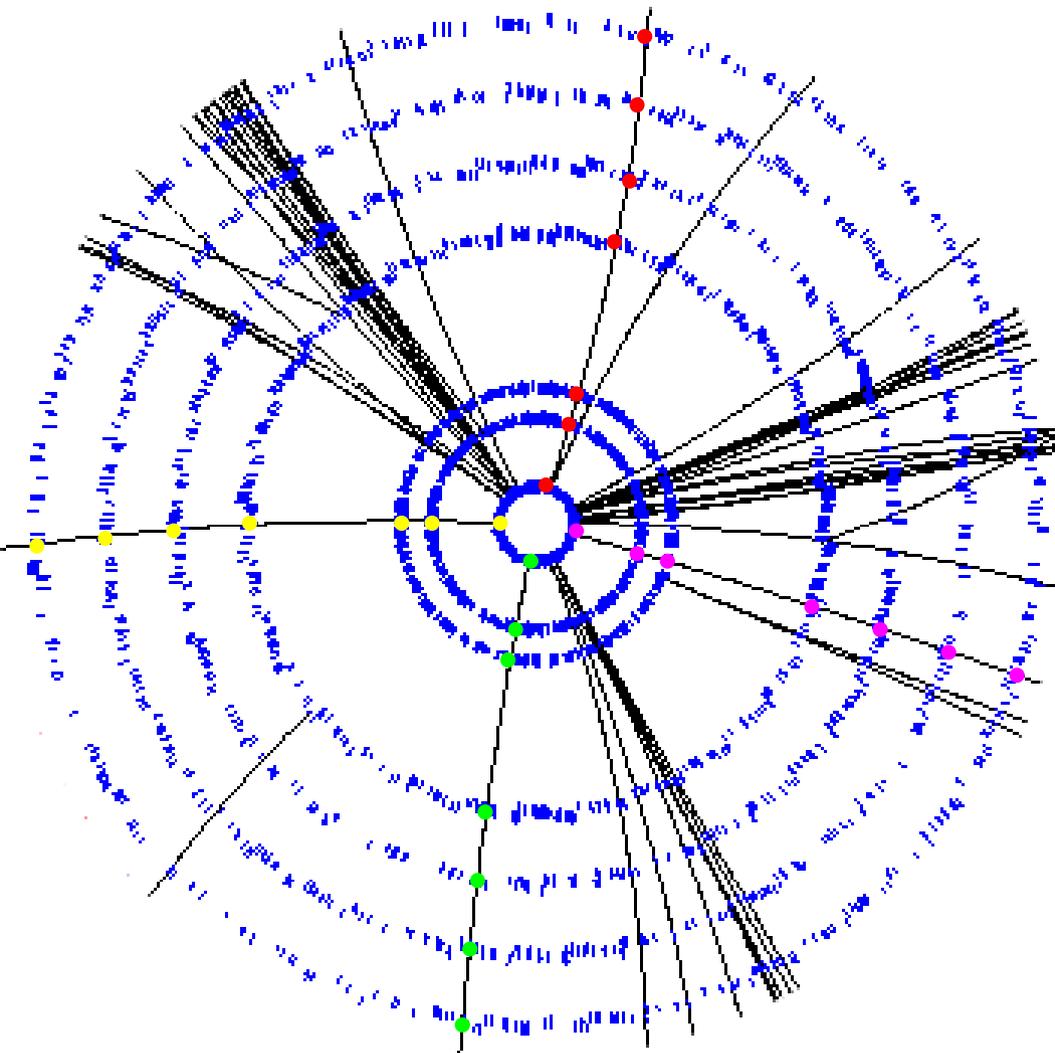
- The pattern matching uses a large bank of patterns and special hardware (associative memory - AM)
- The track's parameters are evaluated using a linear Principal Component Analysis algorithm (DSP MACC units in FPGA)

(j.nima.2003.11.078)

$$p_i = \sum C_{ij} \cdot x_j + q_i$$

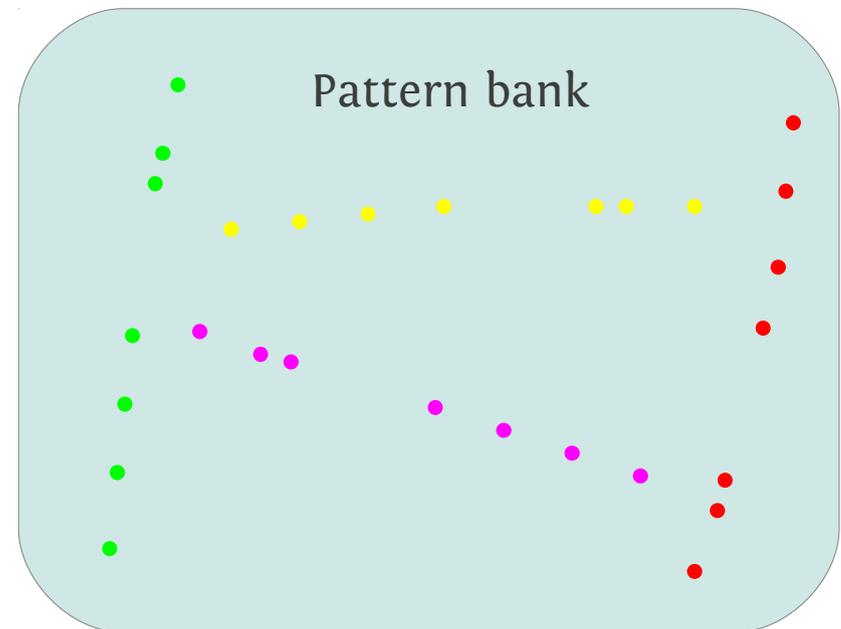


FAST PATTERN RECOGNITION WITH DEDICATED HW



Compare the event with a pre-computed collection of template tracks (pattern bank).

- 1) Find templates that matches hits of the event
- 2) Avoid computing combination of hits → time consuming, does not scale with occupancy



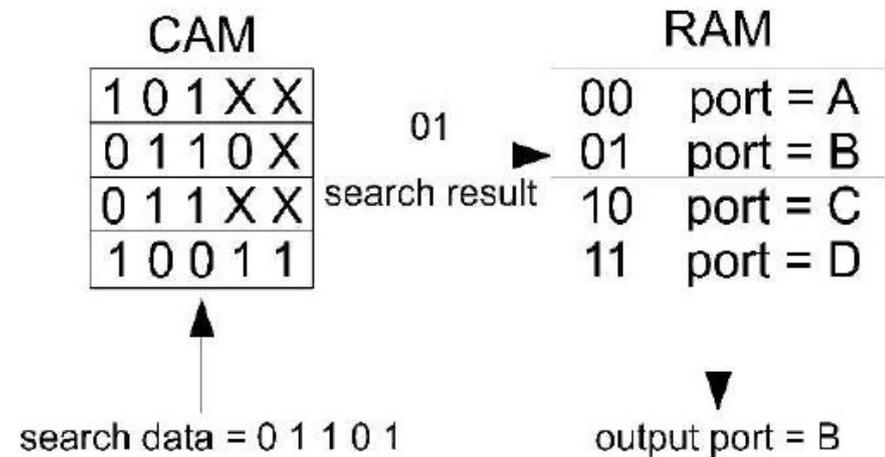
FAST PATTERN RECOGNITION WITH DEDICATED HW

Random Access Memory (RAM)

- Internal storage
- Input: address Output: data

Content Addressable Memory (CAM)

- Internal storage and built-in compare logic
- Input: search data Output: address
- Ternary logic: 0, 1, X “don't care”

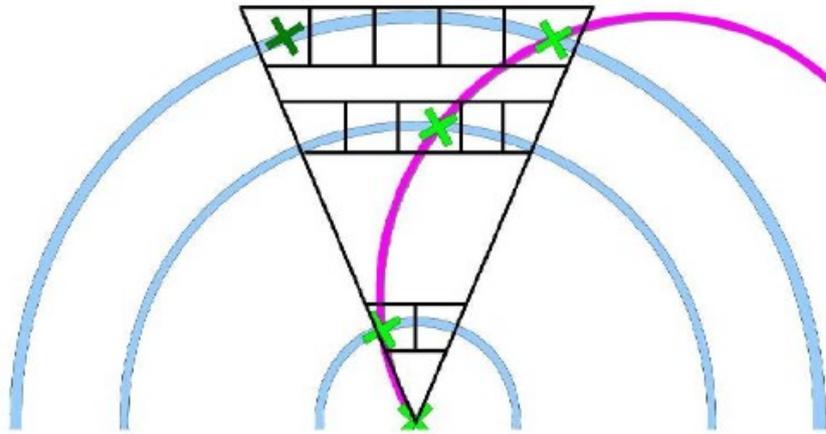


Commercially available

CAM are widely used in network hardware applications.

(also used inside processors: cache, mmu)

FAST PATTERN RECOGNITION WITH DEDICATED HW



Data word: one bit per possible hit location, concatenated layers

Store template tracks using 1s and Xs

Pro: possible with commercially available CAM chips

Cons:

- Very long words in modern silicon detector applications (ie. an ATLAS SCT module → 768 bits)

- Many templates for redundancy (taking into account layer efficiency, store templates with missing layer)

pattern

template

X	X	X	X	1
X	X	1	X	X
1	X			

Template: XXXX1XX1XX1X

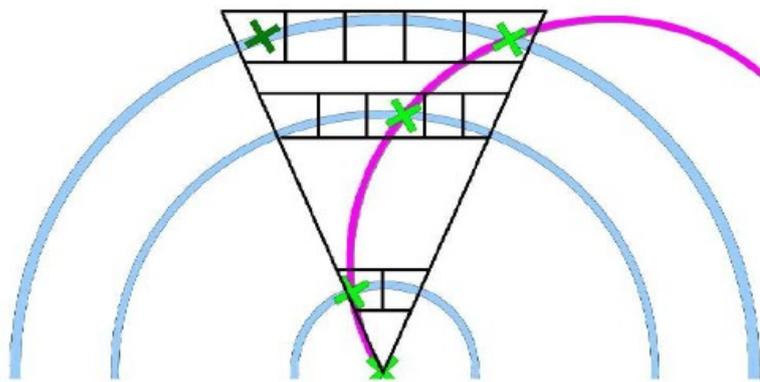
Search data: 100010010010

FAST PATTERN RECOGNITION WITH DEDICATED HW

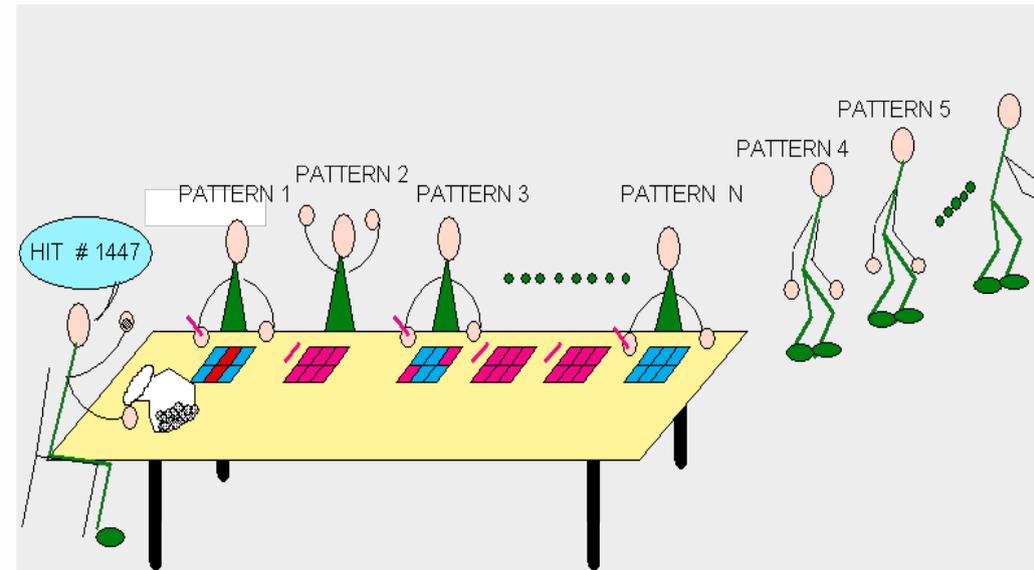
SVT Associative Memory idea: build a custom kind of CAM more suitable for HEP applications.

Key features:

- store templates in encoded form, subdivided by detector layers
- take into account efficiency without template redundancy
- do not wait for complete event readout



	Lay0	Lay1	Lay2
Template:	01	10	000
Search data:	01	10	000 100

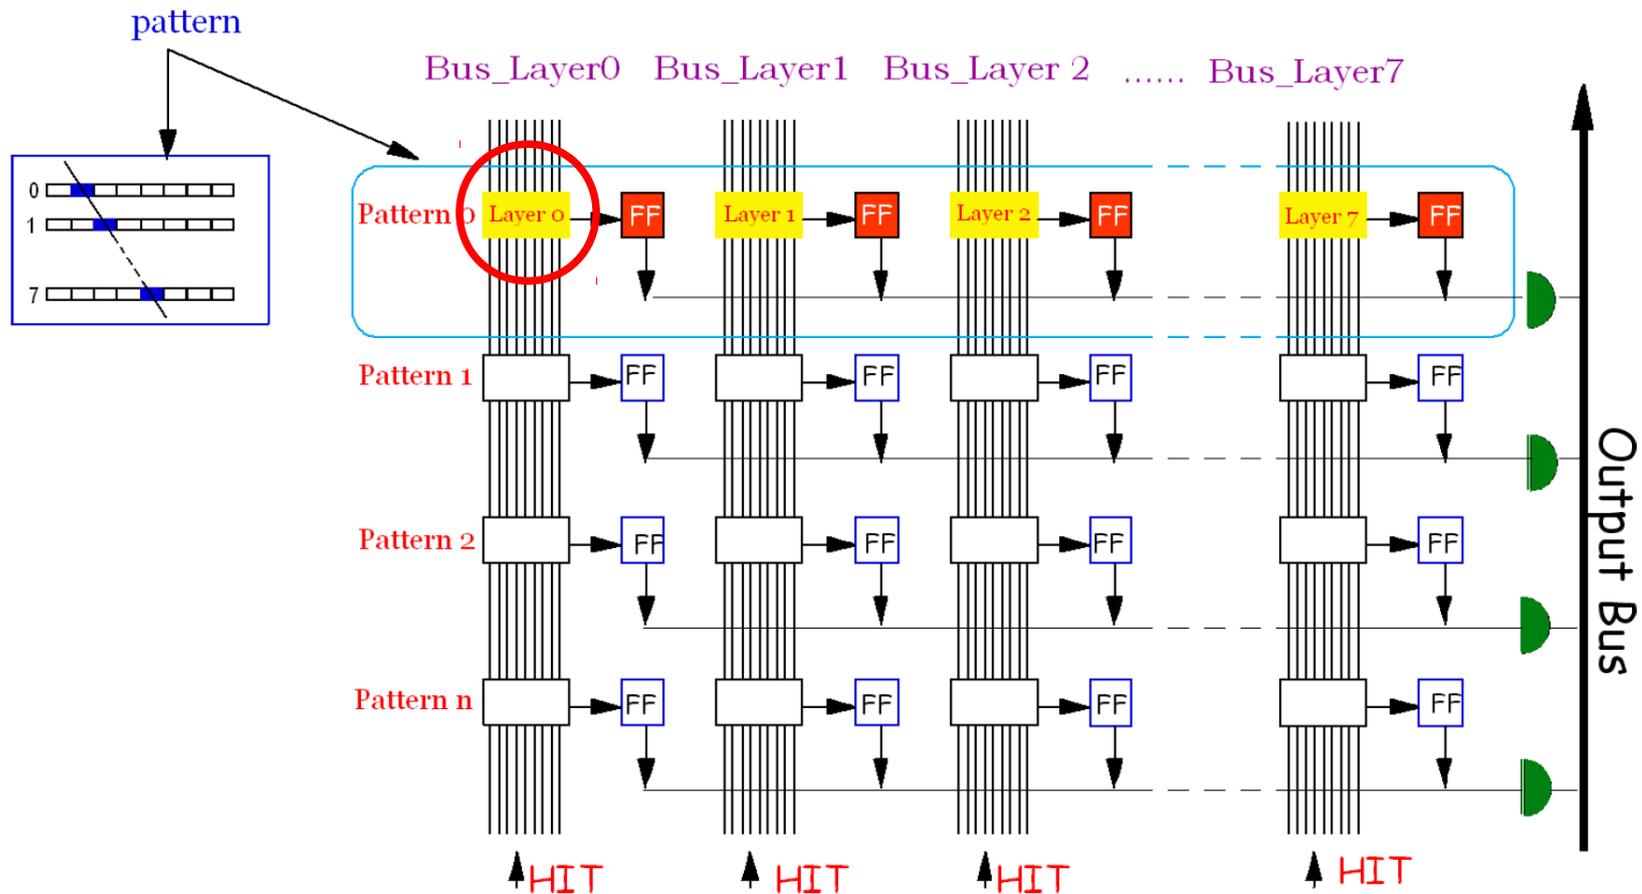


Advanced built-in compare logic:

“Smart” patterns, like bingo players

AMCHIP FEATURES

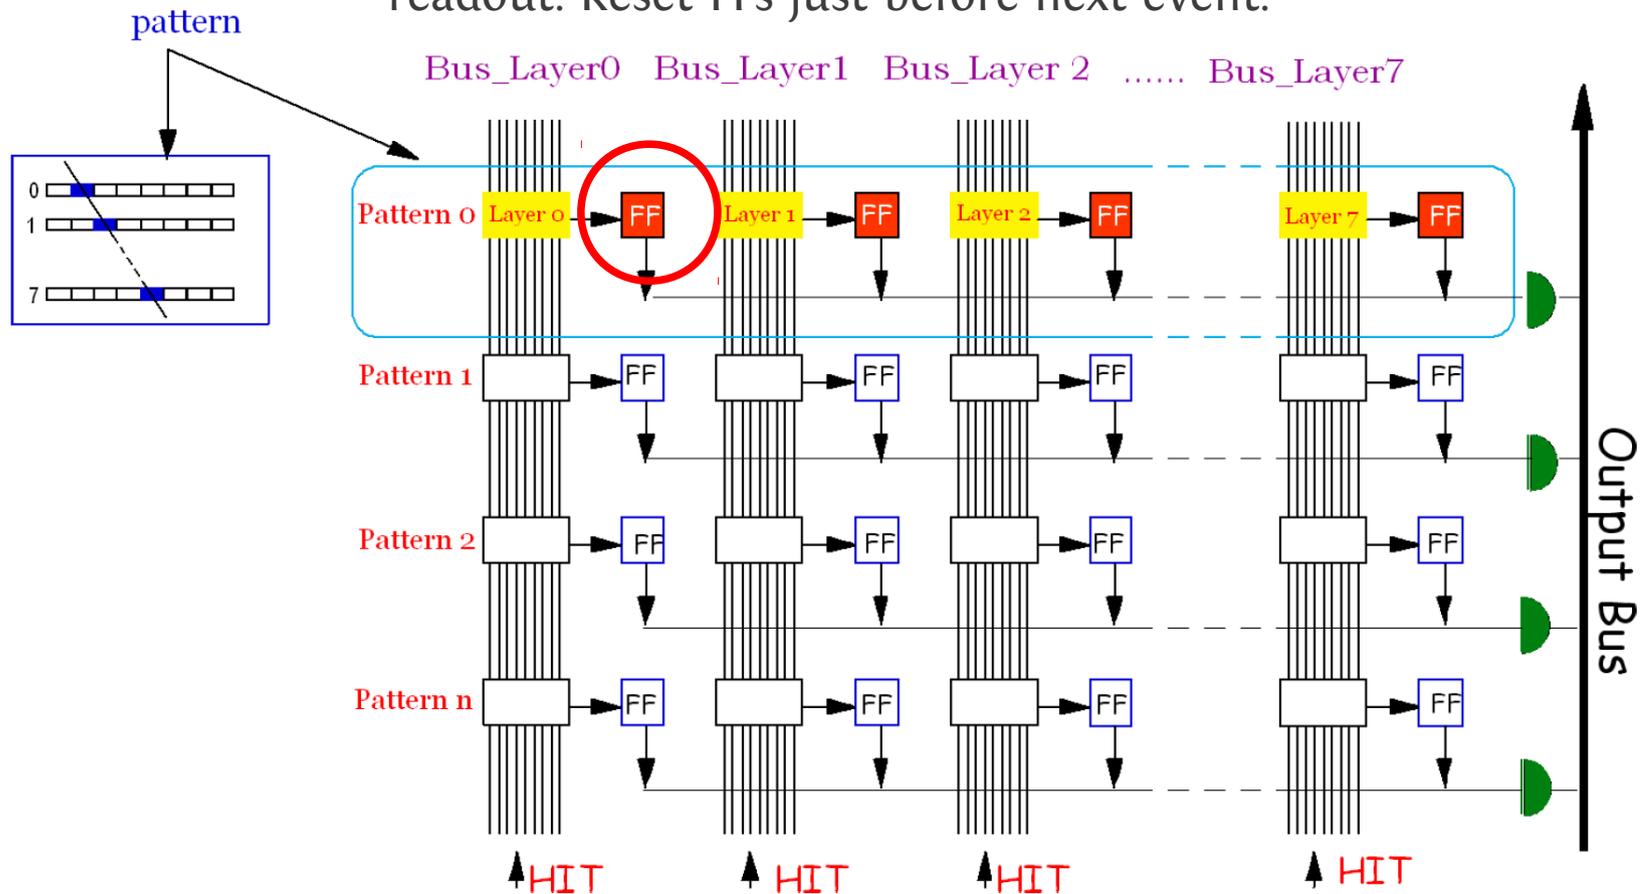
Patterns are divided per layer, a partial CAM for each one



Hit is stored in encoded form. I.e. 15-bit partial CAM can store hit position with 2^{15} resolution
 Wider detector area covered with respect to standard CAM logic.

