

The VMM ASIC

George Iakovidis

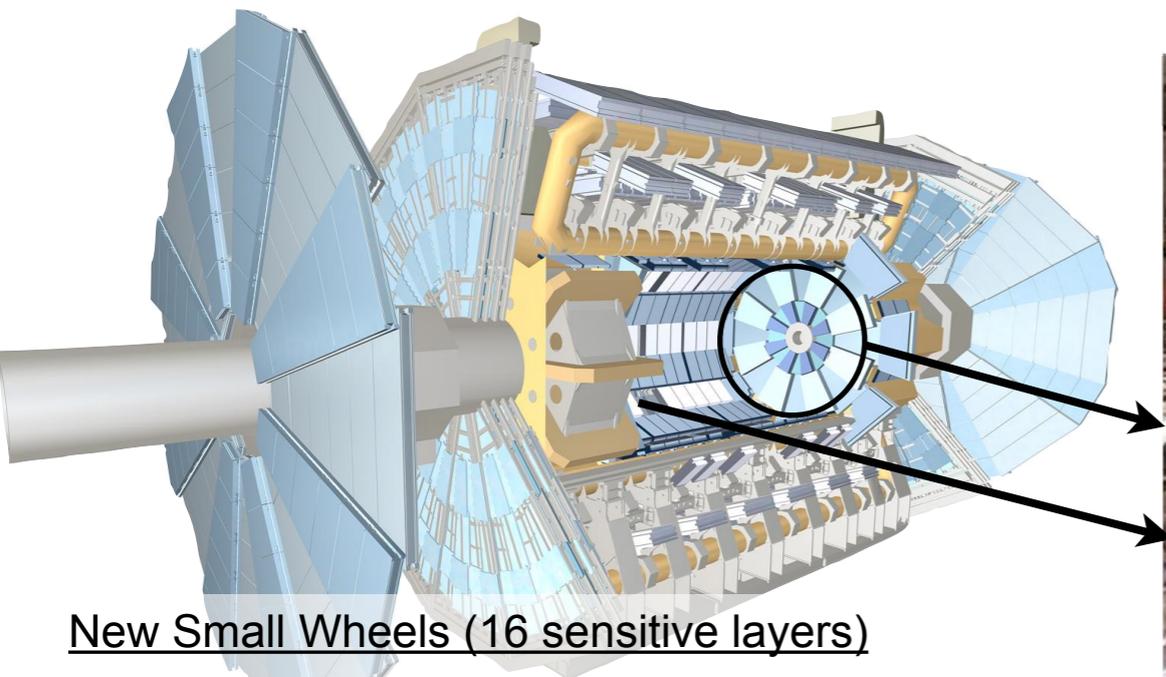
**BROOKHAVEN**
NATIONAL LABORATORY

 U.S. DEPARTMENT OF
ENERGY

BROOKHAVEN SCIENCE ASSOCIATES

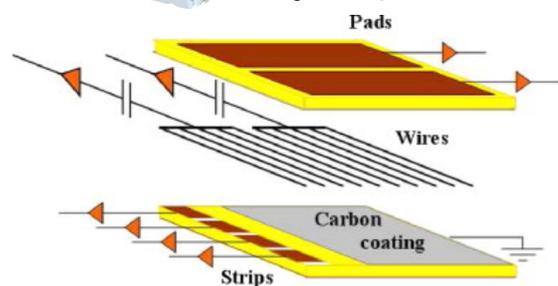
The ATLAS New Small Wheel(s) Upgrade

- The motivation for developing the VMM ASIC was the ATLAS NSW upgrade at CERN
- Biggest MPGD development and the most complex part of the ATLAS muon spectrometer ever built

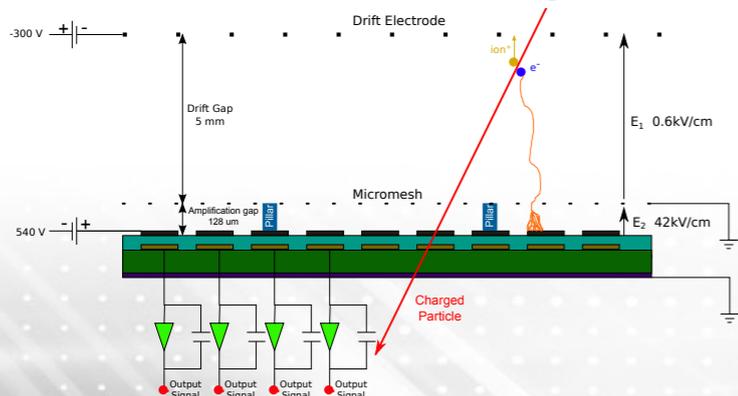
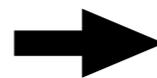


New Small Wheels (16 sensitive layers)

- sTGC (8x layers)



- Resistive strip Micromegas (8x layers)



~10m



Front-end Electronics Requirements (need of custom ASIC)

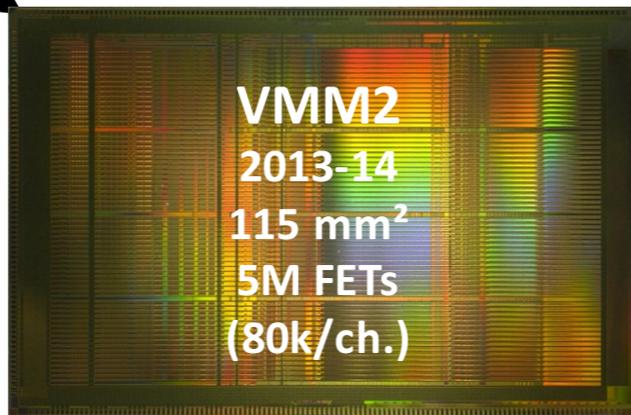
- Challenge of this Project - More than 2.4 million channels total (2.1M for Micromegas and 300k for sTGC) (full MS of ATLAS ~1.6M channels)
- Operate with both charge polarities
- Sensing element capacitance 50-200pF (sTGC Pad up to 3nF)
- Charge measurements up to 2pC @ < 1fC RMS (6pC for sTGC pads)
- Time measurements ~ 100ns @ < 1ns RMS
- Multiple Trigger primitives, complex logic
- Digitisation, deep FIFOs, Low power, programmable
- Space requirements on the detector
- Radiation tolerant

The VMM front-end ASIC - Evolution



- ✓ Mixed-signal
- ✓ 2-phase readout with external ADC
- ✓ **peak and timing** information
- ✓ neighbouring readout
- ✓ sub-hysteresis **discrimination**
- ✓ few timing outputs

- ✓ Mixed-signal
- ✓ **Continuous** readout
- ✓ Current-output peak detector
- ✓ **Increased** range of gains
- ✓ **Three ADCs** per channel
- ✓ FIFOs, **serialised data with DDR**

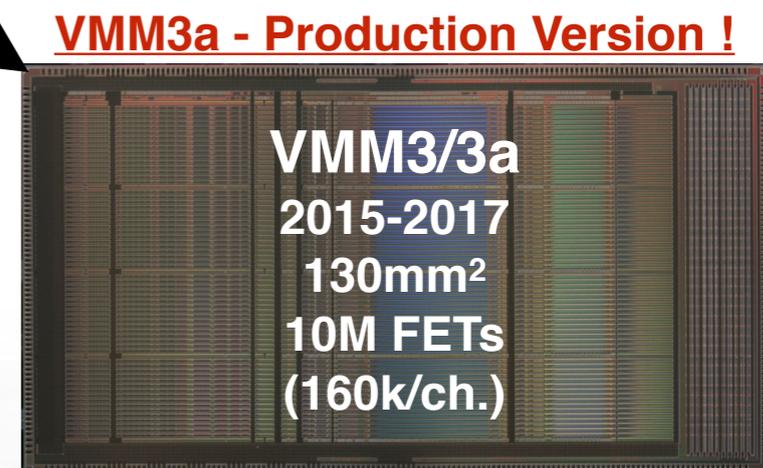


- ✓ Serialised ART with DDR
- ✓ Additional timing modes
- ✓ **64 timing outputs**
- ✓ Additional functions and fixes

- ✓ **LVL0 pipeline** and buffering for ATLAS
- ✓ **SEU-tolerant logic**
- ✓ **Revised front-end** for high charge and capacitance (2nF, 50pC, fast recovery)
- ✓ SLVS signals
- ✓ Reset controls
- ✓ **Timing at threshold**
- ✓ Timing ramp optimisation
- ✓ Ion **tail suppressor** (fast recovery)
- ✓ Int. Pulser range extension
- ✓ ART synchronisation to BC clock
- ✓ **VMM3a fixed open bugs** from VMM3 and introduce some stability fixes on the ADCs and Front-end

- ★ *The VMM was designed at BNL G. De Geronimo in collaboration with IFIN-HH S. Martoiu (L0) (recent version VMM3a, at DG Circuits on behalf of BNL)*
- ★ *It is fabricated in the 130nm Global Foundries 8RF-DM process (former IBM 8RF-DM)*