AMCHIP FEATURES

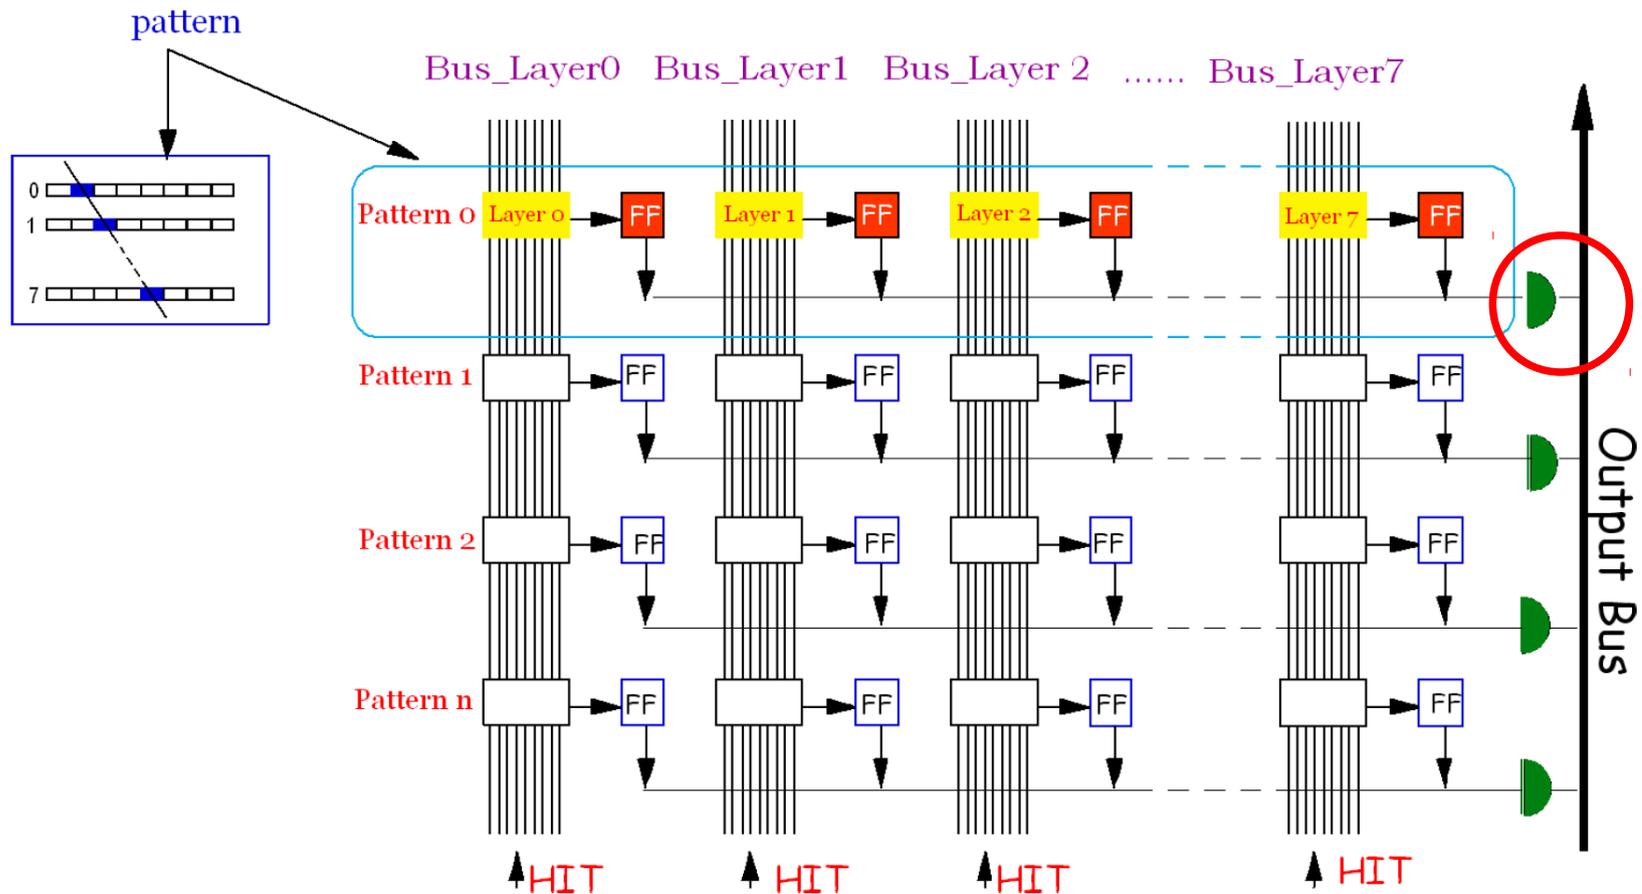
Partial CAM results is stored in a local memory (FF). Compare the event to the templates during detector readout. Reset FFs just before next event.



Separate input buses. No need to build the search data word from different layers, just send hits to the CAM when they are available.

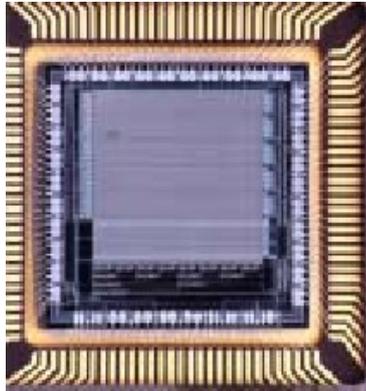
AMCHIP FEATURES

Majority logic: match even if not all layers are matched.
Output the bitmap of matched layer for post-processing.



Majority logic is programmable and dynamical. Lowering threshold during readout is possible for ordered output (full match, one missing, two missing, ...)

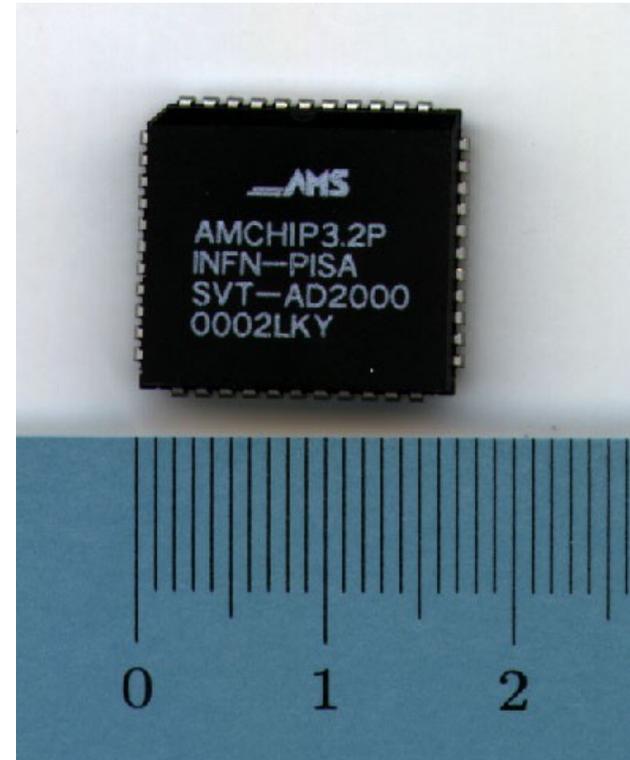
AMCHIP HISTORY



Full-custom VLSI chip
700 nm technology

128 patterns, 6 x 12 bit

Working @ 40 MHz



Development started in the 90s

~5000 chips produced for SVT in 2000

Used by SVT from 2001 to 2006

The development and production proved to be quite hard. many iterations of prototypes, debug and simulations were needed.

What to do for the SVT upgrade?

First article:

S.R. Amendolia, S. Galeotti, F. Morsani, D. Passuello, L. Ristori, N. Turini **The AMchip: a VLSI associative memory for track finding** Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, Volume 315, Issues 1-3, 1 May 1992, Pages 446-448

AMCHIP HISTORY



In the late 90s a functionally compatible version of the full-custom AMchip was developed using programmable devices.

Using a Xilinx FPGA in 350 nm technology it was achieved the same pattern density (128 patt/chip) as the full-custom version.

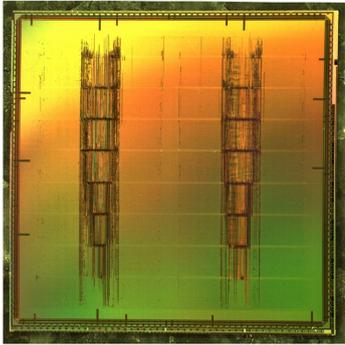
The FPGA version was flexible and a lot faster to develop and debug ...

... but FPGA industry evolved in unpredicted ways (changes in packages, CLB technology, ...) and future versions couldn't achieve required pattern density for the SVT upgrade.

AM FPGA ideas & firmware has been used in low density applications that required flexibility for the SVT upgrade (RoadWarrior)

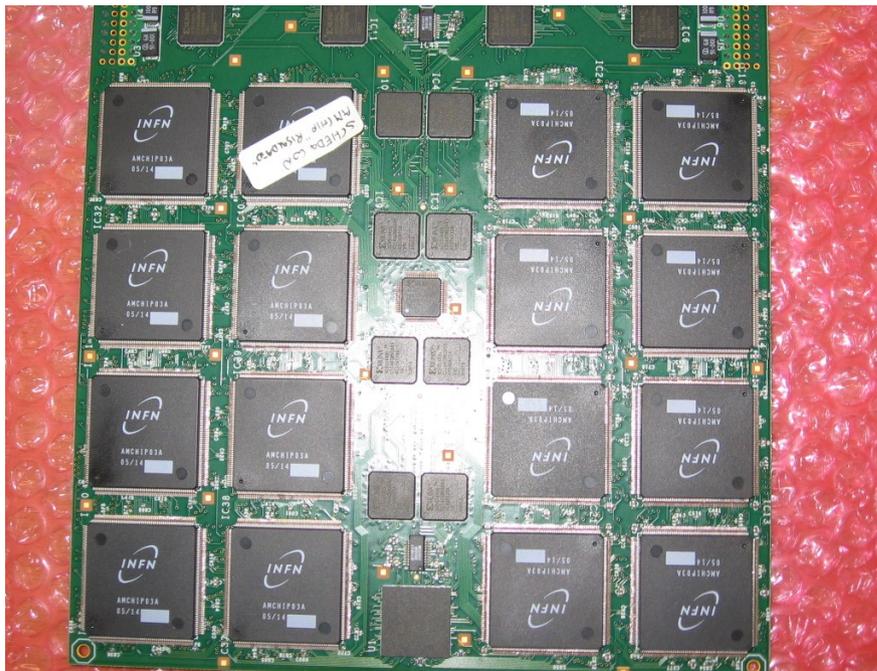
A. Bardi, S. Belforte, A. Cerri, M. Dell'Orso, S. Donati, S. Galeotti, P. Giannetti, , A. Leger, E. Meschi, F. Morsani, D. Passuello, G. Punzi, L. Ristori, T. Speer, F. Spinella, X. Wu **A programmable associative memory for track finding** Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment Volume 413, Issues 2-3, 21 August 1998, Pages 367-373

AMCHIP HISTORY



SVT Upgrade chip AMchip03 (2006): **Standard Cell 180 nm**
5000 pattern/chip for 6-layer patterns,
2500 pattern/chip for 12-layer patterns
40 MHz clock

“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



16x AMchip03 on a LAMB mezzanine

Standard Cell approach was the right compromise between development speed and reliability and pattern density.

2 years from project approval to Installed version

64 chips * 24 boards running in the upgraded SVT at CDF since 2006 until today (literally)

AMCHIP04

- Major improvements:
 - ~80k patt/chip with 1.2x1.2 cm² area
 - Low power consumption
(128 chips/VME board @ 100 MHz)
 - 8 input buses
- New features:
 - Ternary logic for variable resolution
 - Pattern output while next event is loaded
- Technology TSMC 65nm Low Power
- First prototype: 8k patterns, 13x13 mm²

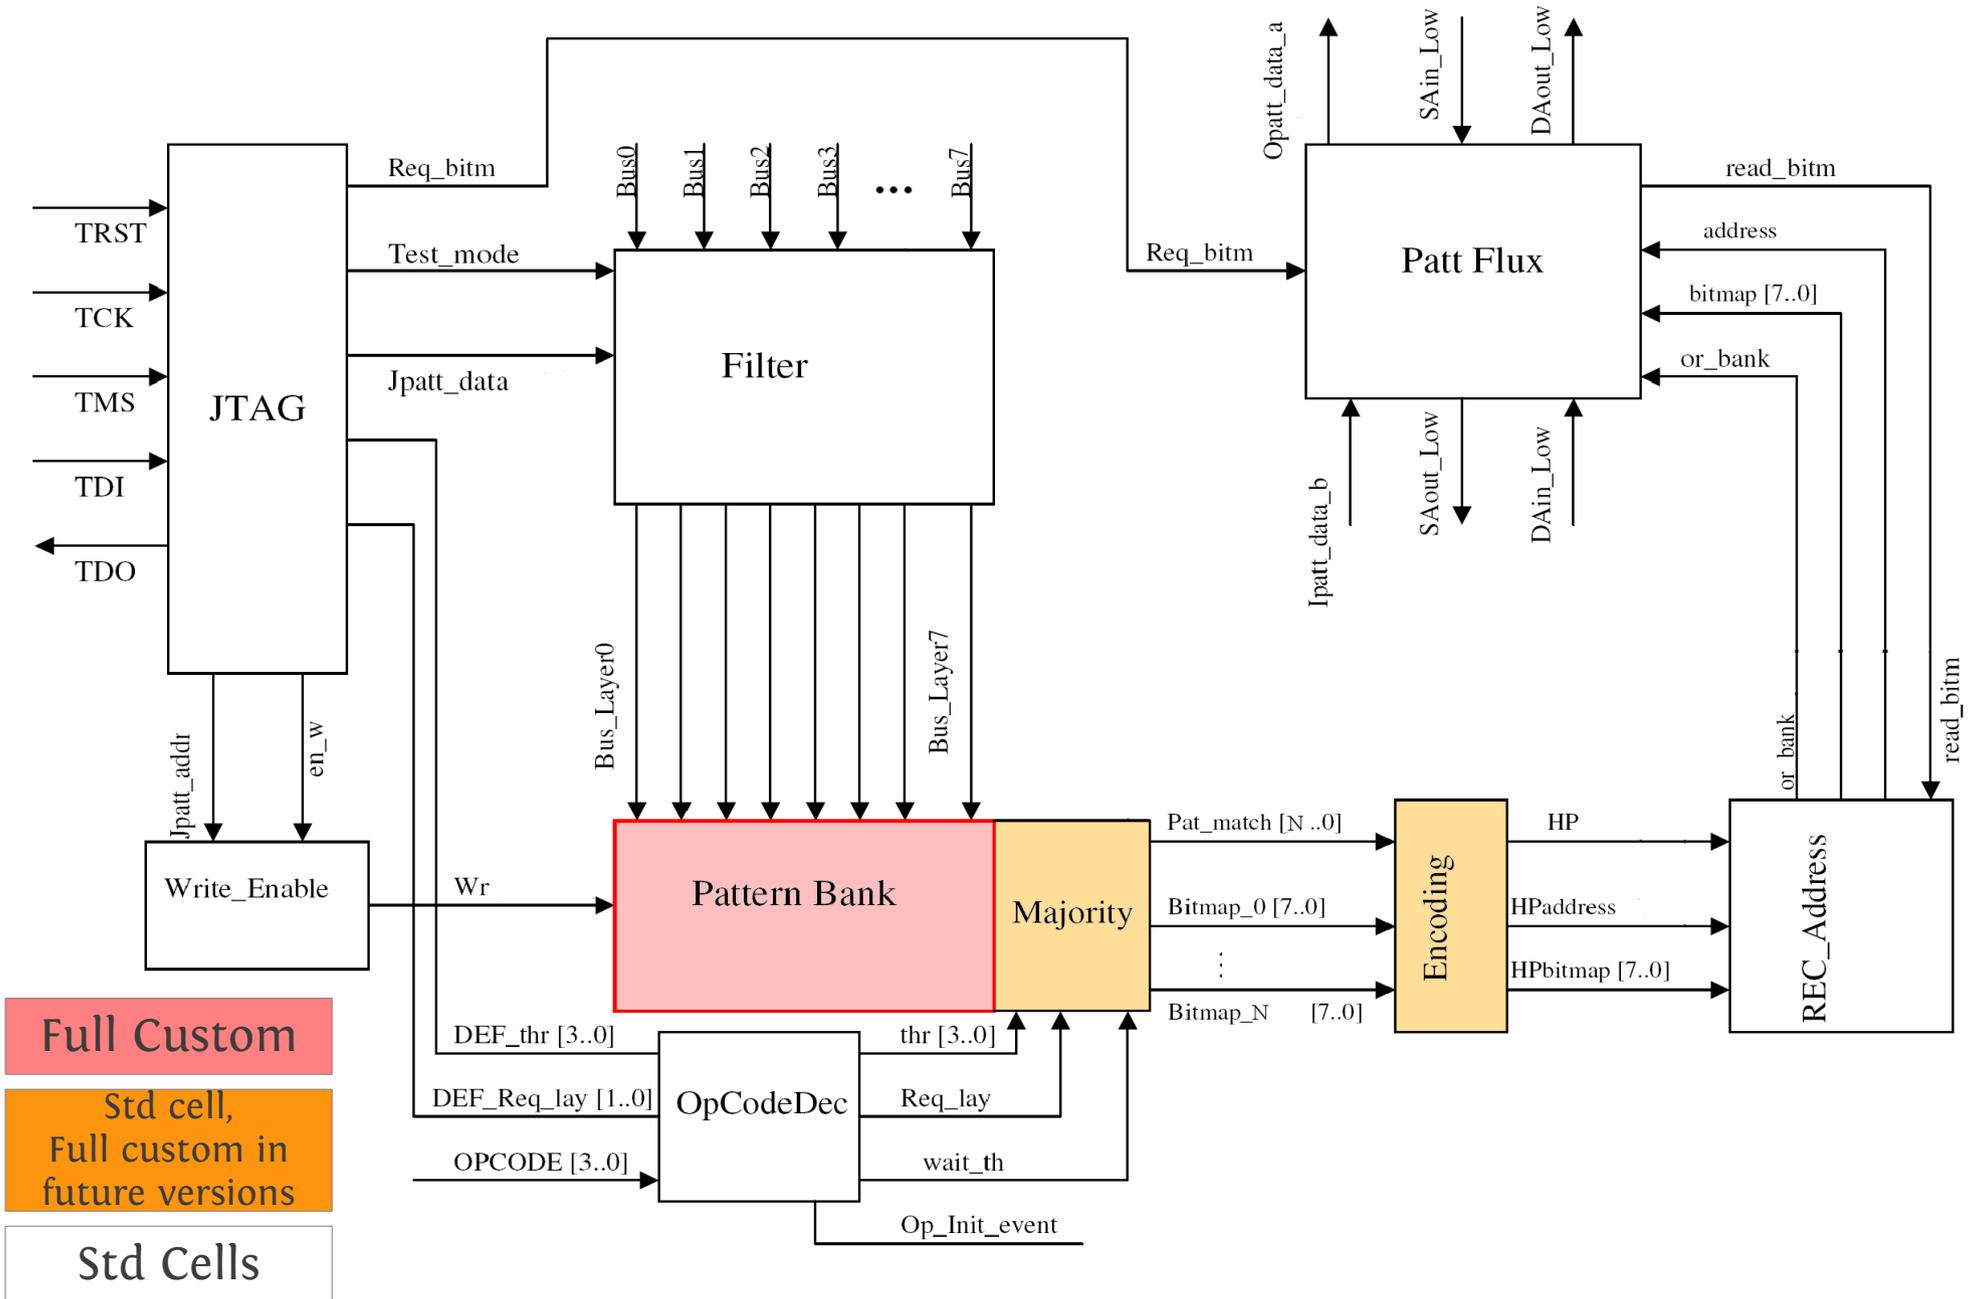
CAM layer full custom

Control logic in
Standard cells

Main dev:
INFN Pisa
INFN Milano
INFN Frascati

Minor contr.
Fermilab
Heidelberg

AMCHIP04

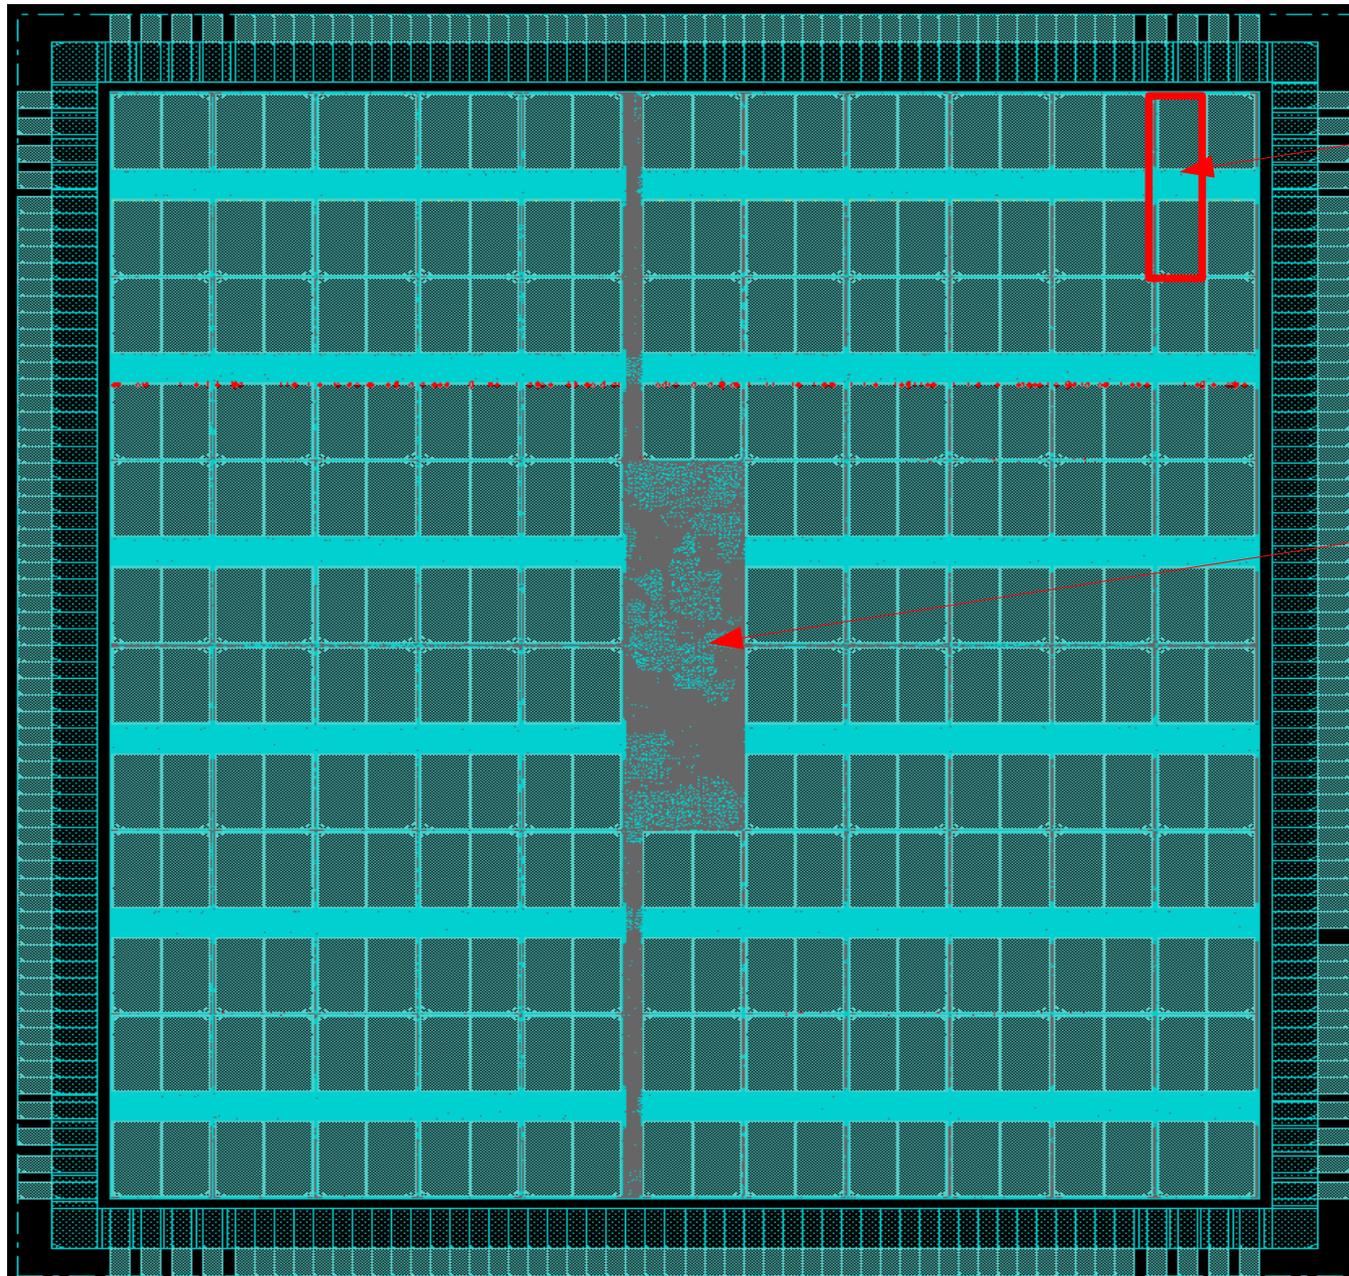


Full Custom

Std cell,
Full custom in
future versions

Std Cells

AMCHIP04: FLOORPLAN (8K)

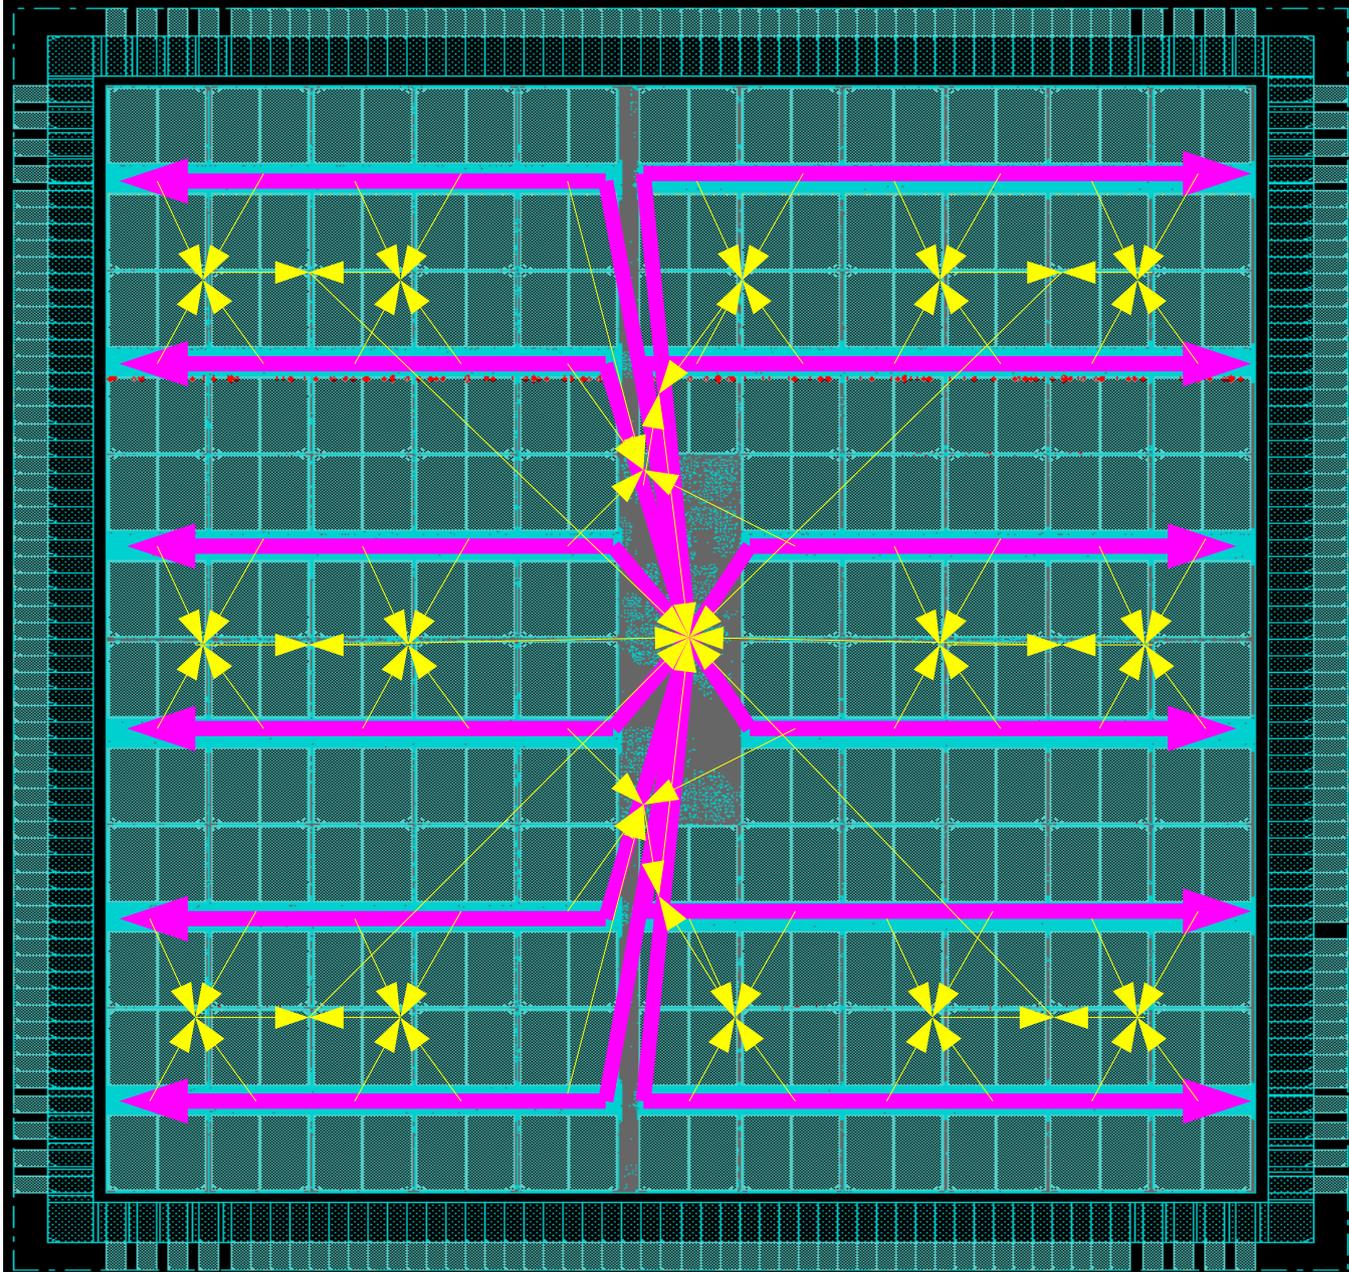


64 patterns CAM
layers block +
Local majority logic
and readout tree +
Write enable row
decoders

All remaining control
Logic (JTAG, bus filters,
pattern flux control,
opcode control)

The 8k design & floorplan
have been conceived to
minimize the work for the
final 80k version.

AMCHIP04: FLOORPLAN (8K)



Hit buses distribution

From the center control logic to 12 rows of AM memory.

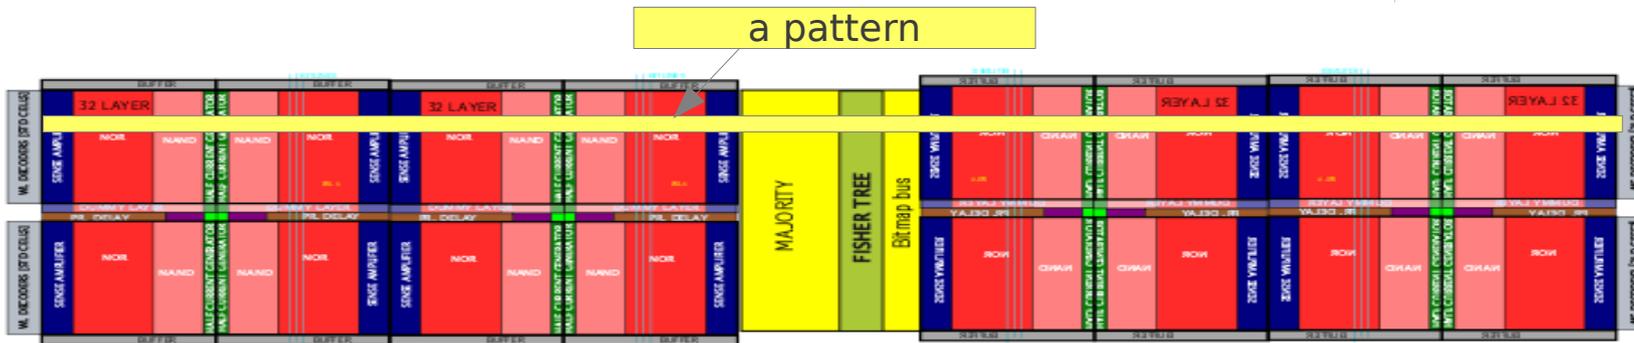
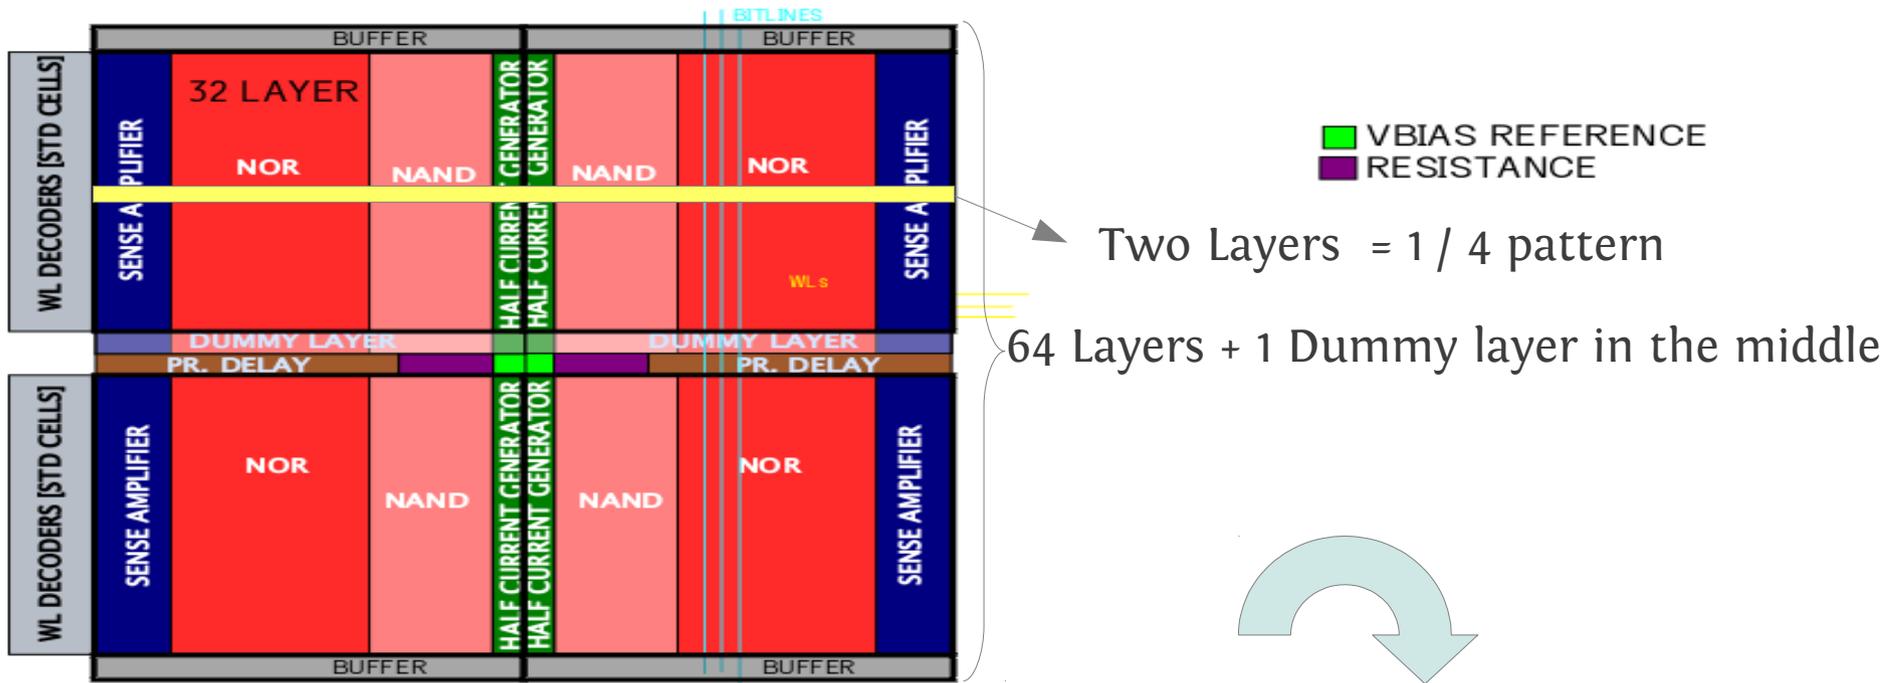
Designed to keep net delays uniform and under control without burden too much the routing (18 x 8 x 2 bits bus)

Matched patterns tree

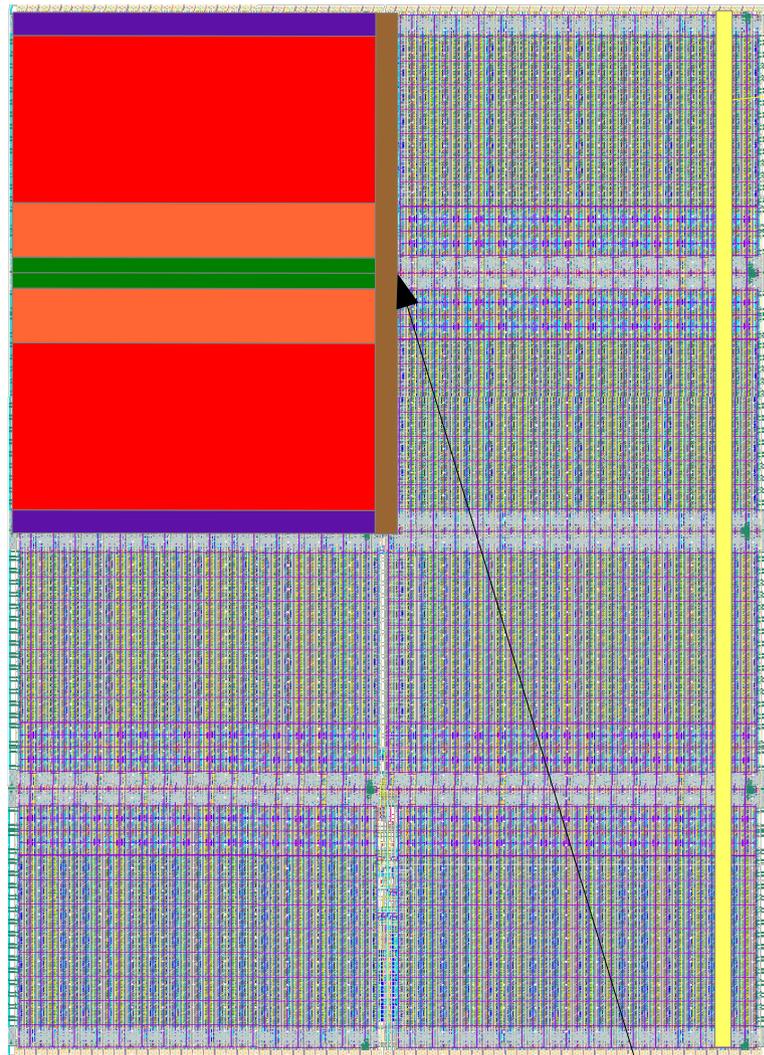
Blocks are grouped in 128 → 256 → 512 → 1024 Patterns and then routed To the center logic.