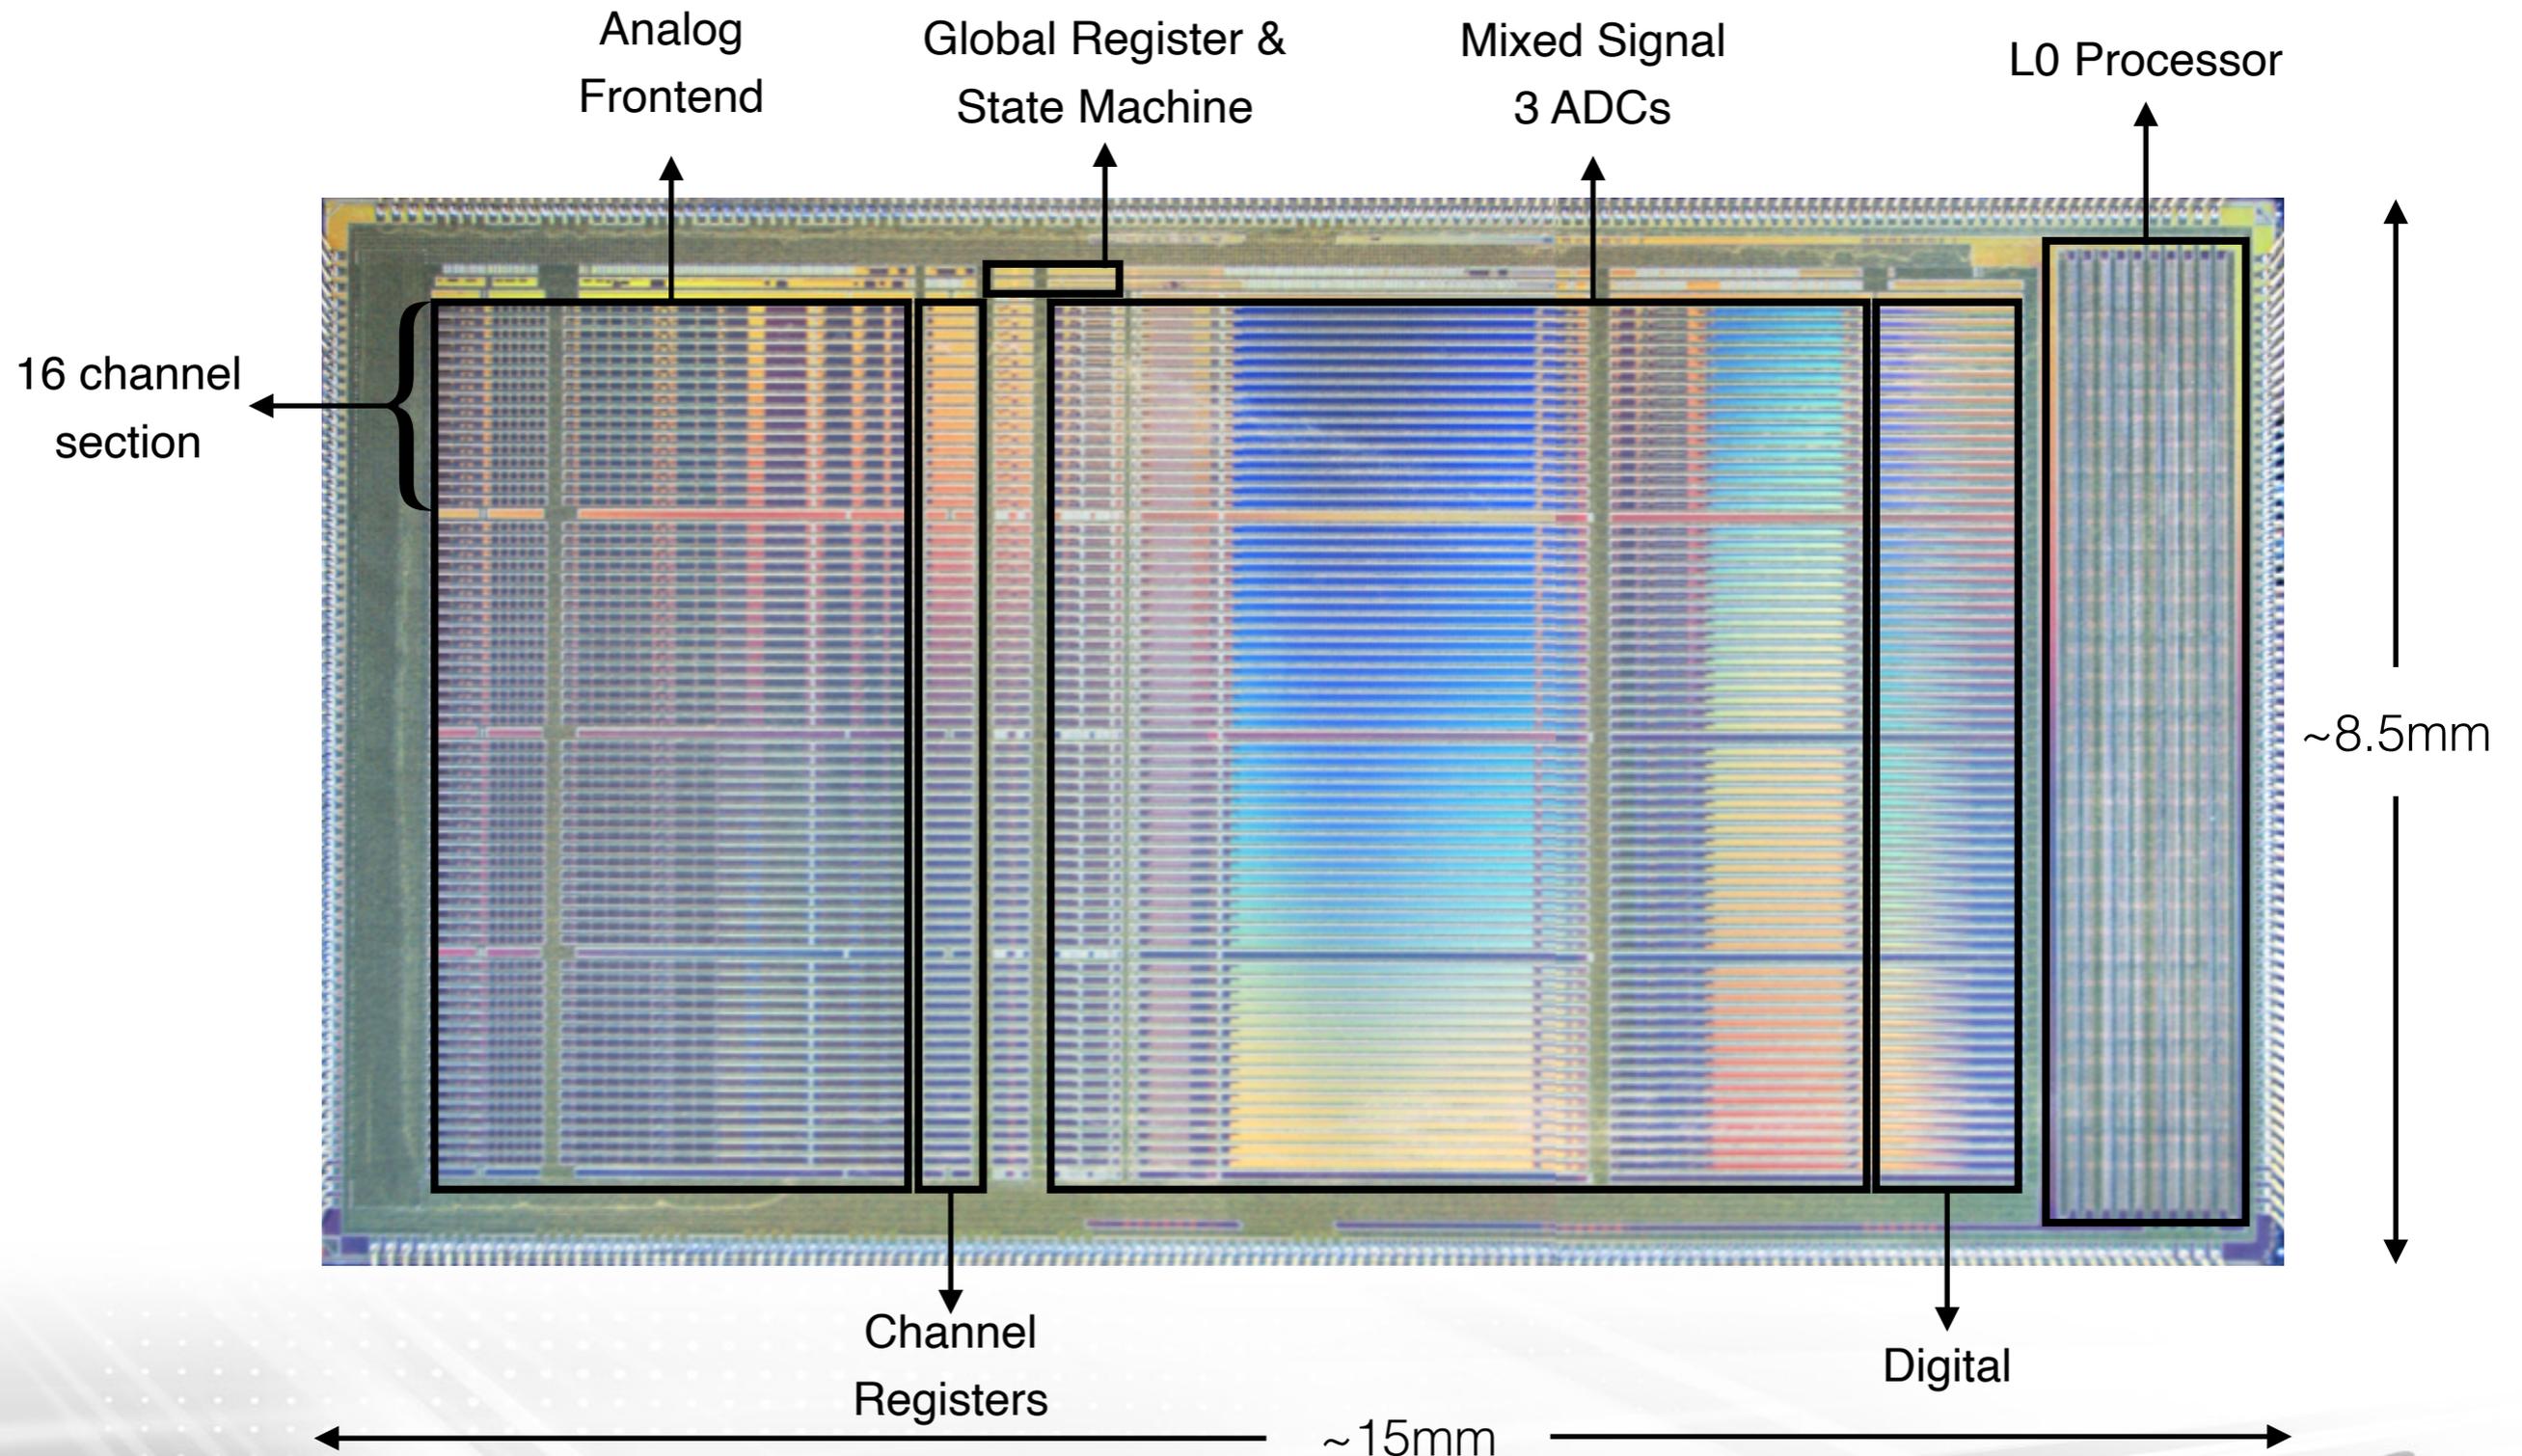
BGA 400, 1mm



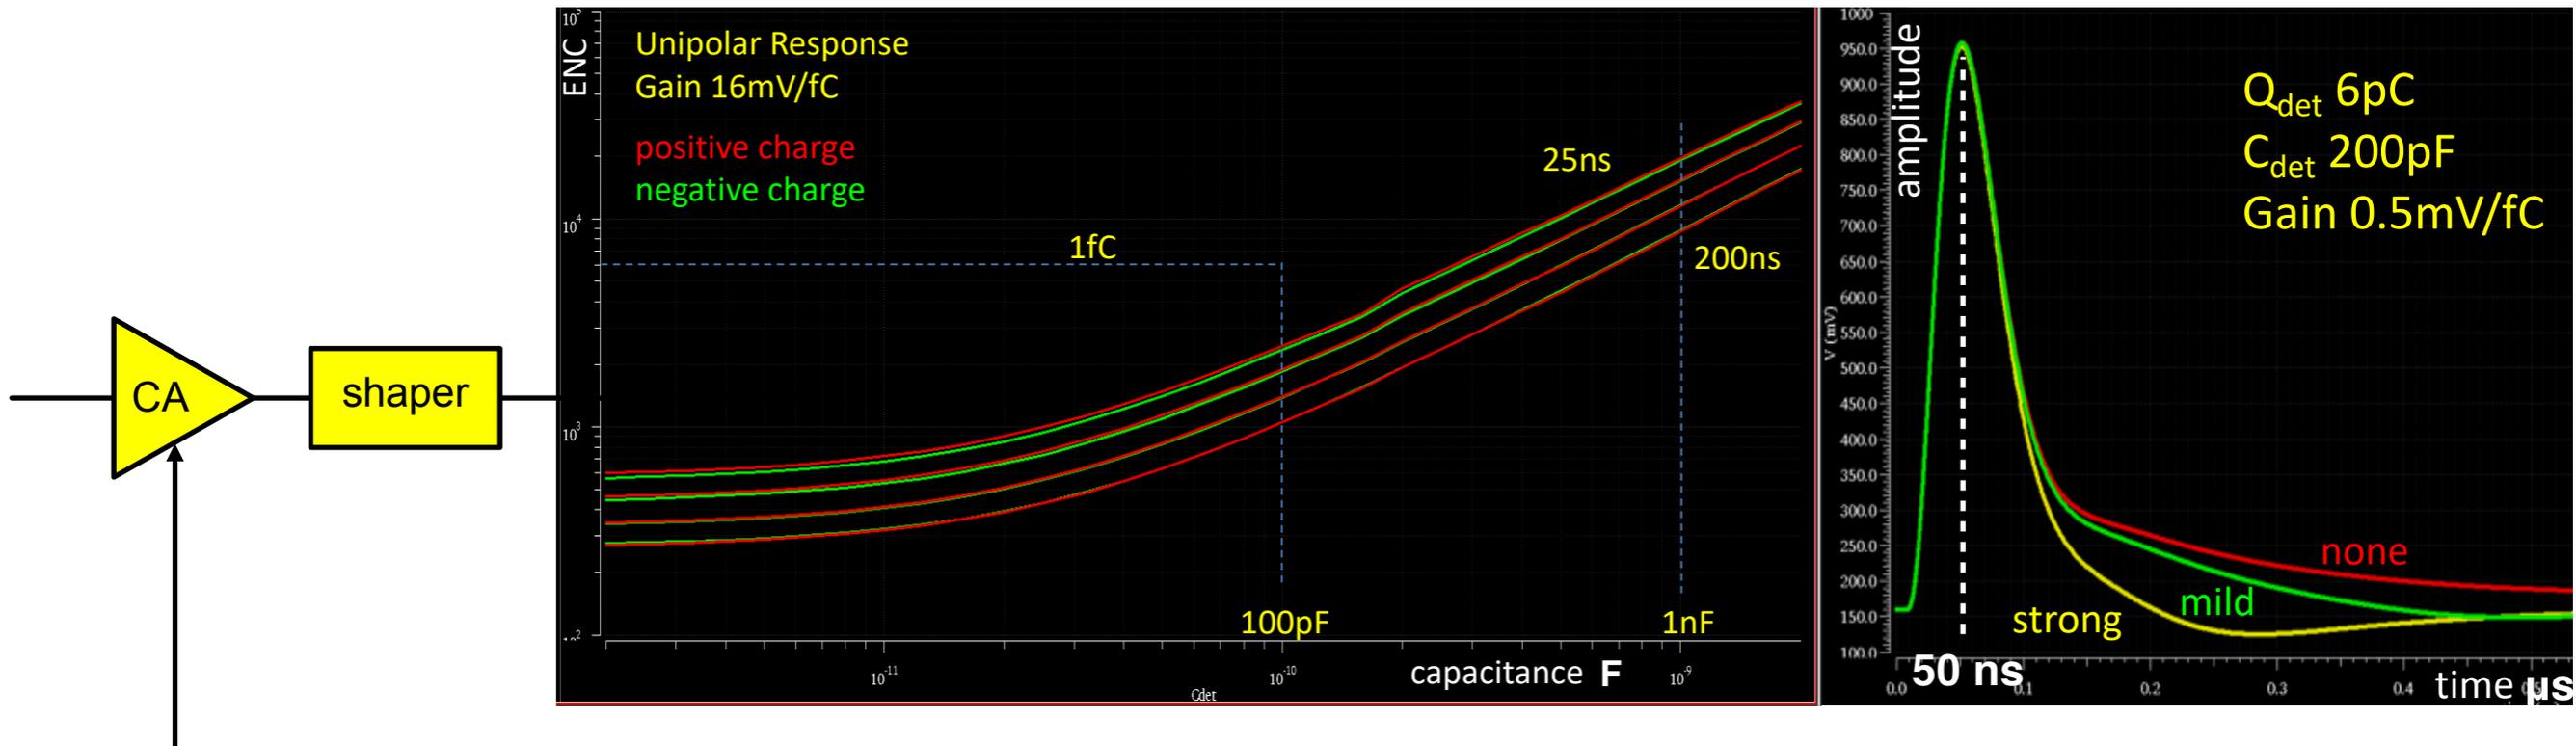
~10mW/channel

An actual photo of the ASIC

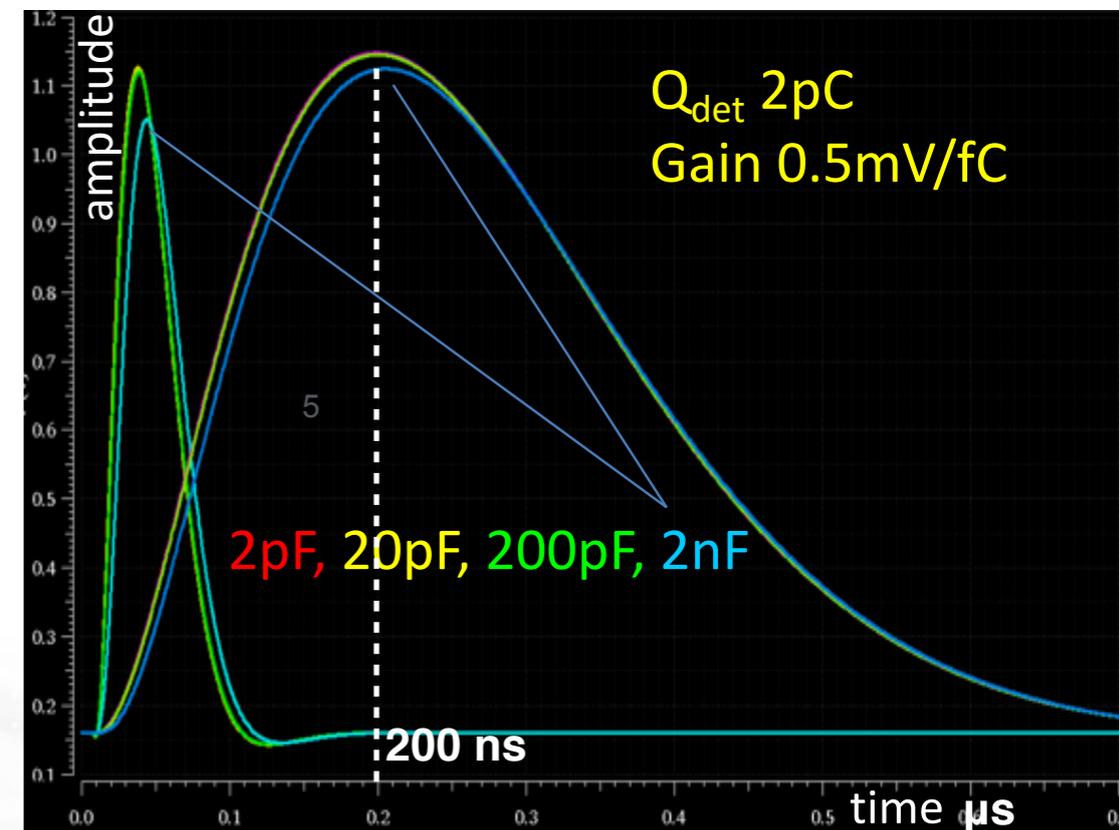
- The ASIC features **64 channels** that extend along the size of the die. At the end the L0 section (explained in later slides) is separated to isolate the noise from the digital activity