This layout is natural with our readout tree structure and not a problem for routing (at each step: $\log_2(\text{patterns})+2$ bits)

AMCHIP04: 64 PATTERNS BLOCK



AMCHIP04: 64 PATTERNS FULL-CUSTOM LAYOUT



4 Layers = 1/2 pattern

Full custom Layout of 64 x 4 CAM layers (half pattern):
w= $\sim 226 \mu\text{m}$ X
h= $\sim 123 \mu\text{m}$

without including

- major logic
- readout logic
- control logic

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

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

32 CAM layer memory block

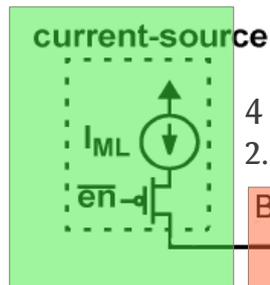
Dummy layer and programmable delay

AMCHIP04: POWER SAVING

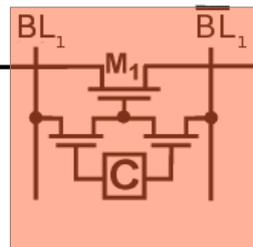
To save power we have used two different match line driving scheme:

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

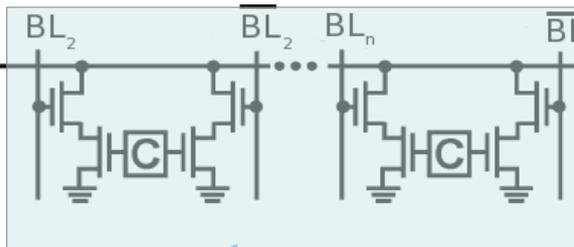
Current source:
3.7 x 1.8 μm each



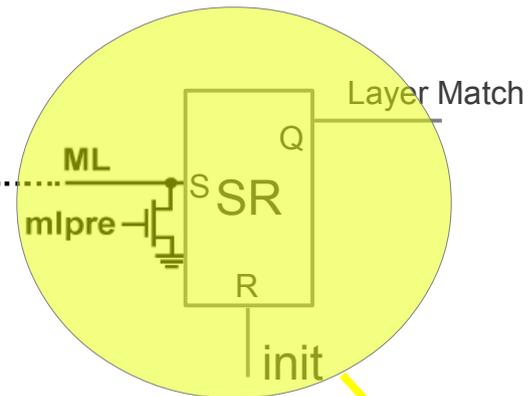
4 NAND Cells:
2.6 x 1.8 μm each



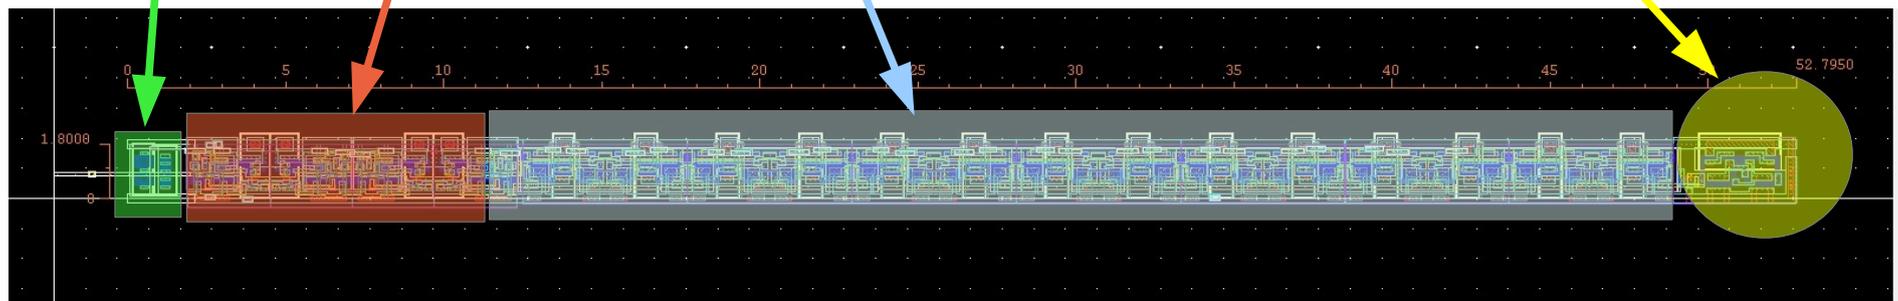
14 NOR Cells:
2.6 x 1.8 μm each



Latch + ML discharge:
4.7 x 1.8 μm each



Full Layer layout : $w \sim 53 \mu\text{m}$ X $h = 1.8 \mu\text{m}$



AMCHIP04: POWER SAVING

Current race for a layer CAMs column:

- 1) When a new hit is to be compared an enable signal (MLEN) is asserted for half a clock cycle
- 2) The column enable is the AND of MLEN and the negated output ($\overline{\text{MLOFF}}$) of an always matching layer (dummy layer)
- 3) The dummy layer “measures” the local matching time (+ a programmable delay) optimizing the time the current generator is on (power consuming)

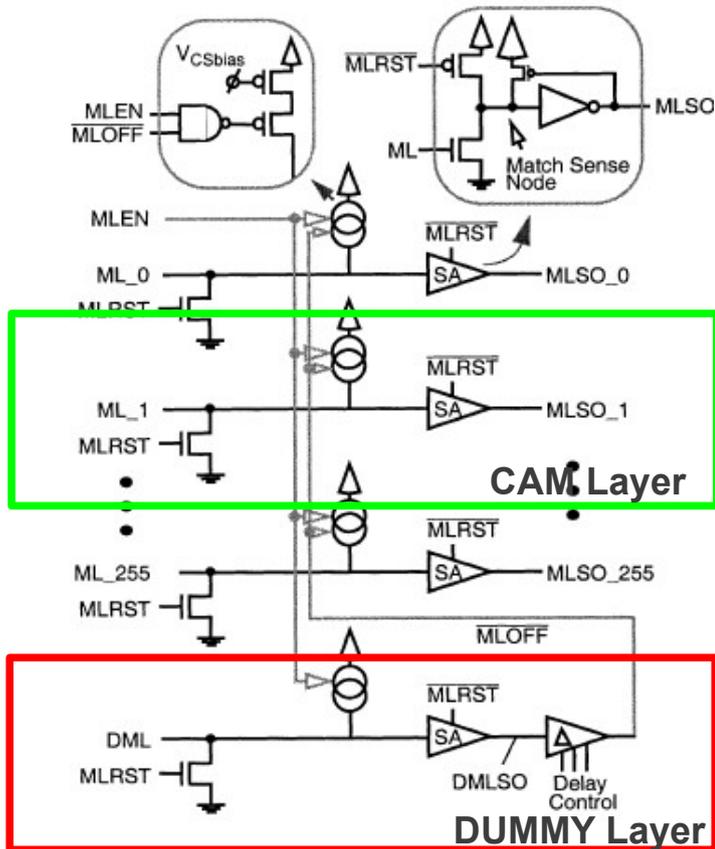


Fig. 5. Current-race ML sensing scheme.

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

AMchip04 has one dummy layer every 64 patterns in a layer CAMs column

The dummy layer is placed in the middle to optimize $\overline{\text{MLOFF}}$ distribution to the column

A four steps programmable delay is available to optimize $\overline{\text{MLOFF}}$ timing and measure current race effect on power consumption in the first prototype

AMCHIP04: POWER SAVING

Selective precharge:

- 1) The first bits of the pattern word are implemented with NAND type cells
- 2) If a NAND type cell doesn't match the match-line is not charged, saving power

NAND cells are slow, a good balance between NAND and NOR cells is needed to have both low power consumption and operational speed

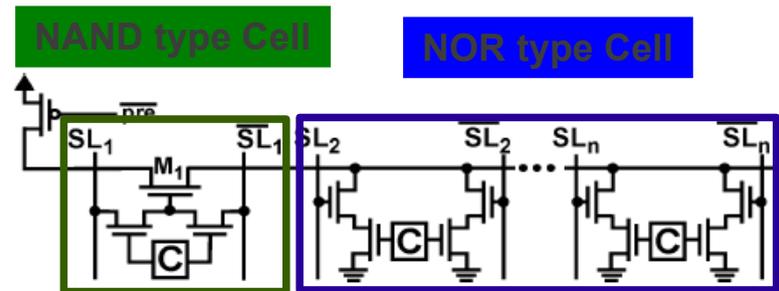


Fig. 16. Sample implementation of the selective-precharge matchline technique [43]. The first cell on the matchline is a NAND cell, while the other cells are NOR cells. Precharge occurs only in the case where there is a match in the first cell. If there is no match in the first cell, the precharge transistor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

AMchip04 layer CAMs are made by 4 NAND cells and 14 NOR cells

AMCHIP04: TIMING PERFORMANCES

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

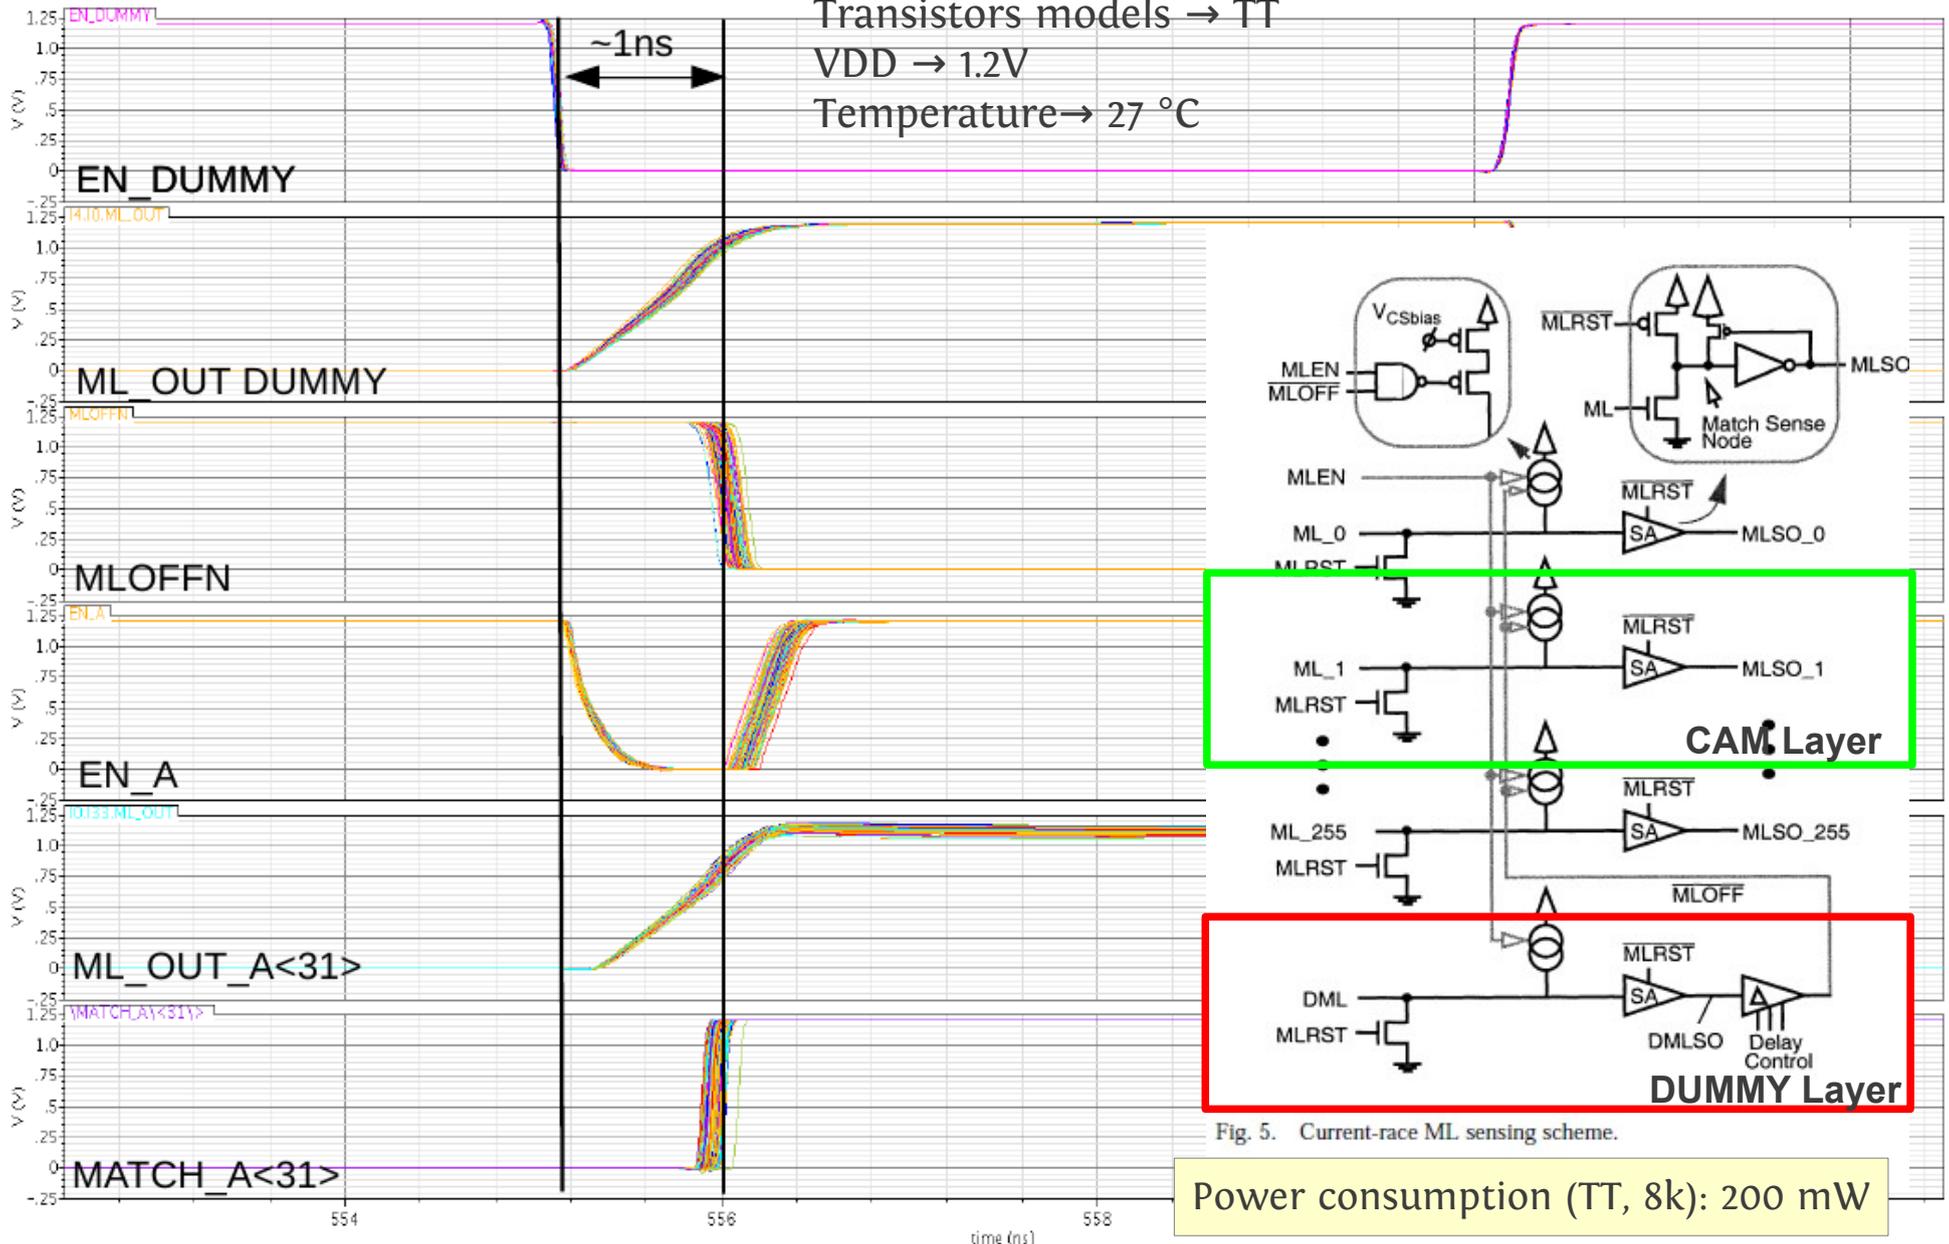
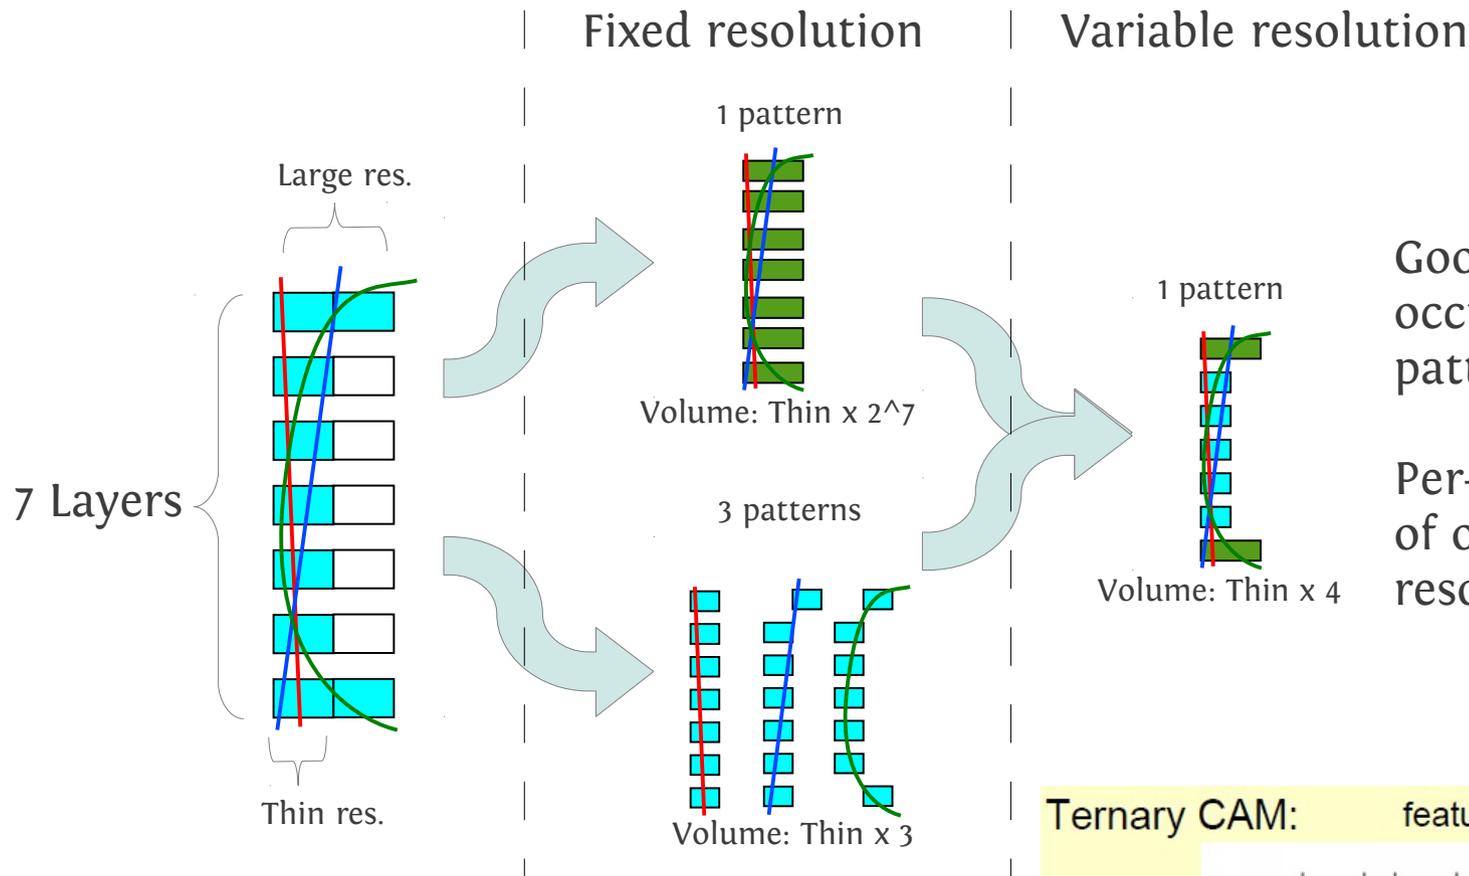


Fig. 5. Current-race ML sensing scheme.

Power consumption (TT, 8k): 200 mW

AMCHIP04: VARIABLE RESOLUTION

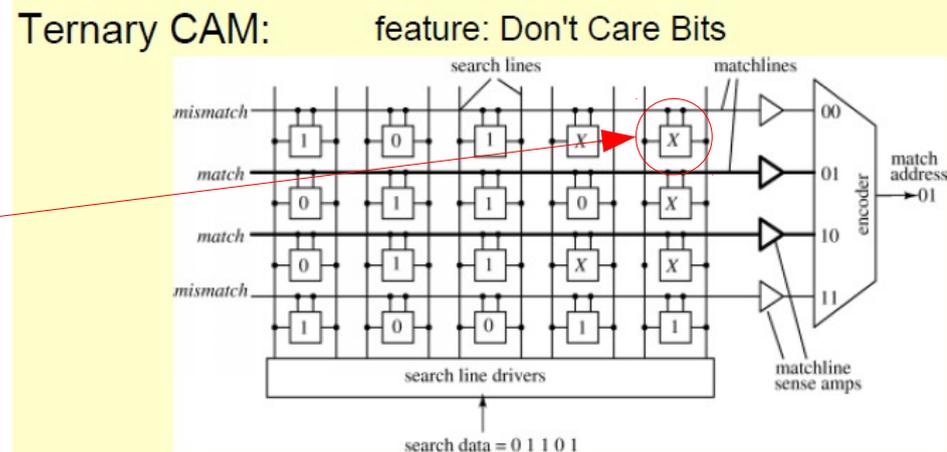


Good rejection and occupy only one pattern location.

Per-pattern choice of optimal resolution.

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



AMCHIP04: TERNARY CELLS

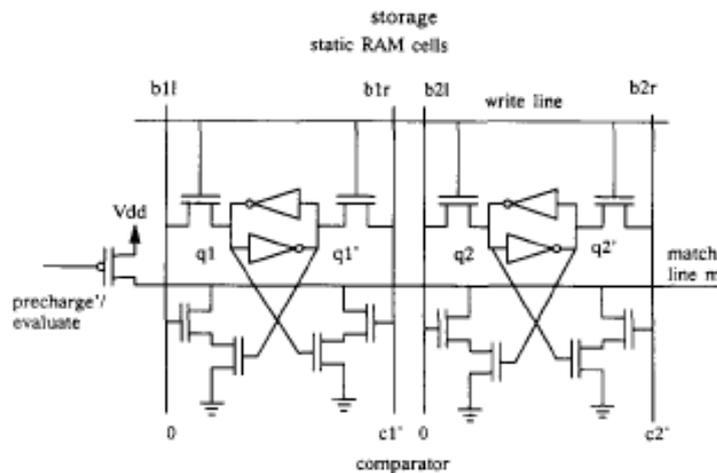


Fig. 8. Two adjacent static binary CAM cells.

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

AMchip04 has 14 NOR cells
Up to 7 ternary logic bits

No difference in CAM layer:
the number of ternary bits is
configurable via JTAG

Two standard NOR cells implement a ternary logic cell out-of-the-box.

Special drive of BL lines is needed.

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.

binary CAM
equivalent
operation

	c1c2	b1l	b1r	b2l	b2r	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.

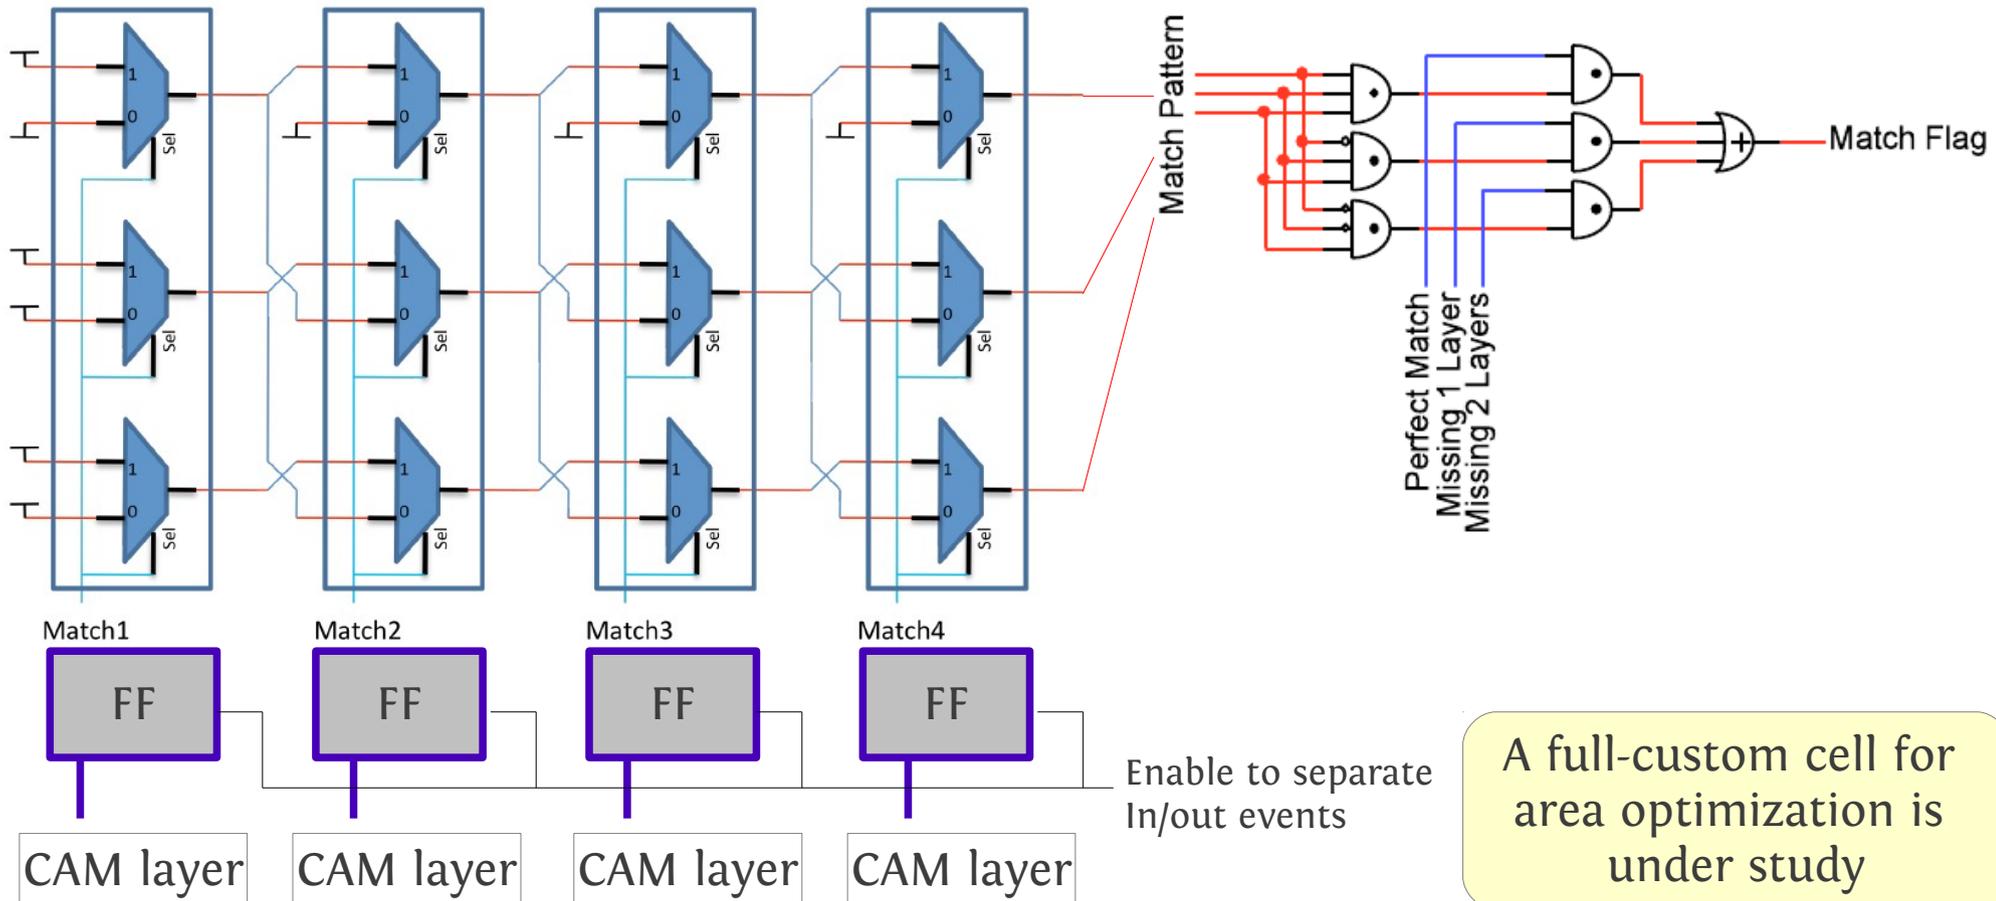
(b)

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.

AMCHIP04: MAJORITY LOGIC

Pattern match logic is made by identical logic for each layer:
Receives in input 3 bits, if not matching shift down the output.

At the end of the chain the 3 bits are compared with the majority requirements:
perfect match, 1 or 2 missing



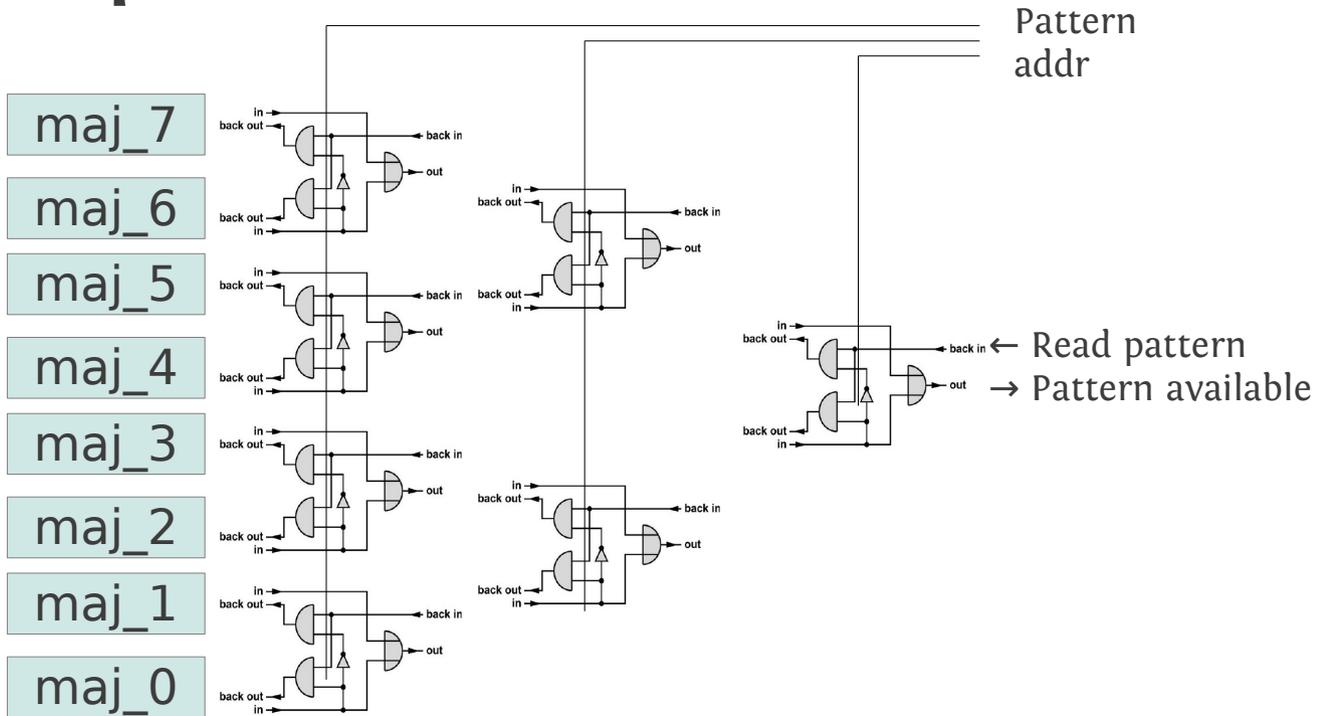
Layer matches from CAM layers are stored in a FF just before resetting the CAM layers.
We can load an event in the CAM layers while we are reading the patterns found in the previous event.

AMCHIP04: READOUT TREE

The readout tree is made by identical cells that selects priority between two branches of the tree.

Address is built by an OR of all lesser-priority selection lines at each tree level.

This design is simple and modular, it is easy to merge pattern sub-banks readout.



Simple OR-based logic
For priority encoding and readout

A full-custom cell for area
optimization is under study

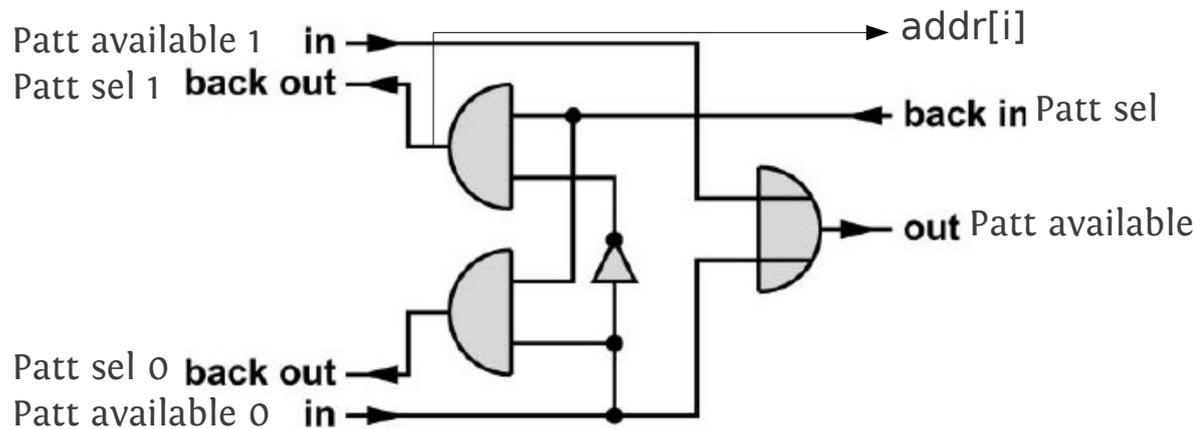


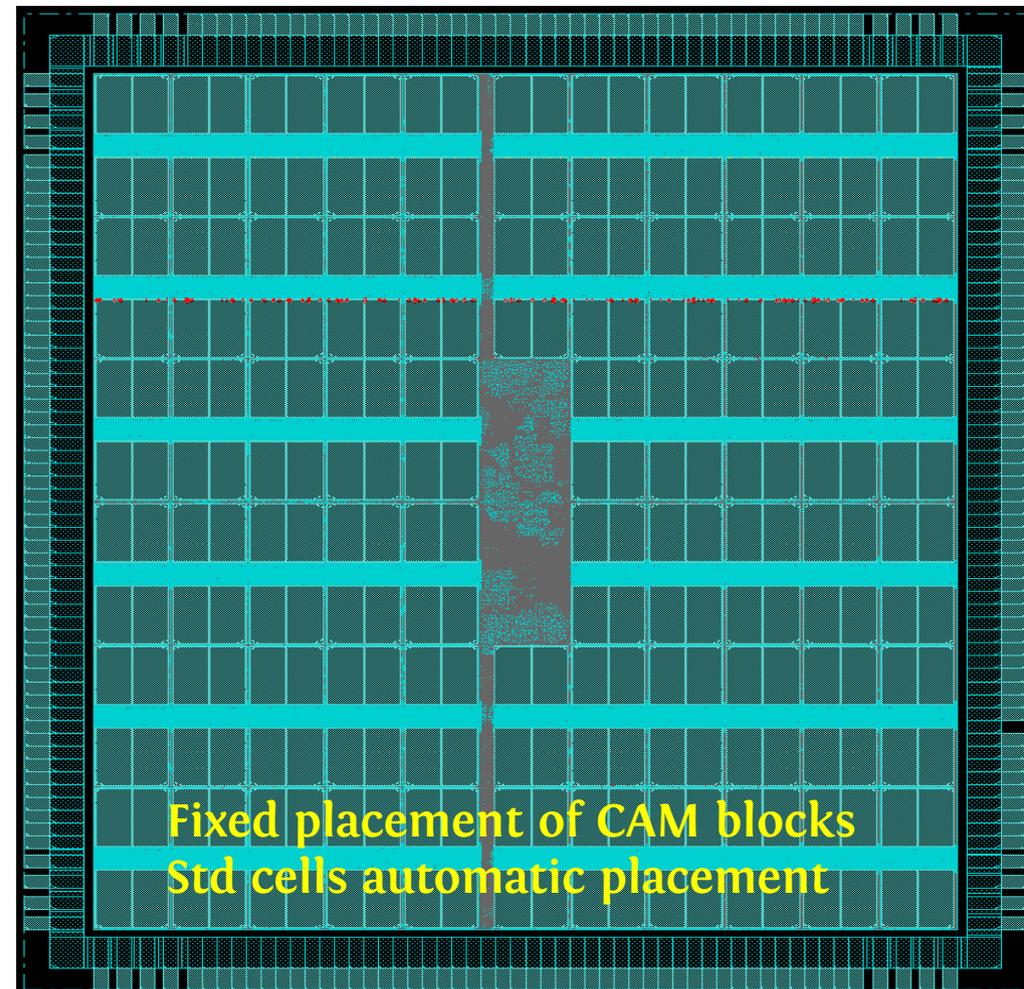
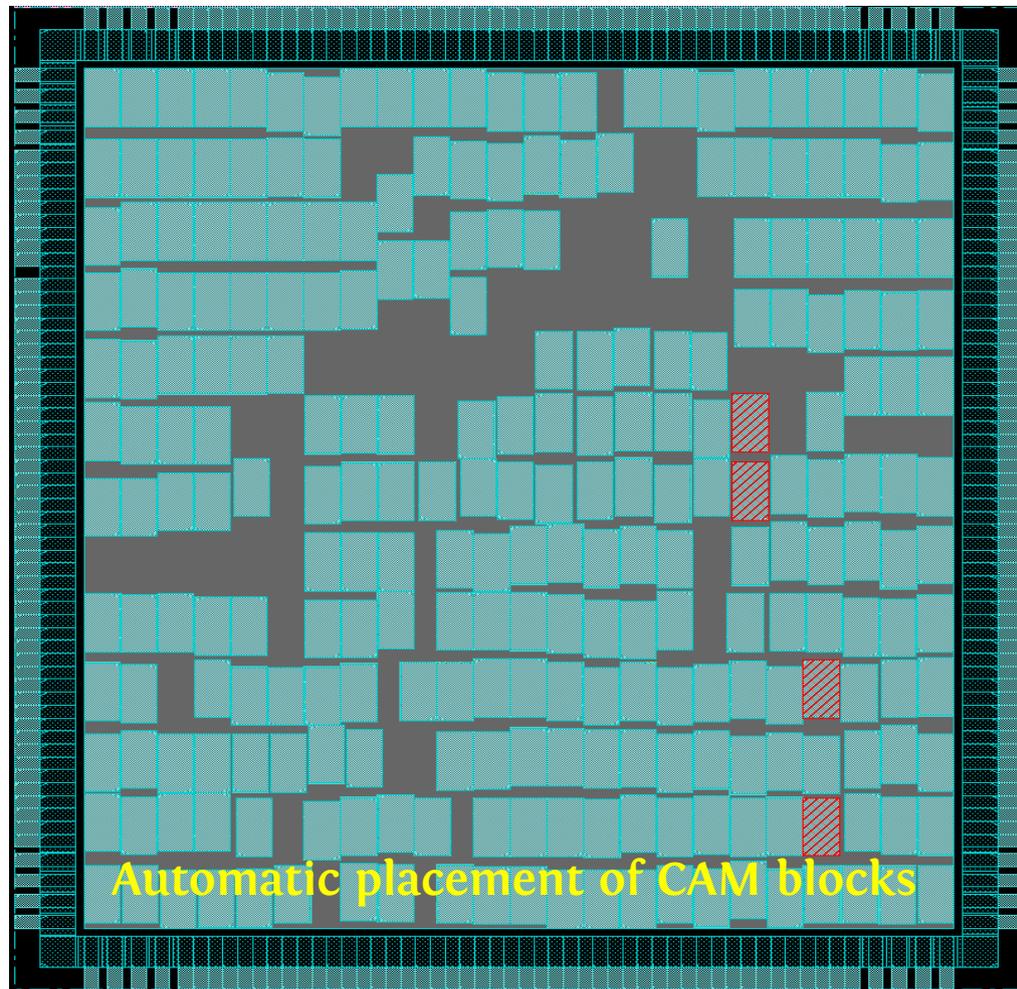
Image & cell scheme from:

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

AMCHIP04: BACK END

The AMchip04 is a challenging design from the place&route point of view.

The 64 patterns half-CAM block is placed to follow the desired general routing idea. Standard cells are placed by the automatic placer of Cadence EDI 9.1



AMCHIP04: BACK END

Fixed positions of 64 patterns CAM blocks helped are essential for complete the routing process.

The automatic router (nanoroute in Cadence EDI 9.1) is constrained to speed up and optimize the routing process. It follows the pre-planned routing.