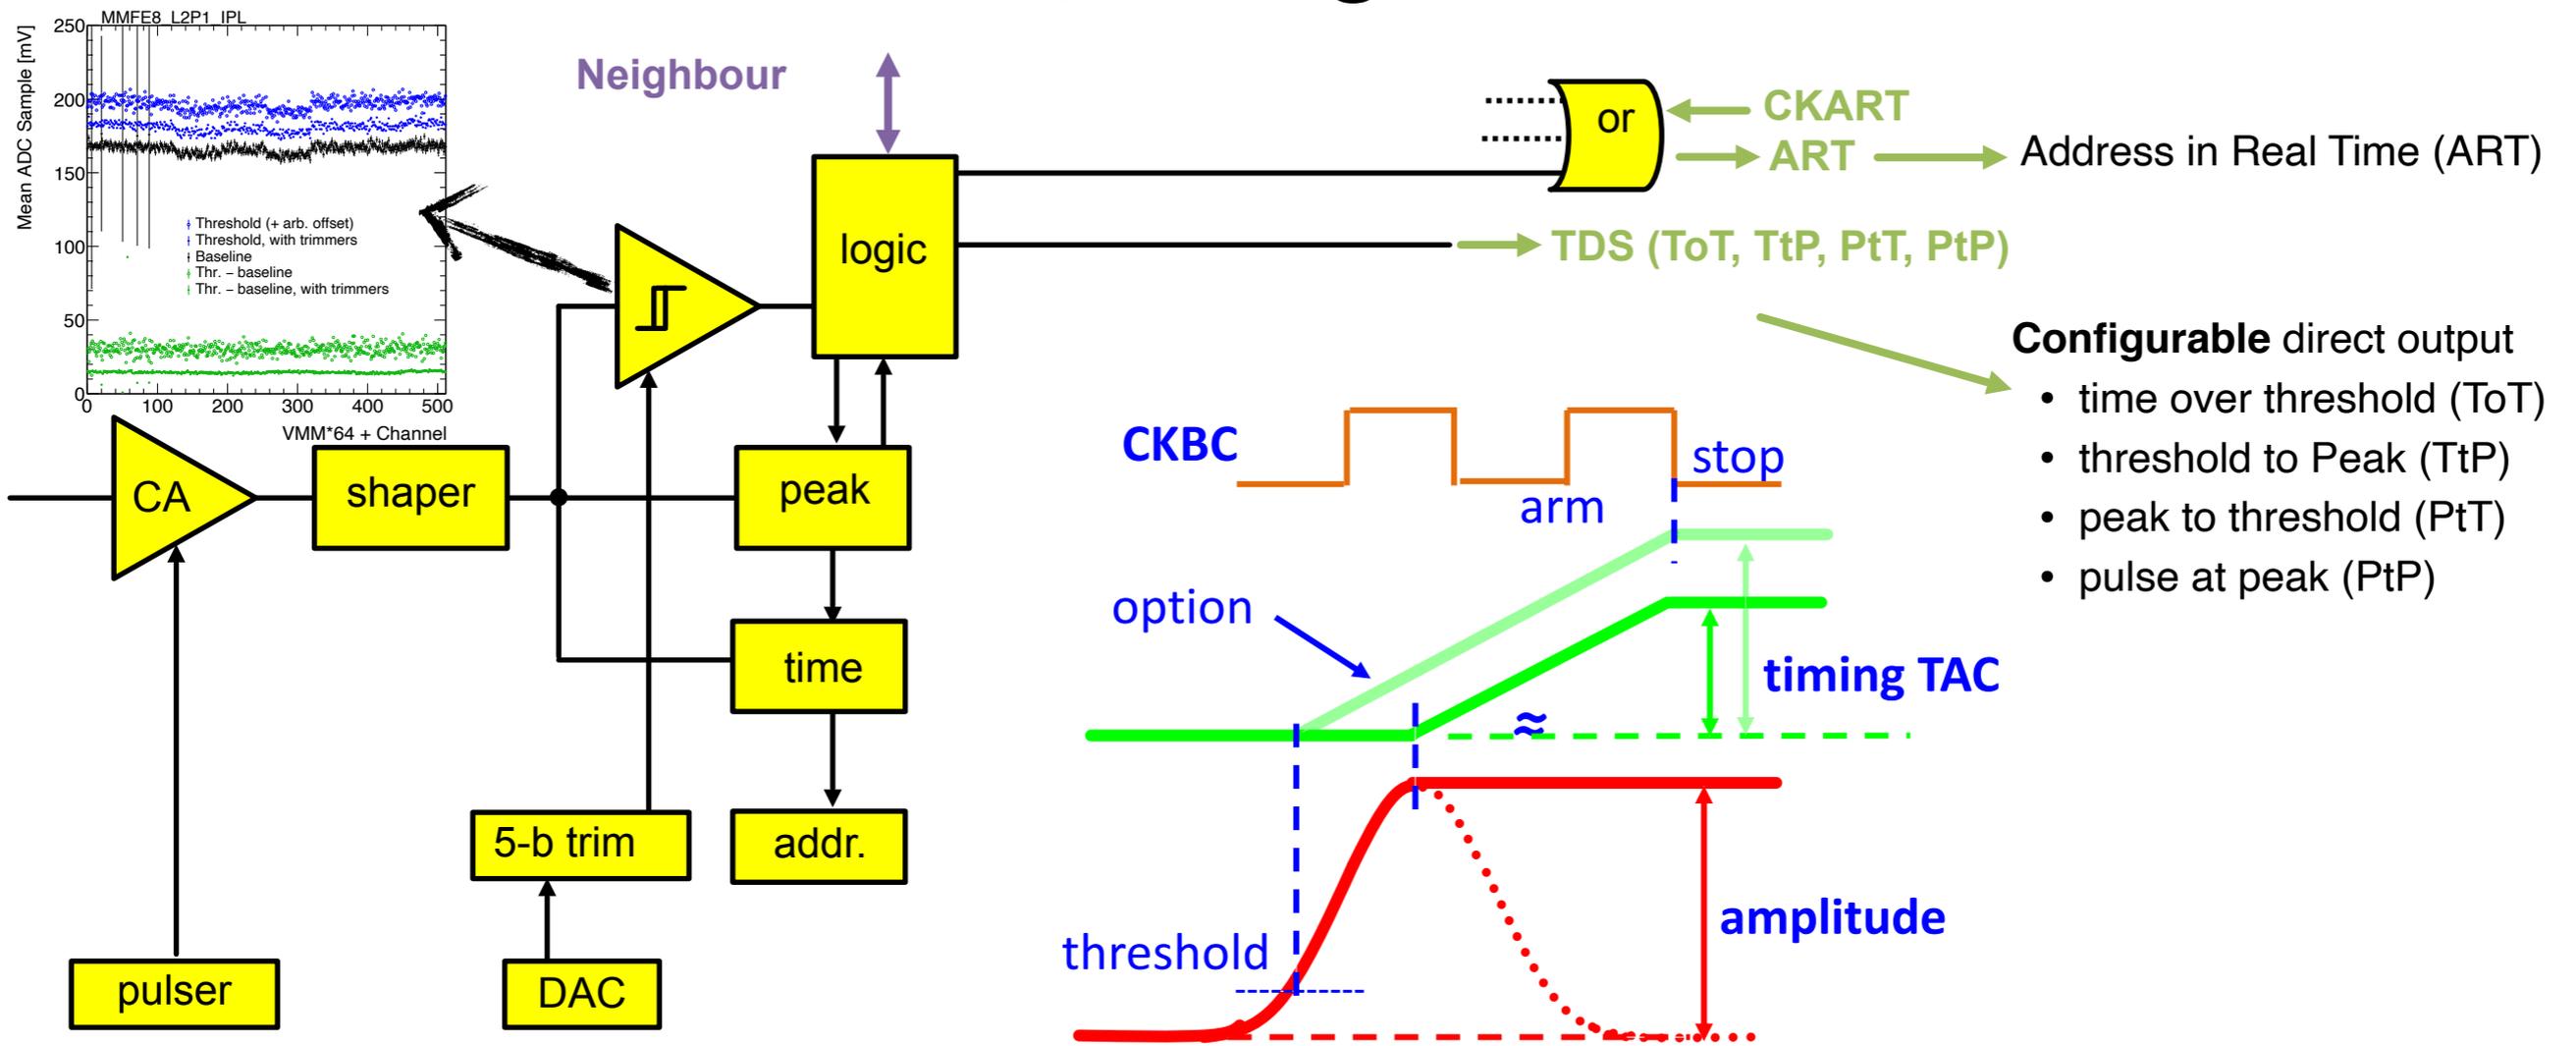
VMM3a Amplifier & Shaper



- Input transistor: **PMOS**, 3 stage amplifier,
 - 2 stages used for **adjustable gain**: 0.5, 1, 3, 4.5, 6, 9, 12, 16 mV/fC
 - last stage for **charge inversion** (positive or negative)
- Input **capacitance**: can operate from sub-pF to several nF
- Maximum **charge**: 2 pC in **linear range**, **fast recovery** from 50 pC
- Semi gaussian DDF c-shaper **3rd order**
 - **Configurable** ion tail **suppression**: none, mild or strong
 - **Adjustable peaking** time: 25, 50, 100, 200 ns
 - Leakage-adaptive, **BGR-stabilised baseline**