Metal 3 (H)

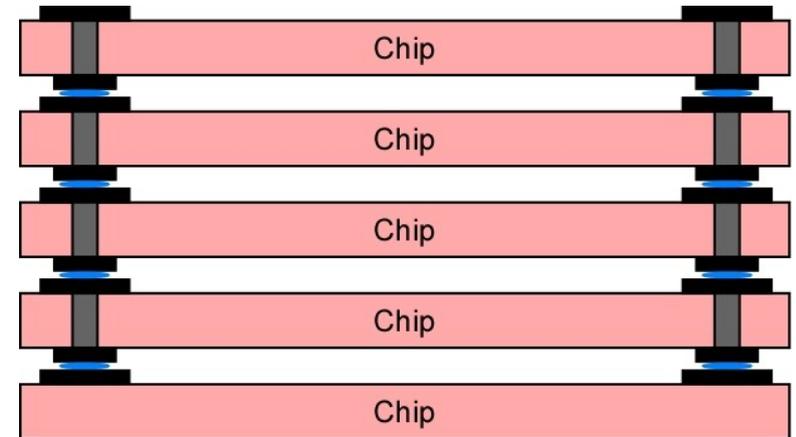
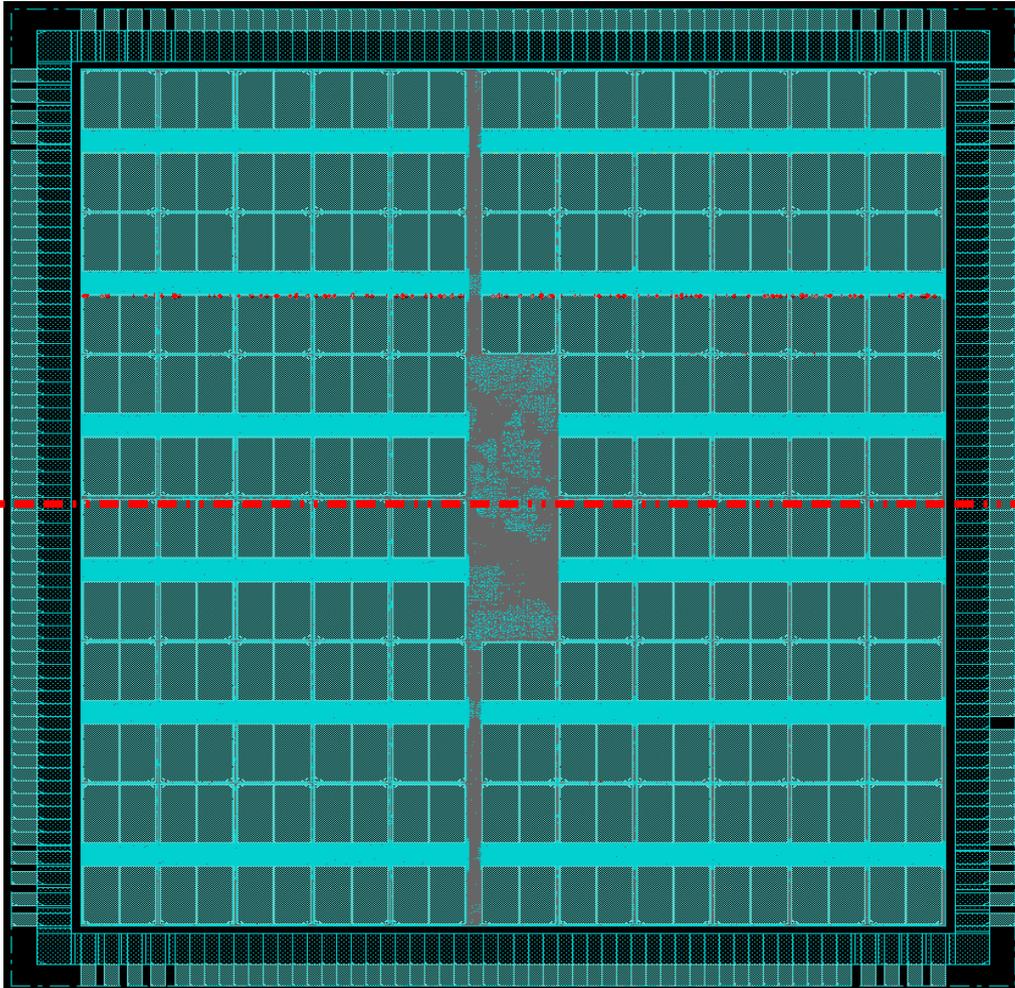
PAD_IN_BUS3_8 IN_BUS3_8
PAD_INOUT_SAIN INOUT_SAIN
FE_FILLER_W_14
FE_FILLER_W_4
PAD_IN_BUS6_11 IN_BUS6_11
FE_FILLER_W_13
FE_FILLER_W_3
PWR_VDDPST_22
FE_FILLER_W_12
FE_FILLER_W_2
PAD_PWR_VSS_23 PWR_VSS_23
FE_FILLER_W_11
FE_FILLER_W_1
PAD_PWR_VSS_24 PWR_VSS_24
FE_FILLER_W_10
FE_FILLER_W_0

/TOP2_TR
/TOP2_BR
/TOP2_TR
/TOP2_BR
/TOP2_TR
/TOP2_TL
/TOP2_BL
/TOP2_TL
/TOP2_BL
/TOP2_TL

AMCHIP04 NEAR FUTURE PLANS

- New features, pattern density and simulated power consumption are within the stated goals of the chip
- FrontEnd & BackEnd almost finished
- First prototype of the new FTK AMchip04 will be submitted to the foundry in autumn
- A second prototype is foreseen to test new technology (full-custom majority and readout tree) and to correct possible bugs/problems

FUTURE EVOLUTIONS IN 2.5D



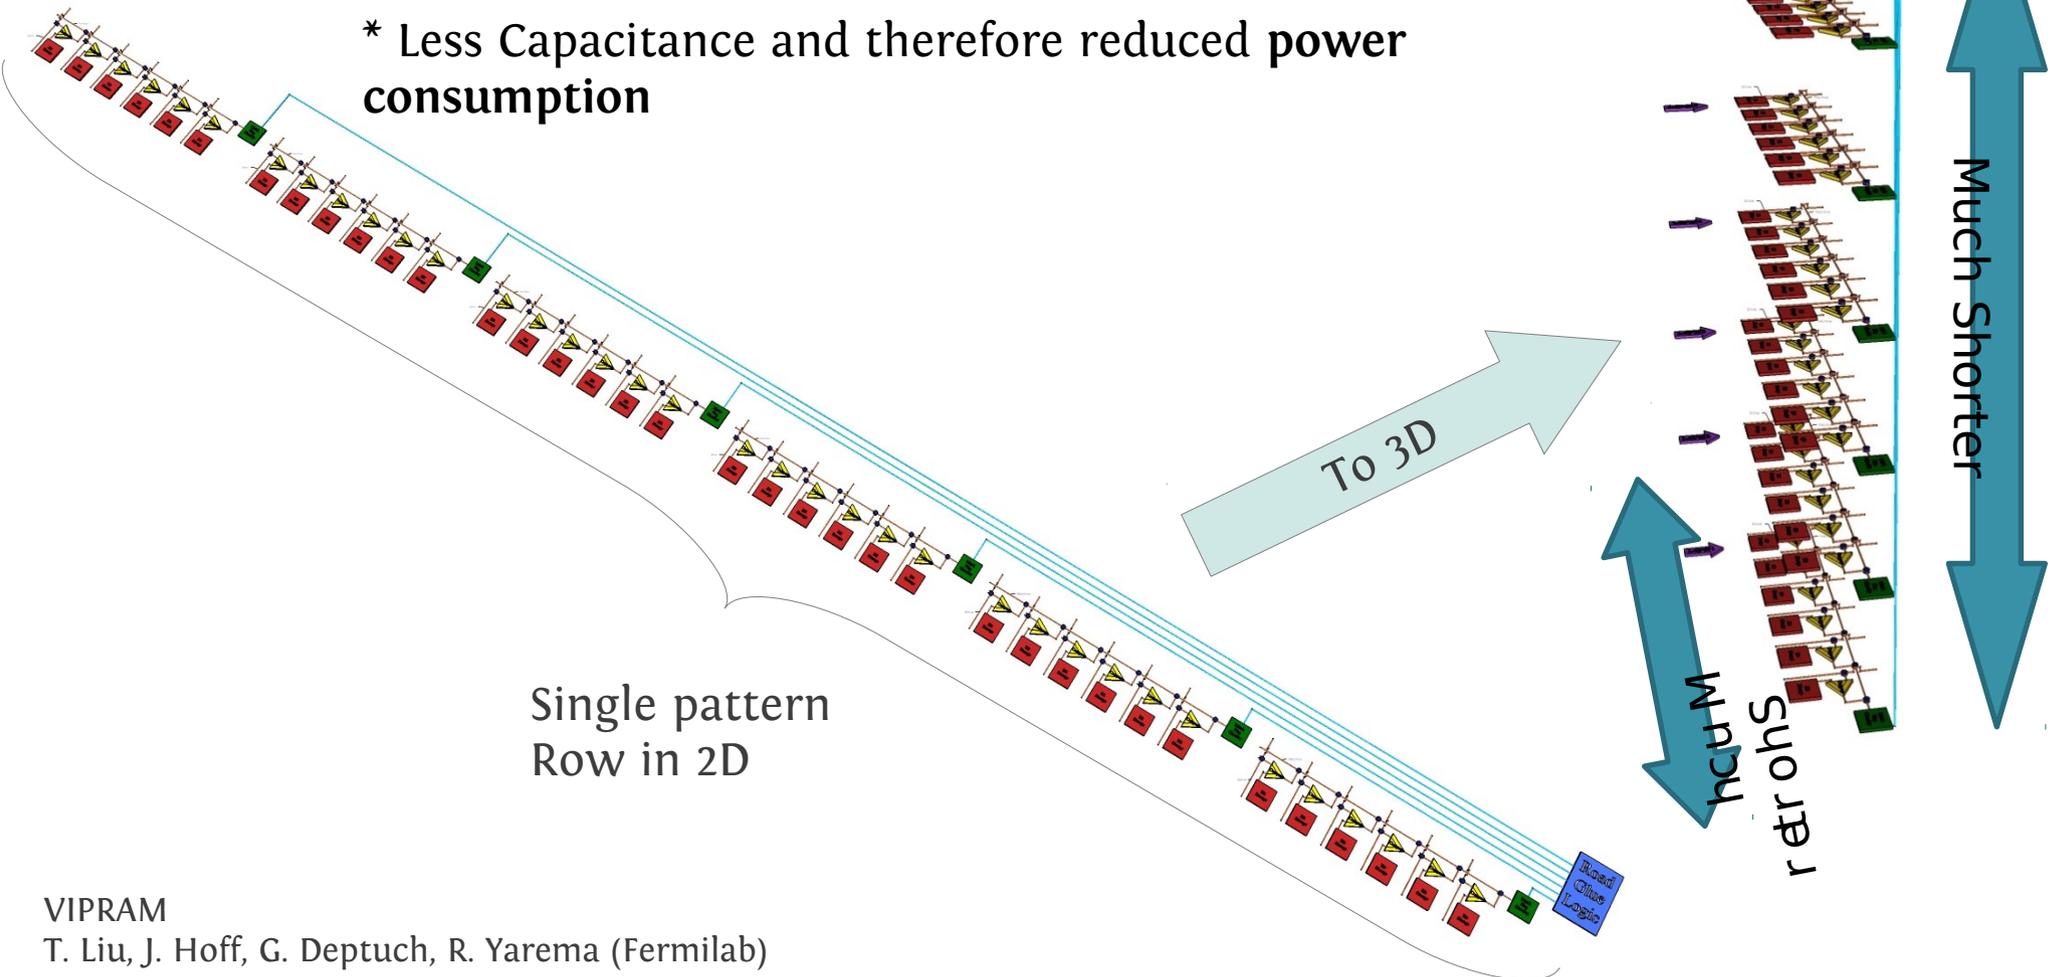
AMchip04 has been designed to be horizontally symmetric.

- In/out buses for pattern output pipeline can change direction
- Buses are swapped internally to maintain consistency

Symmetry helps in designing and routing mezzanines for 2D chips, but also enables vertical stacking:

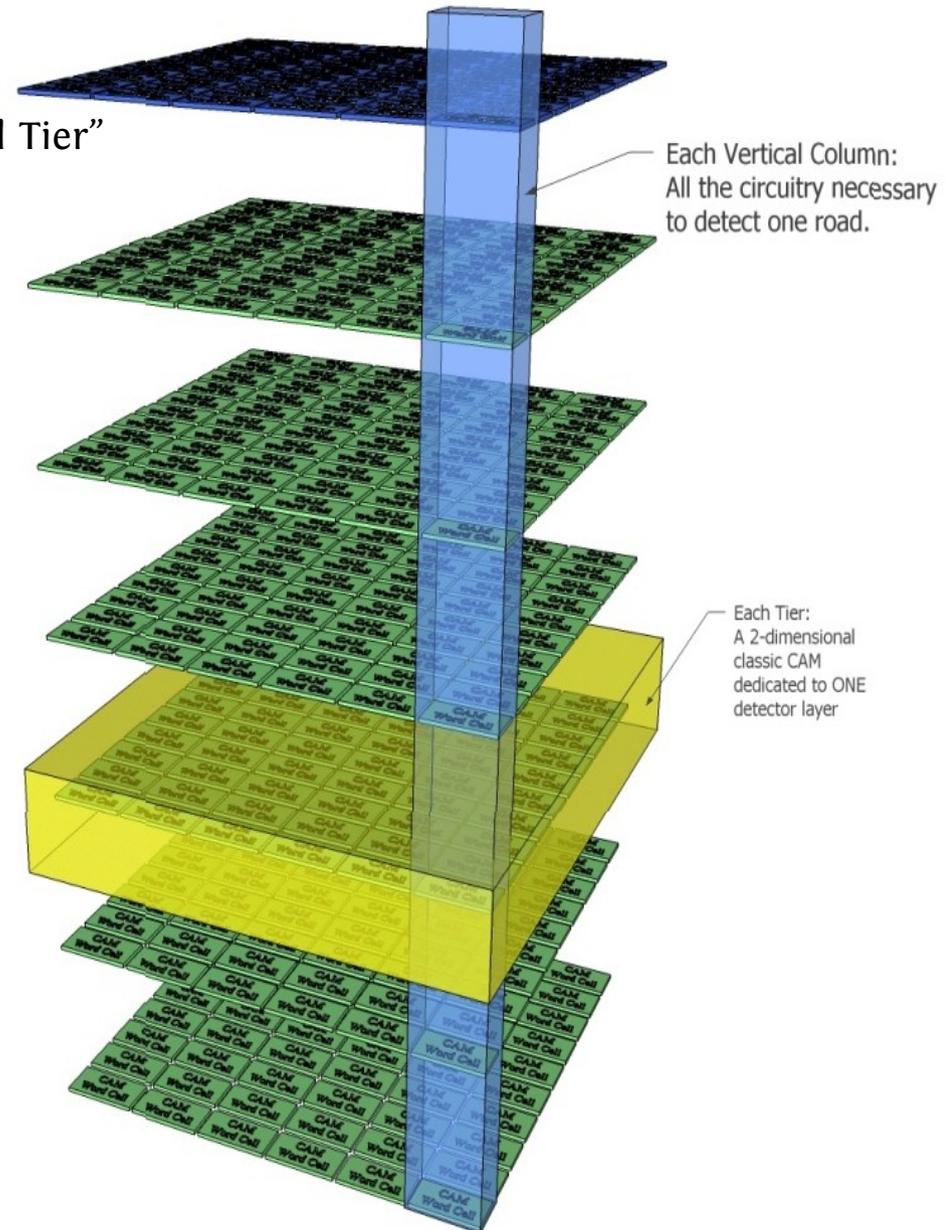
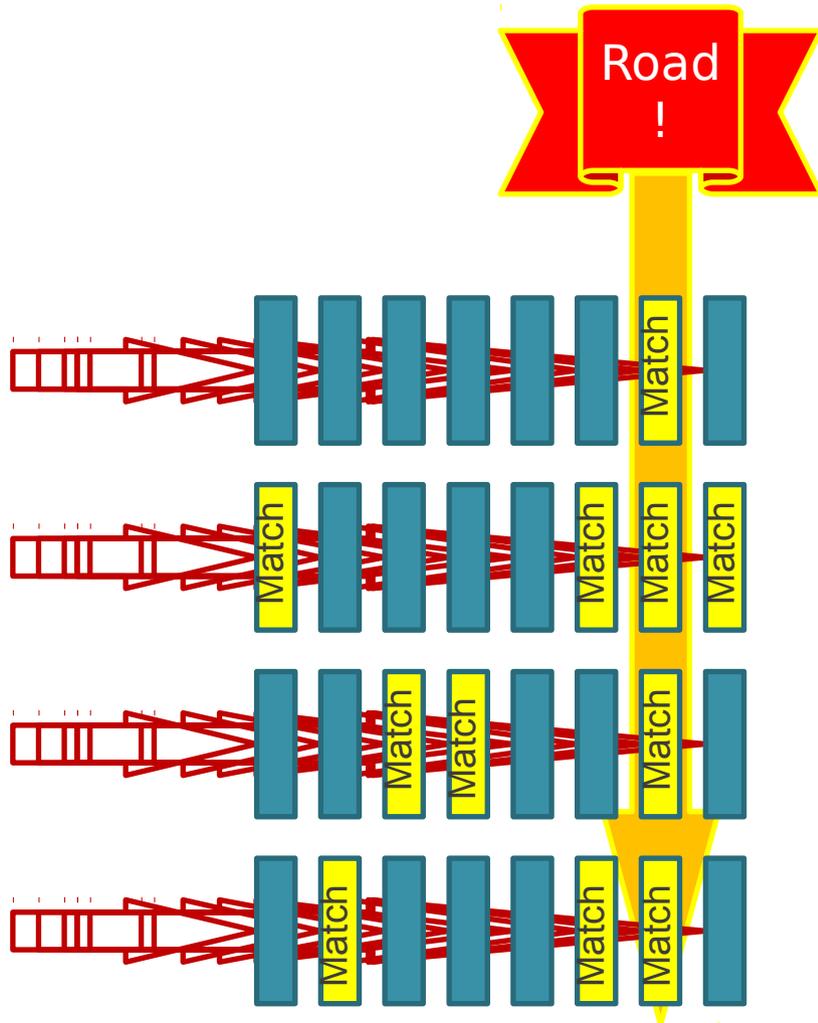
FUTURE EVOLUTIONS IN 3D

- * A Reduced Footprint and therefore greater **pattern density**.
- * Shorter Match lines and therefore greater **speed**.
- * Less Capacitance and therefore reduced **power consumption**



FUTURE EVOLUTIONS IN 3D

- * Each detector layer corresponds to a single tier
- * All communication from “CAM Tiers” to the single “Control Tier”



VIPRAM
T. Liu, J. Hoff, G. Deptuch, R. Yarema (Fermilab)

29-09-2011

Francesco Crescioli (INFN Pisa/CERN) ~ LPNHE seminar

38

CONCLUSIONS

- AMchip is a device for solving pattern recognition problems in HEP applications
 - Conceived in the late 80s, developed since 90s
 - AMchip (SVT), AMFPGA, AMchip03 (SVT Upgrade), **AMchip04 (FTK)**
 - Similar to a CAM, but with extended functionalities
 - Separate buses, majority logic, variable resolution
 - It should be considered as a different kind of memory
- AMchip has been developed for use in specific online hardware tracking processors ...
 - SVT & Upgraded SVT at the CDF experiment
 - FTK at the ATLAS experiment
- ... but it's a general purpose device
 - (other HEP tracking applications: L1 triggers, future 3D devices; outside HEP applications: computational vision)

CONCLUSIONS

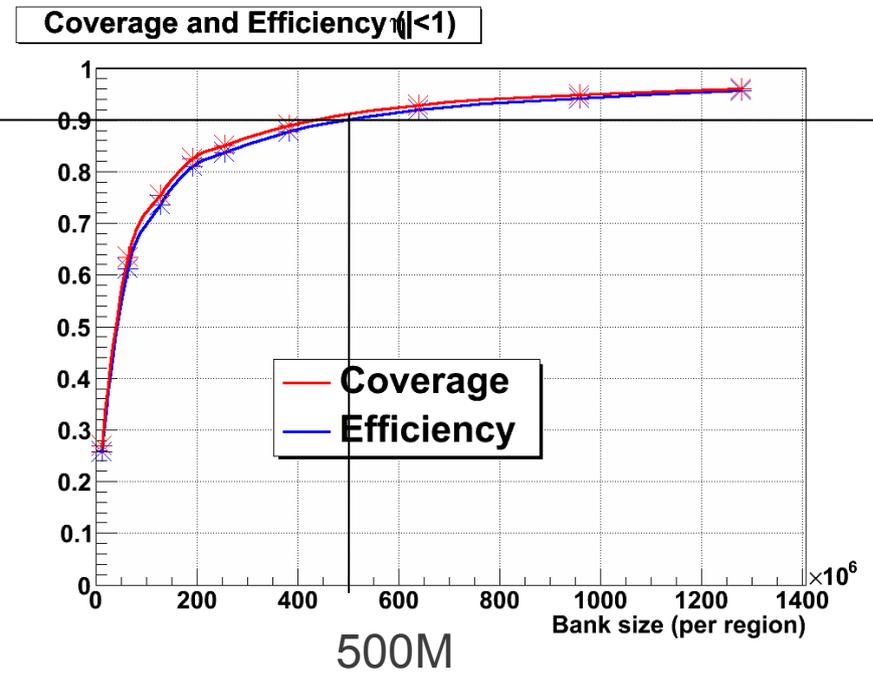
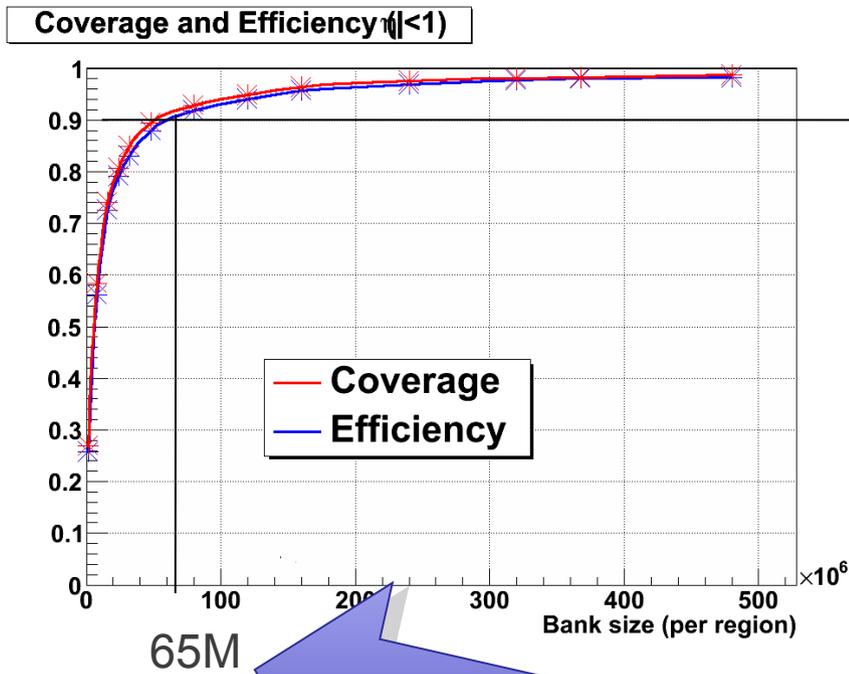
- AMchip04 represents a major improvement in the AMchip family
 - Pattern by pattern variable resolution (“don’t care” bits)
 - Matched patterns readout / next event loading at the same time
 - Mixed full-custom / standard cells design
 - CAM blocks as full-custom hard blocks, optimized for low power consumption and area efficiency
 - Control logic in standard cells (easy to develop & debug)
 - New design of majority and readout tree (simple, modular, full-custom version is foreseen)
 - Performance improvements wrt AMchip03: $>10\times$ pattern density, $<1/10$ W/pattern, 100 MHz, 8 buses
- Designed with future evolutions in mind (vertical stacking, full 3D design, ...)
- First prototype (8k patterns) will be available by Q1 2012

Backups

Pattern efficiency

Pattern size
 $r-\phi$: 24 pixel, 20 SCT
 36 pix z

Pattern size
 $r-\phi$: 12 pixel, 10 SCT
 36 pix z



90%

of channels (barrel only, 45 ϕ degrees)

Want this

<# roads/event @ 3E34> = 342k

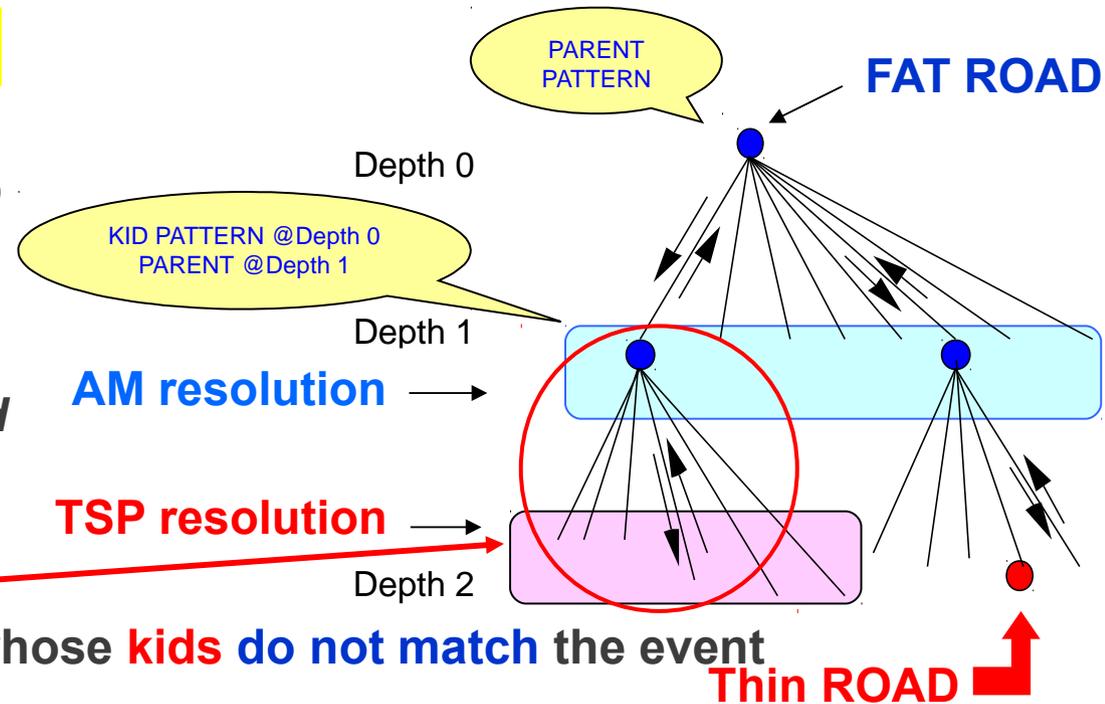
<# roads/event @ 3E34> = 40k

TSP simulation & varying-resolution pattern banks

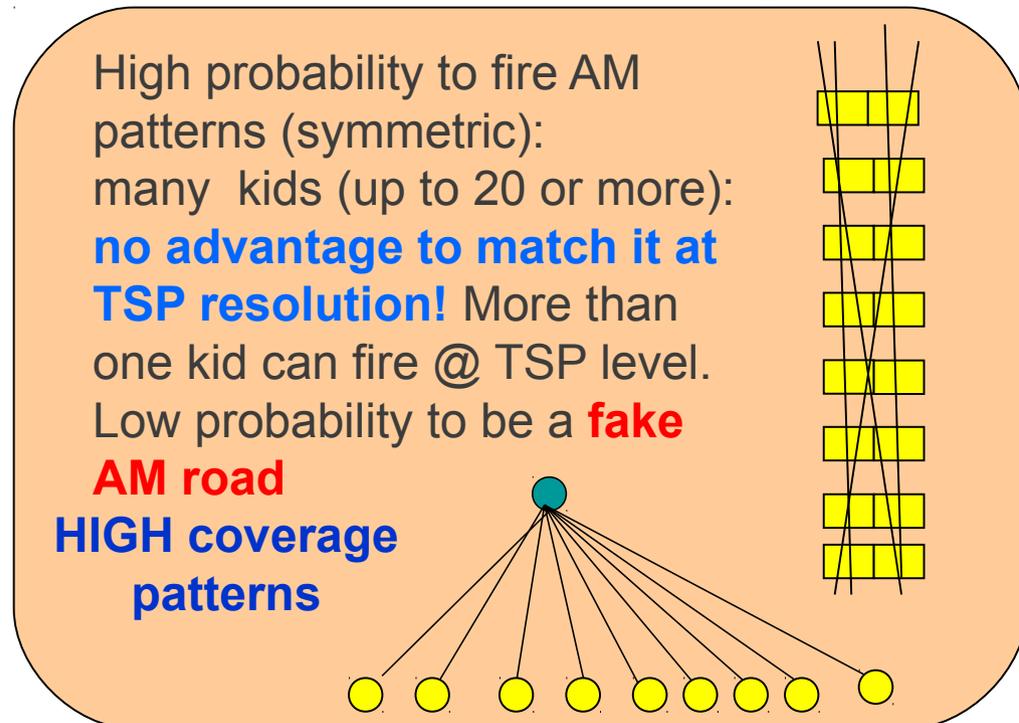
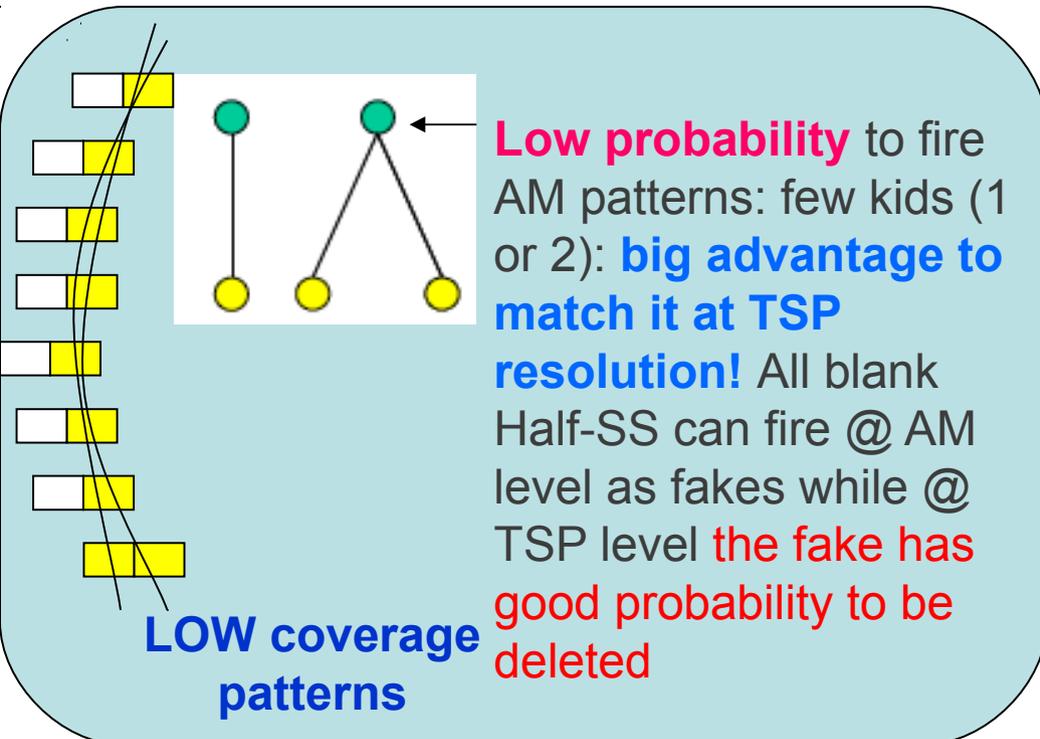
Guido Volpi & Roberto Vitillo - Pisa

We do have now a **structured** “**pattern bank**”, where each **thin road** is connected to its **parent pattern** in FTKsim.

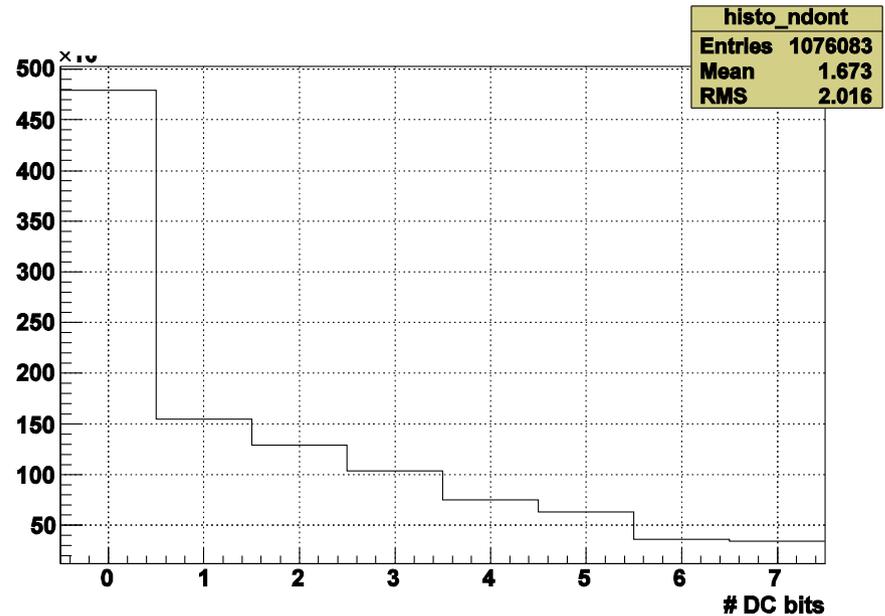
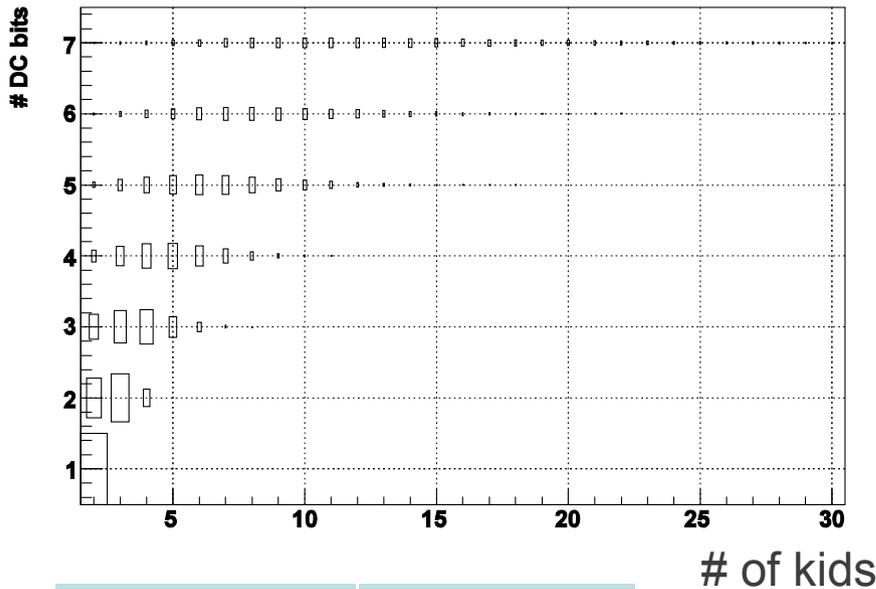
Ongoing tests for **TSP algo** after the **RoadFinder** (AMsim) in FTKsim; we have studied the **bank composition** and **AM FAKE roads**.



AM Fake road is a **AM matched pattern** whose **kids do not match the event**



AM with care / don't care



	<roads/event >
TSP	38000
AM@TSP	28000
AM@DC	44000
AM	342000

Care/don't care very effective to reduce the number of roads.

Area cost on the chip approx. 1 extra cell for each DC bit.

Now 15 cells/layers.

With 1 DC bit area increases by 1/15 ~ 7%.

For comparison going to TSP resolution would require 3x patterns.

MAIN DEVELOPERS

- Full custom design and layout: M. Beretta, A. Stabile
 - Tools: Virtuoso
- Back end implementation: A. Stabile, F. Crescioli
 - Tools: EDI 9.1
- Front end implementation: F. Crescioli
 - Tools: Design Compiler
- C++/Verilog simulation: F. Crescioli, I. Sacco
 - Tools: custom C++ emulator, NC Sim
- Coordination & ideation: A. Annovi, P. Giannetti, V. Liberali

FAST PATTERN RECOGNITION WITH DEDICATED HW

Coarse Template Finding

Refinement Step

Contant Addressable Memory

(unencoded coarse hit position)

template				
addr 1	1 X X	1 X X	X 1 X	X 1 X
addr N-1	X X 1	X X 1	X X 1	X X 1
addr N	X X 1	X X 1	X 1 X	X 1 X
addr N+1	X X 1	X 1 X	X 1 X	X 1 X

Tag RAM

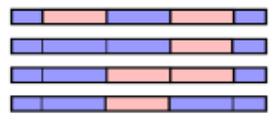
encoded coarse hit position	base address Validator RAM
00 00 01 01	base 1
10 10 10 10	base N-1
10 10 01 01	base N
10 01 01 01	base N+1

template
addr 1
addr N-1
addr N
addr N+1

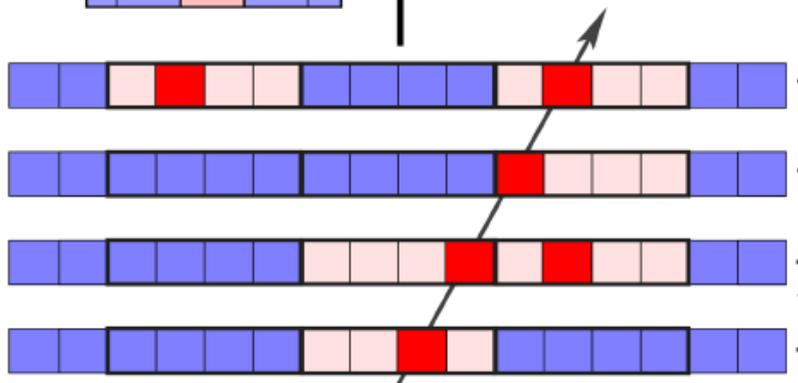
matching

1	0	1	0	0	1	0	1	1	0	1	0
---	---	---	---	---	---	---	---	---	---	---	---

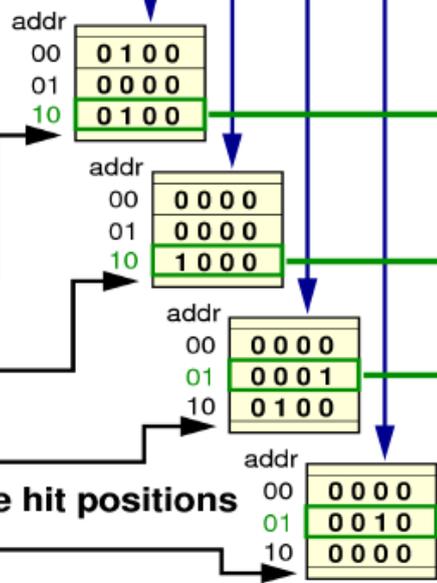
coarse hit positions



Coarsening



Buffers

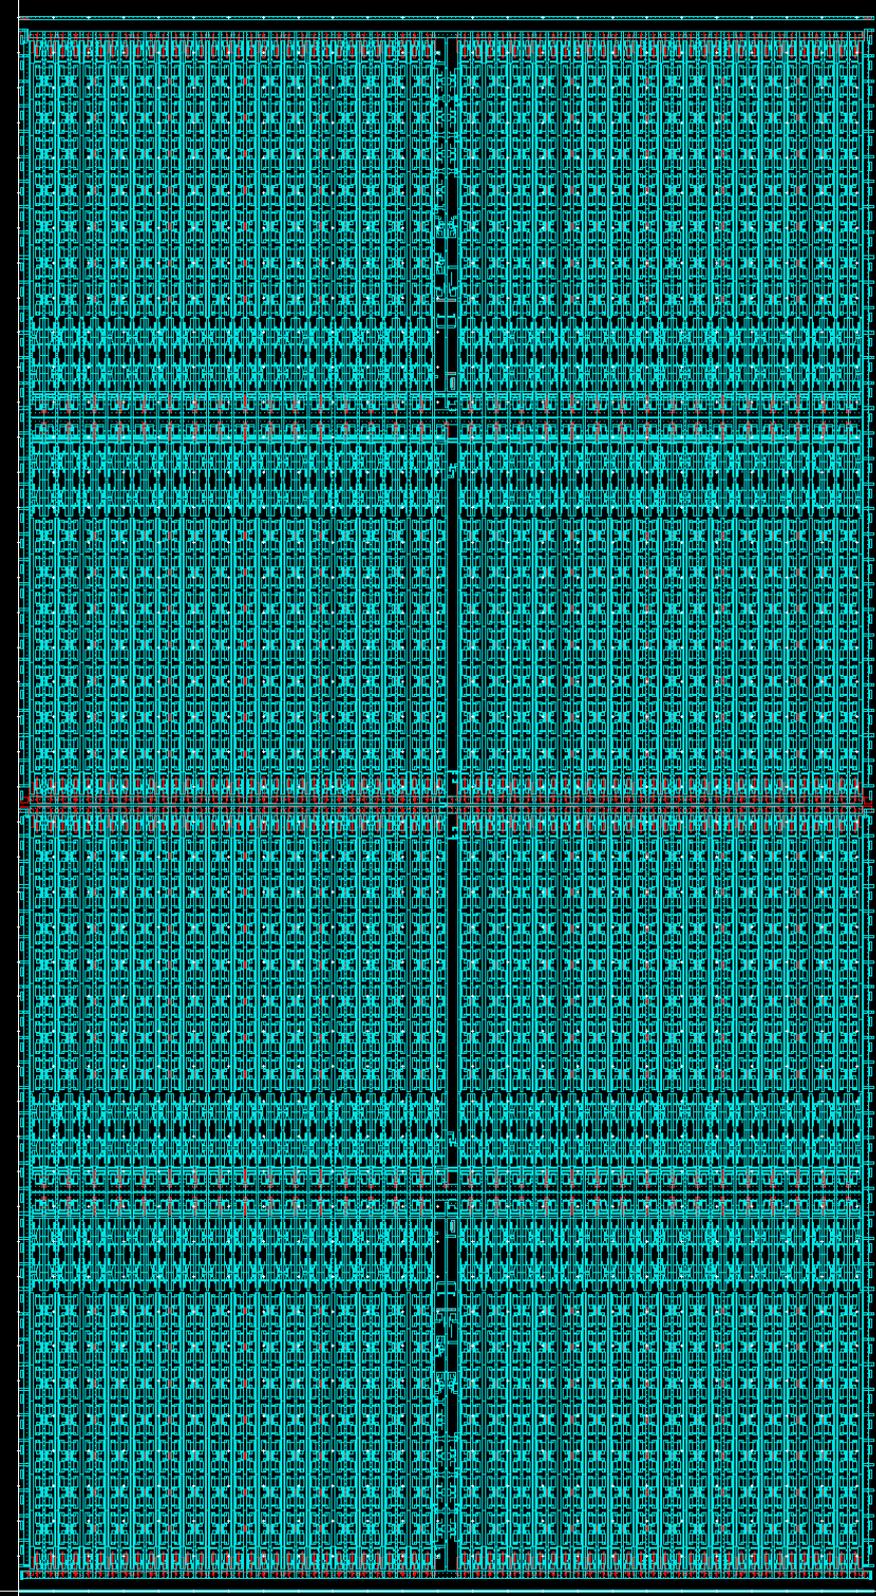


Validator RAM

lookup address:			valid bit
base	refined hit pattern		
N-1	1111 1111 1111 1111		0
N	0000 0000 0000 0000		0
N	0100 1000 0001 0001		0
N	0100 1000 0001 0010		1
N	0100 1000 0001 0001		1
N	1111 1111 1111 1111		0
N+1	0000 0000 0000 0000		0