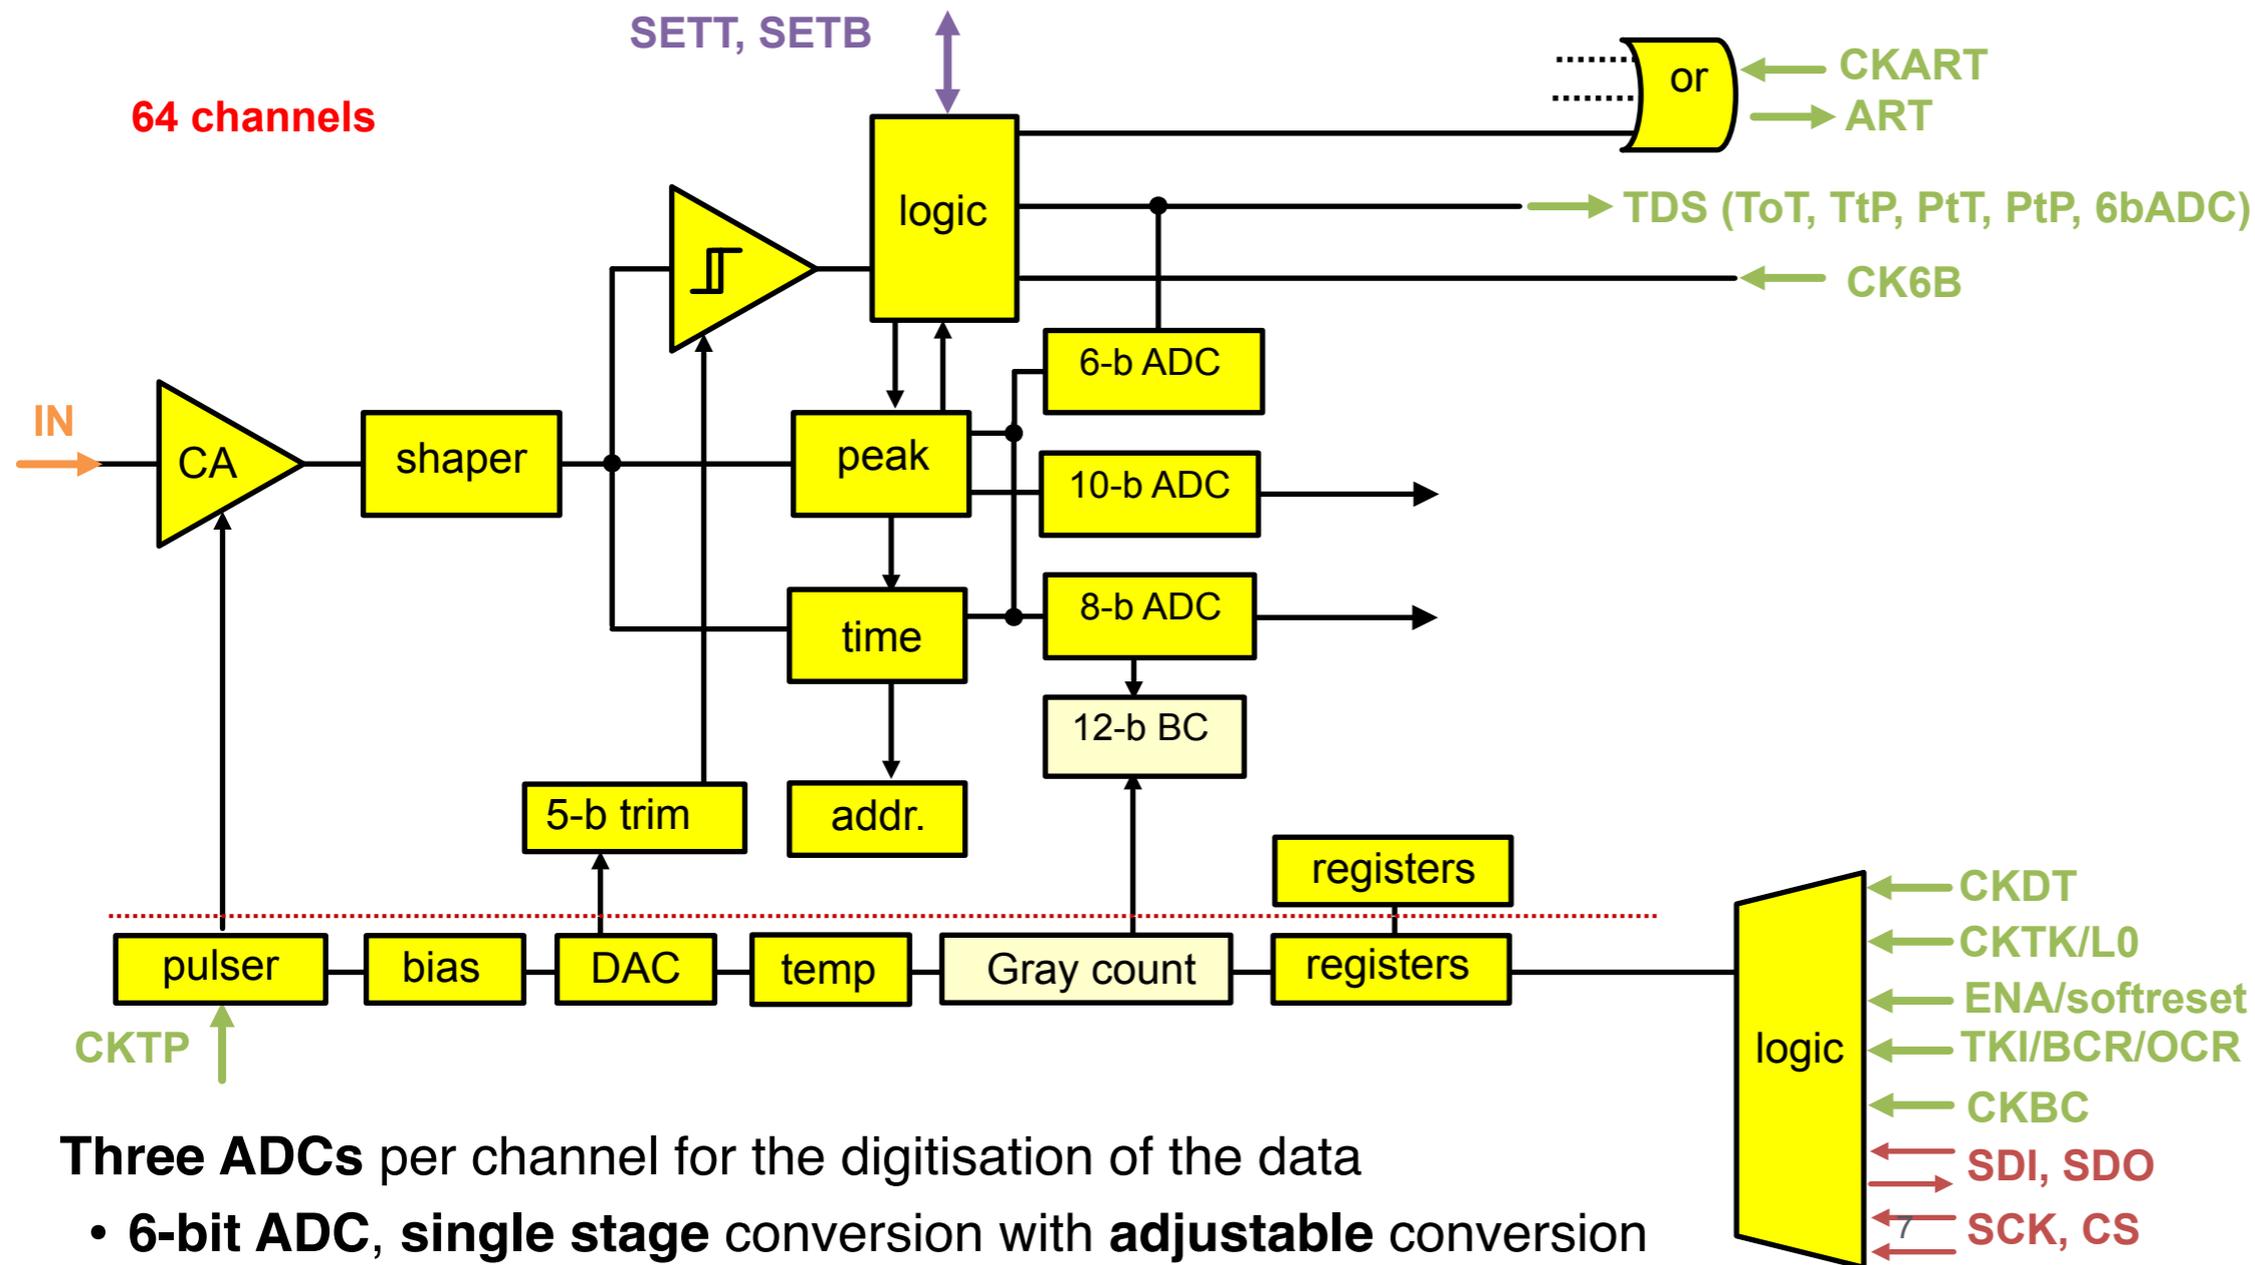


VMM3a Discrimination, Charge and Time



- Global 10-bit DAC for adjusting the **threshold - Discrimination** with sub-hysteresis (effective 2mV)
- Adjustable **5-bit discrimination** threshold **per channel** to adjust at ~mV level
- **Neighbour** logic to trigger sub-threshold channels with inter-chip communication
- Configurable **direct output** per channel and serial fast output of address as an OR of all channels
- **Peak detection**: measurement of peak **amplitude** and storage in analog memory
- **Time detection**: measurement of **peak/threshold** timing through a configurable time to amplitude converter (**TAC**: 60, 100, 350, 650 ns) and storage in analog memory
 - Clock working mode on **synchronous** machines but also as strobe for **asynchronous** operations

VMM3a ADCs

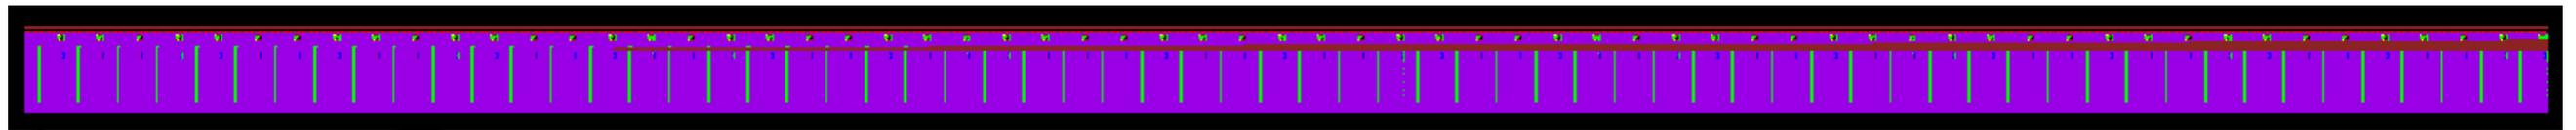


Three ADCs per channel for the digitisation of the data

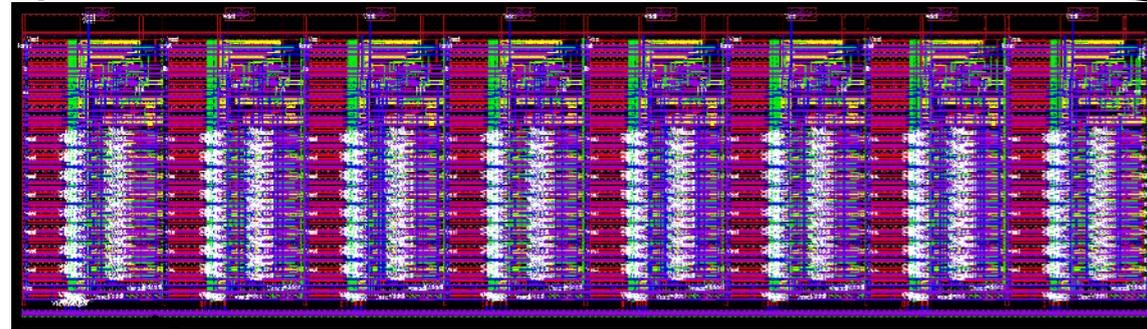
- **6-bit ADC, single stage** conversion with **adjustable** conversion time and **offset**, completes within **25 ns** from **peak**
- **10-bit ADC, 200 ns adjustable** conversion time/offset, for peak **amplitude** conversion
- **20-bit timing information** with 8-bit ADC, **100 ns** conversion time + 12-bit Gray-code counter, BC clock
- **2 step mode conversion for 10-bit & 8-bit ADCs** - First stage the comparison identifies one of the macro-cells and at the second stage the micro-cell is identified, possibility to jump through macro-cells

10-bit Current-Mode Domino ADC - Functionality

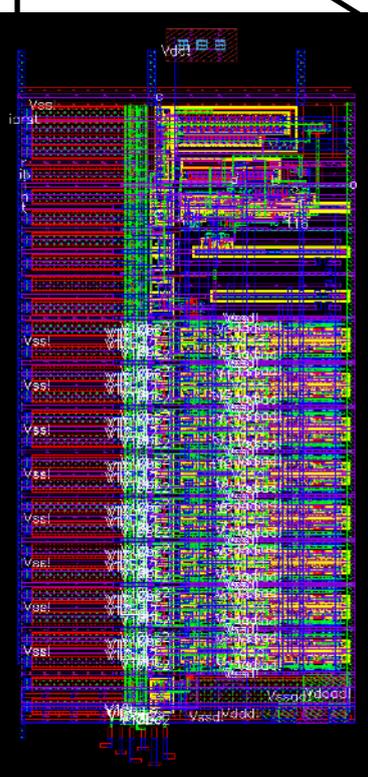
ADC Cells



8x macro-cells

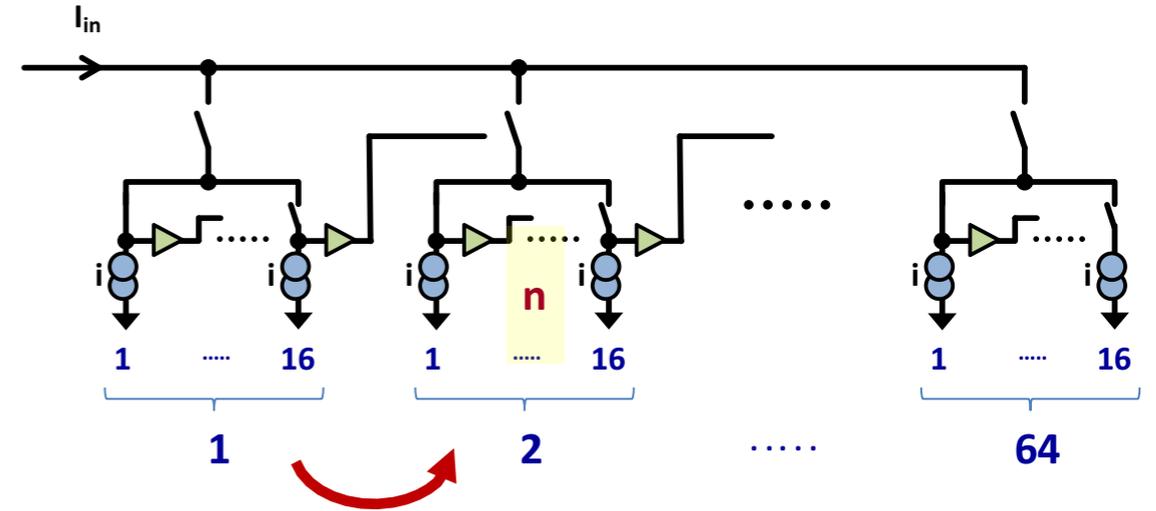
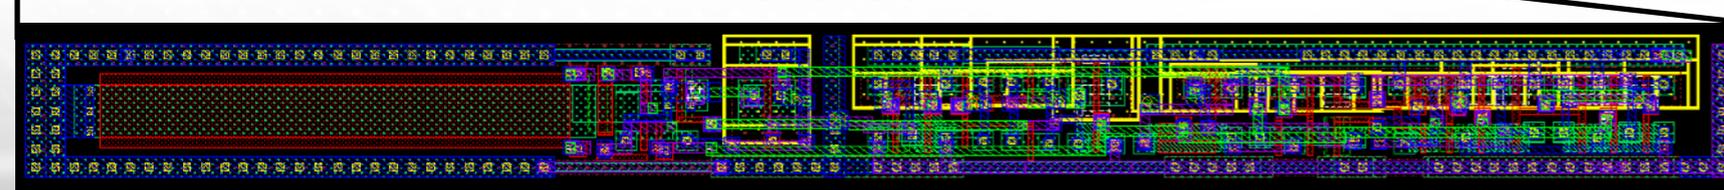