lookup

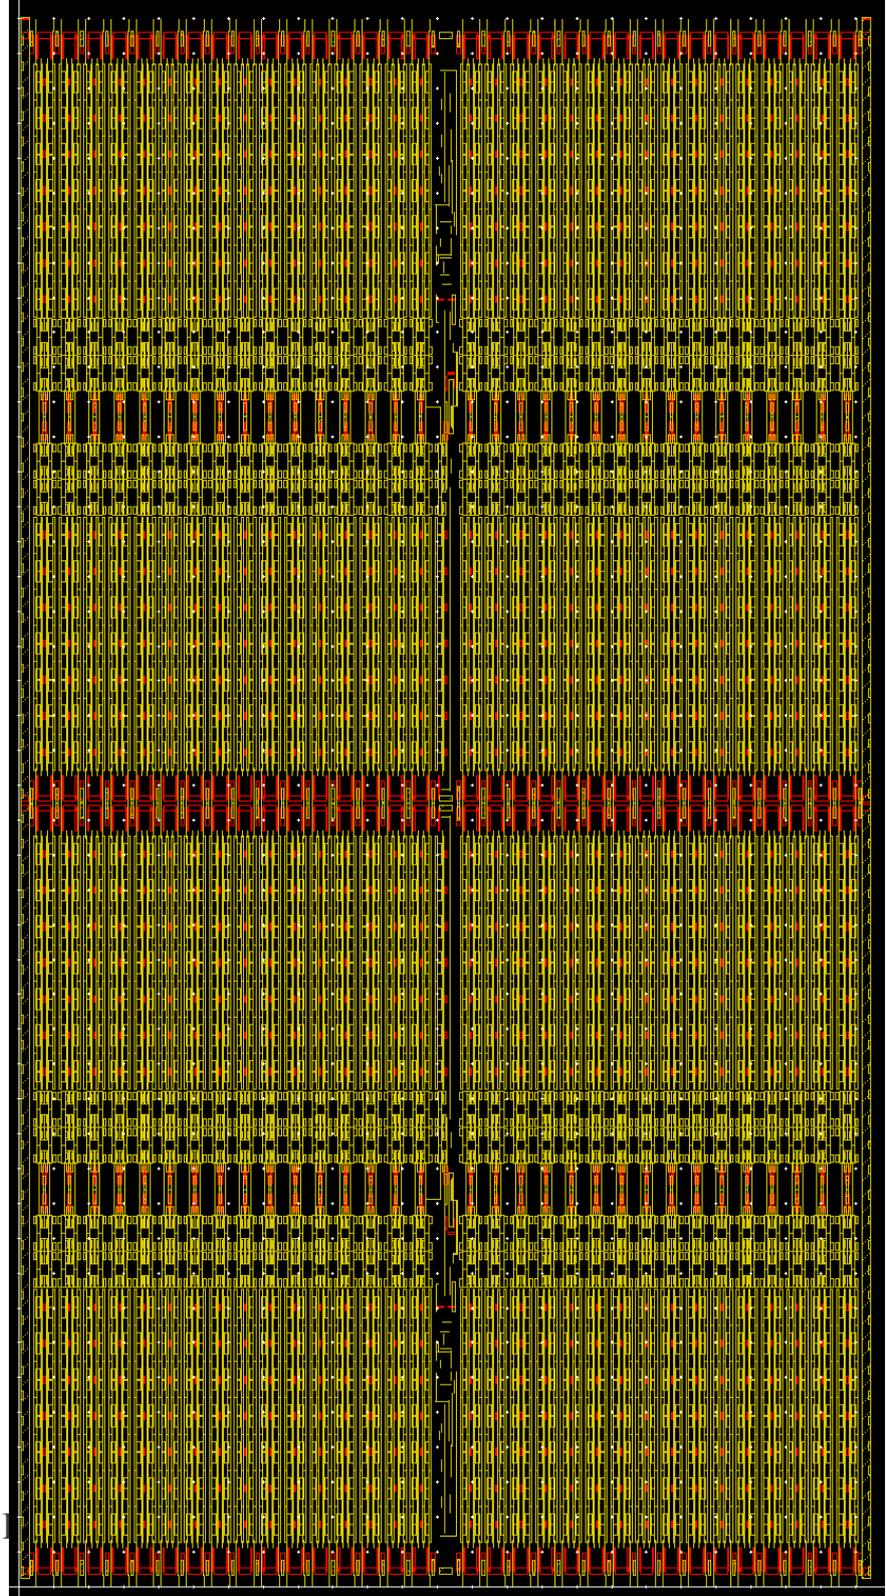
(from A. Schoening, S. Schmitt ideas for ATLAS L1 tracking trigger)

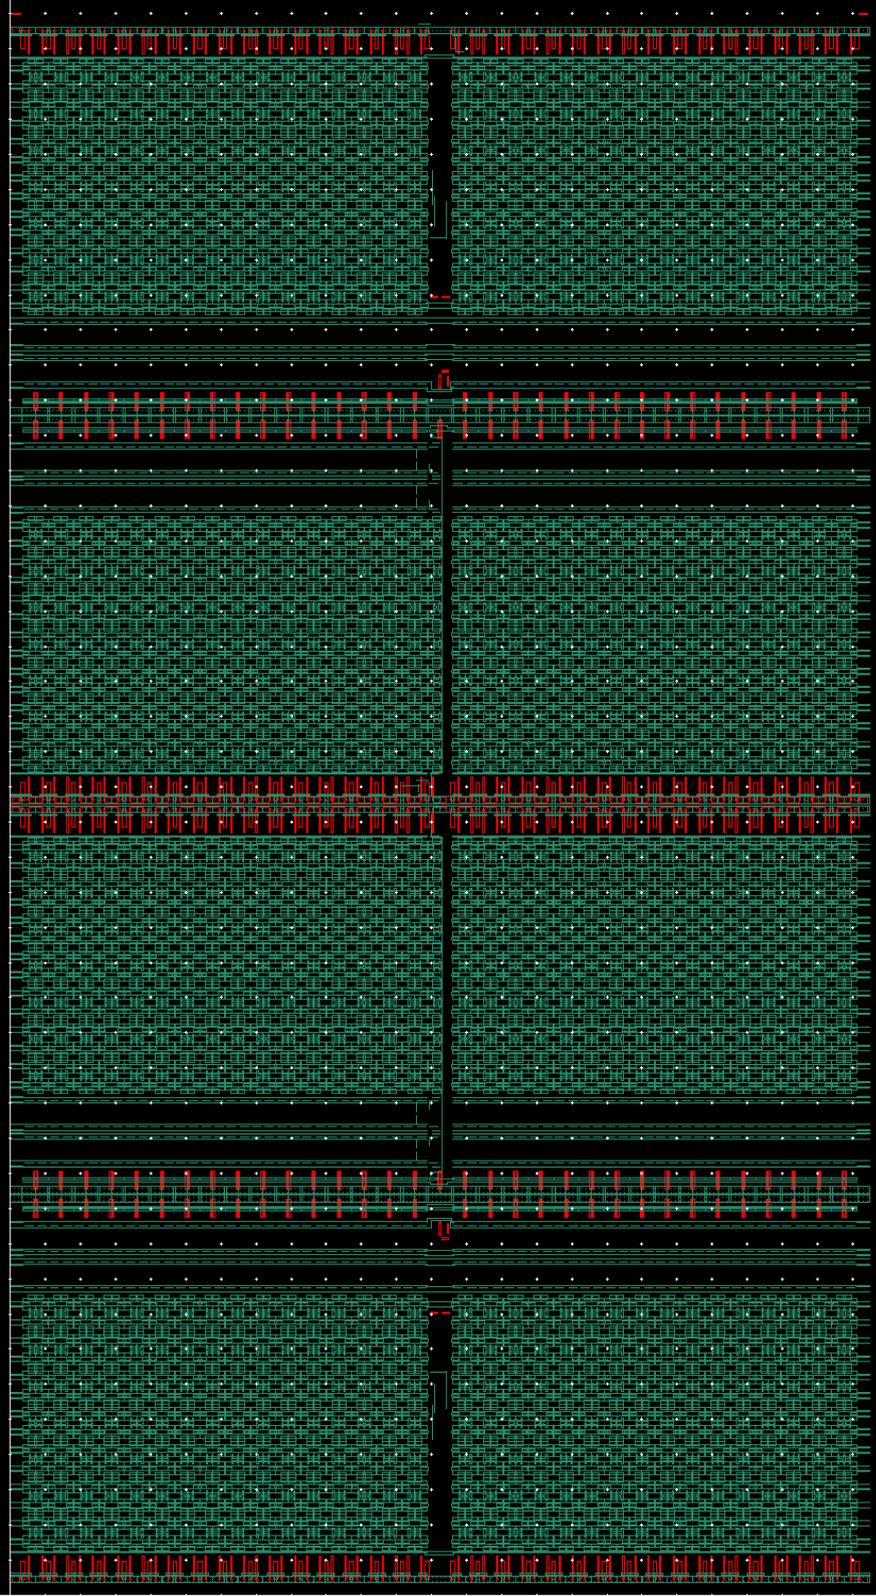


64 patt
CAM block

← metal 1

Metal 2 →



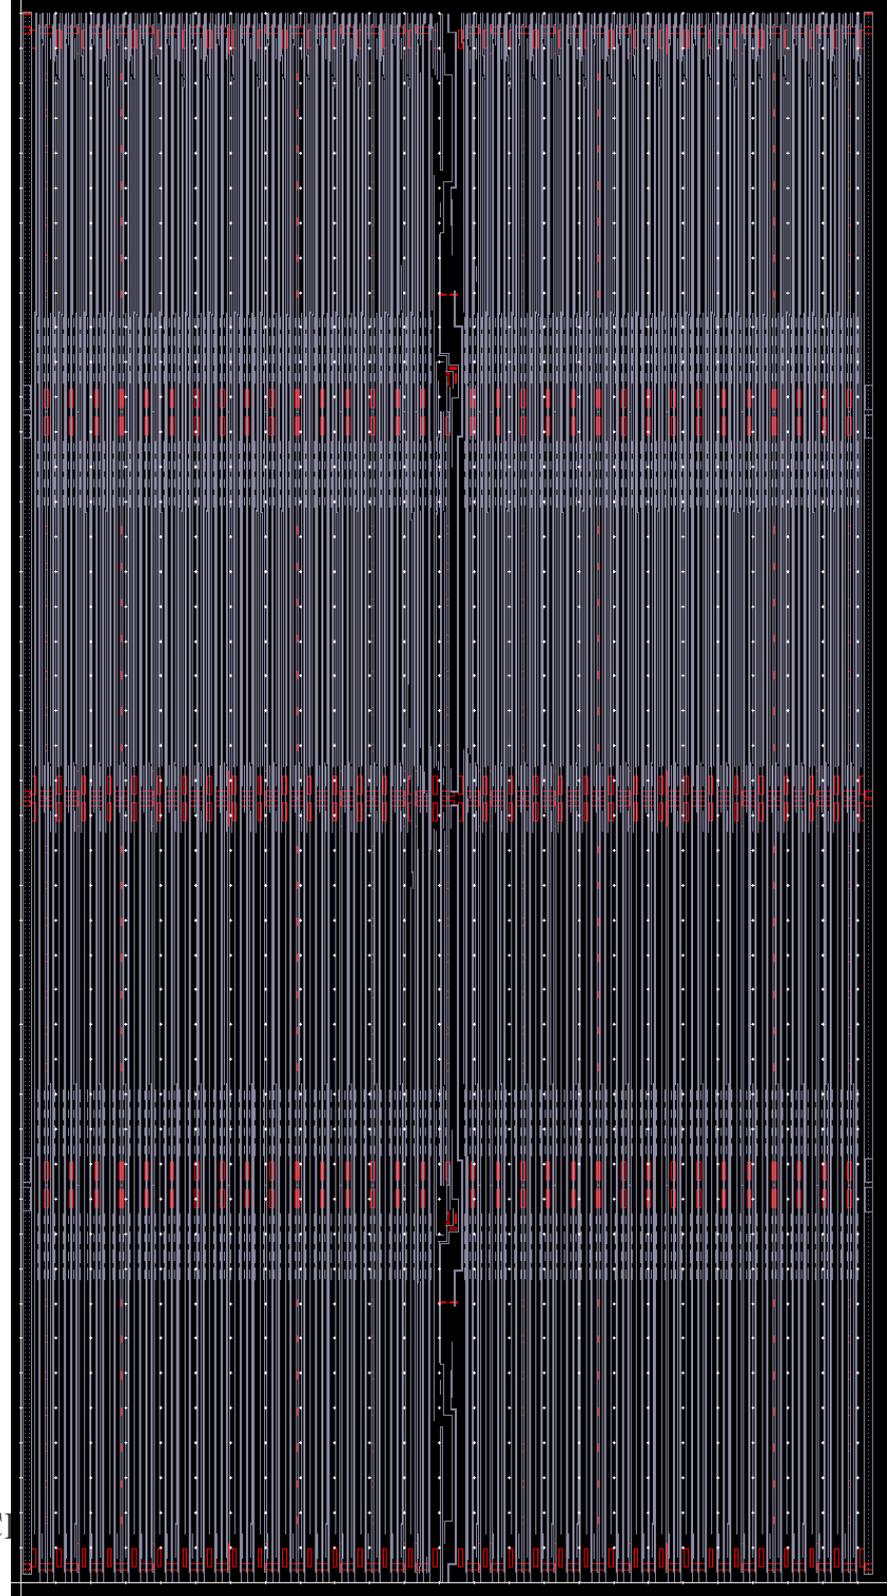


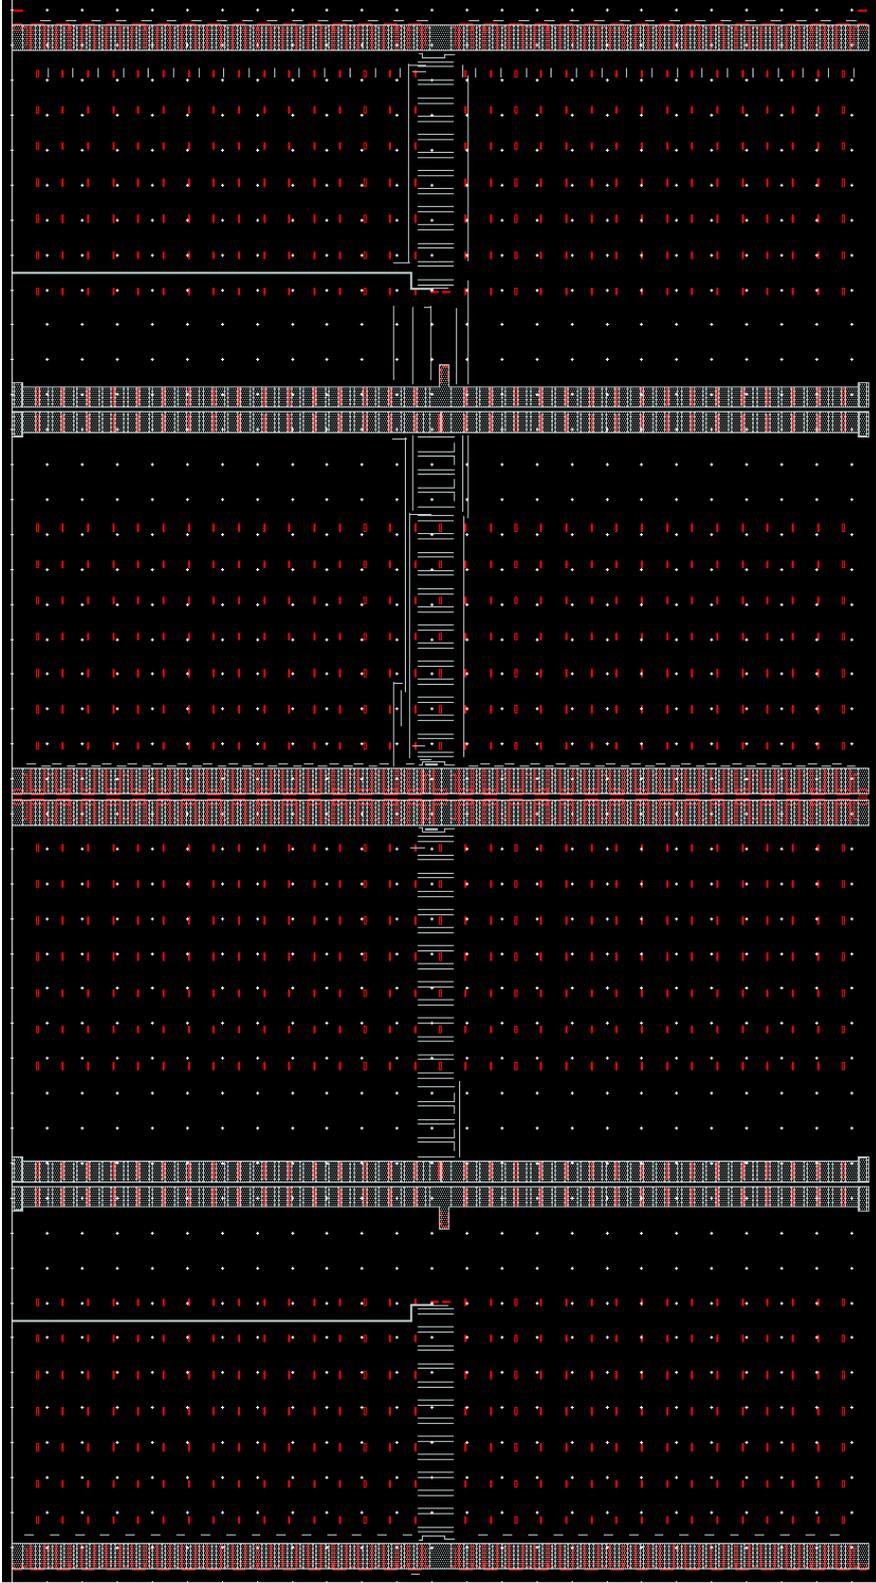
64 patt
CAM block

← metal 3

Metal 4 →

Crescioli (INFN Pisa/CI)





64 patt
CAM block

← metal 5

trans →