1 macro-cell



16x micro-cells

1x micro-cell

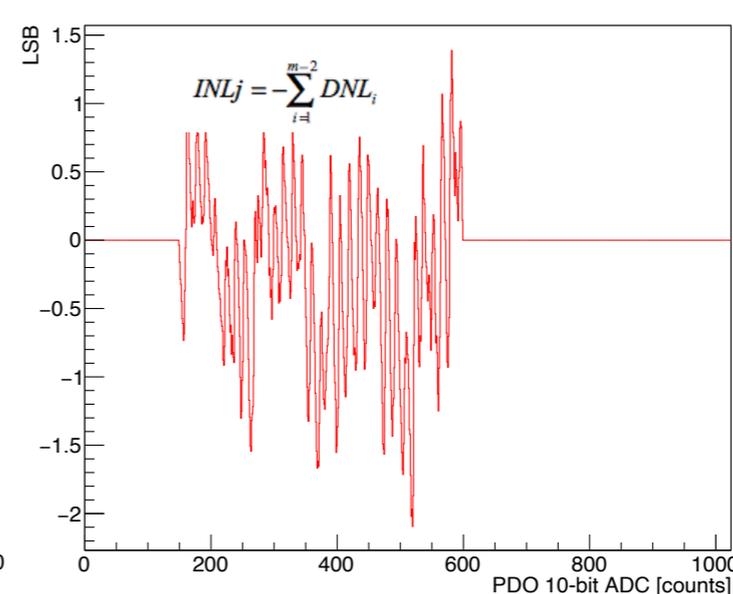
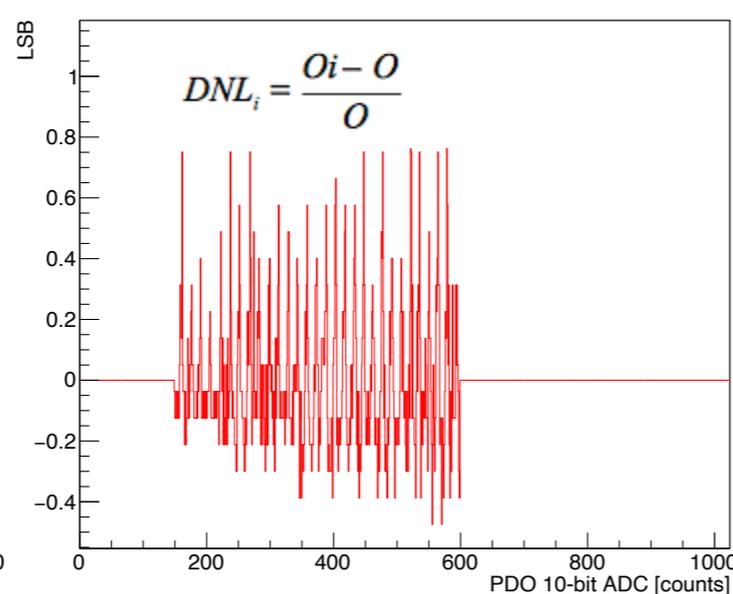
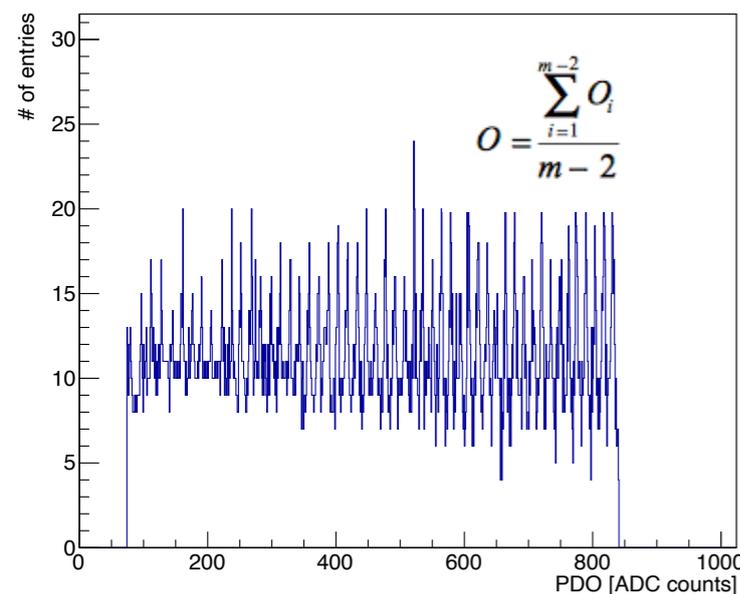
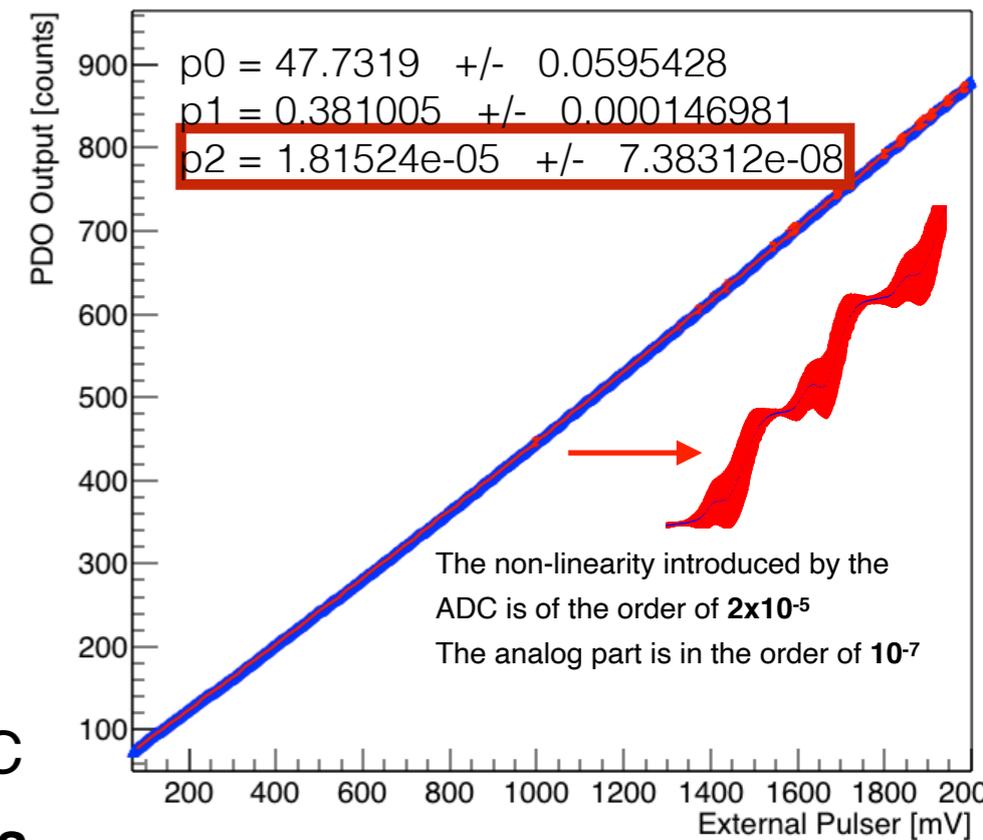


ADC Core (per channel)

- **1024** current sources (similar to a digital thermometer)
- **64** macro-shells (6 upper bits xxxxxx0000), **16** micro-shells (4 lower bits 000000xxxx)
- **2 step mode** - First the comparison identifies one of the 64 macro-cells. Then on second step the lower 4 bits are identified (250ns adj conversion time)
- **8 bit ADC** is built in the similar way (5+3)
- **6 bit ADC** is a **single stage conversion** similar to the 64 macro shells with fast digitisation (50ns)

ADC Performance (eg 10-bit ADC)

- Even from the VMM2, the ADCs have been proven one of the most difficult points on VMM
- ADCs on VMM2 showed missing codes, high accumulation and degradation of performance
- Situation was significantly improved on VMM3a which addressed several problem but still not perfect
- Using **DNL** and **INL** is calculated and used to estimate the **ENOB** of the 10-bit ADC
- **Equivalent number of bits ~7.5 (noise free)** for the 10-bit ADC
- **Performance was considered enough for gaseous detectors (schedule constrains as well) - moved to production**

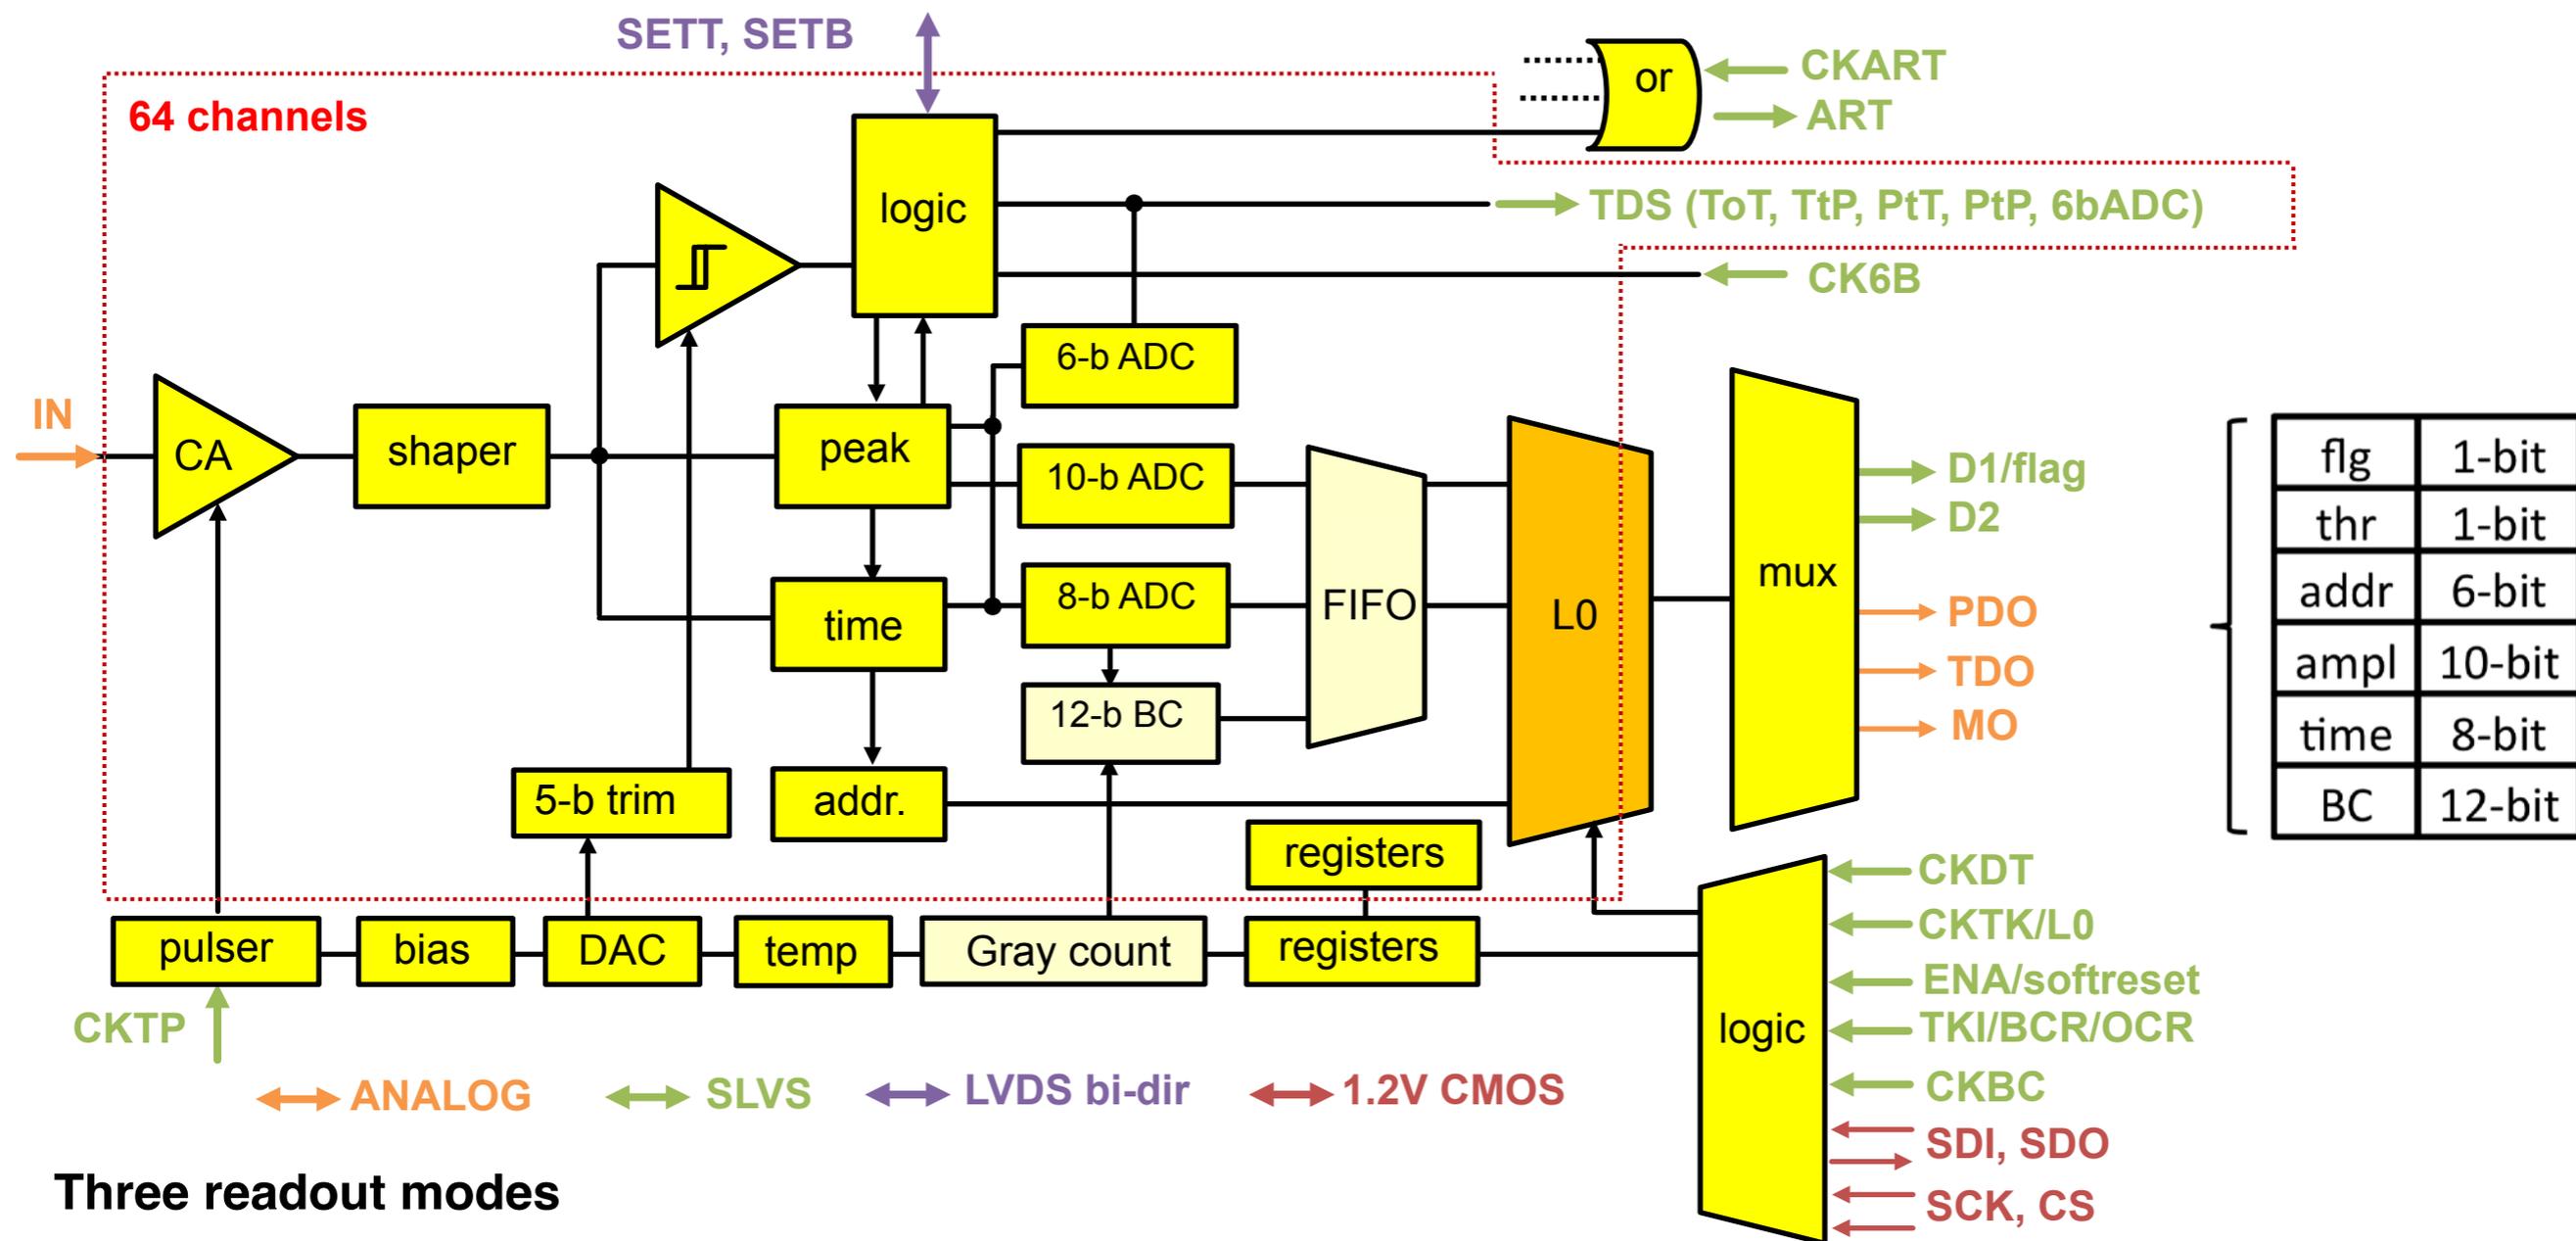


$$\sigma_c = \sqrt{\frac{1}{12} + \frac{1}{m-2} \sum_{i=1}^{m-2} INL_i^2}$$

Then, ENOB can be calculated as

$$ENOB = \log_2 \frac{m}{\sigma_c \sqrt{12}}$$

VMM3a Readout & Overall Architecture

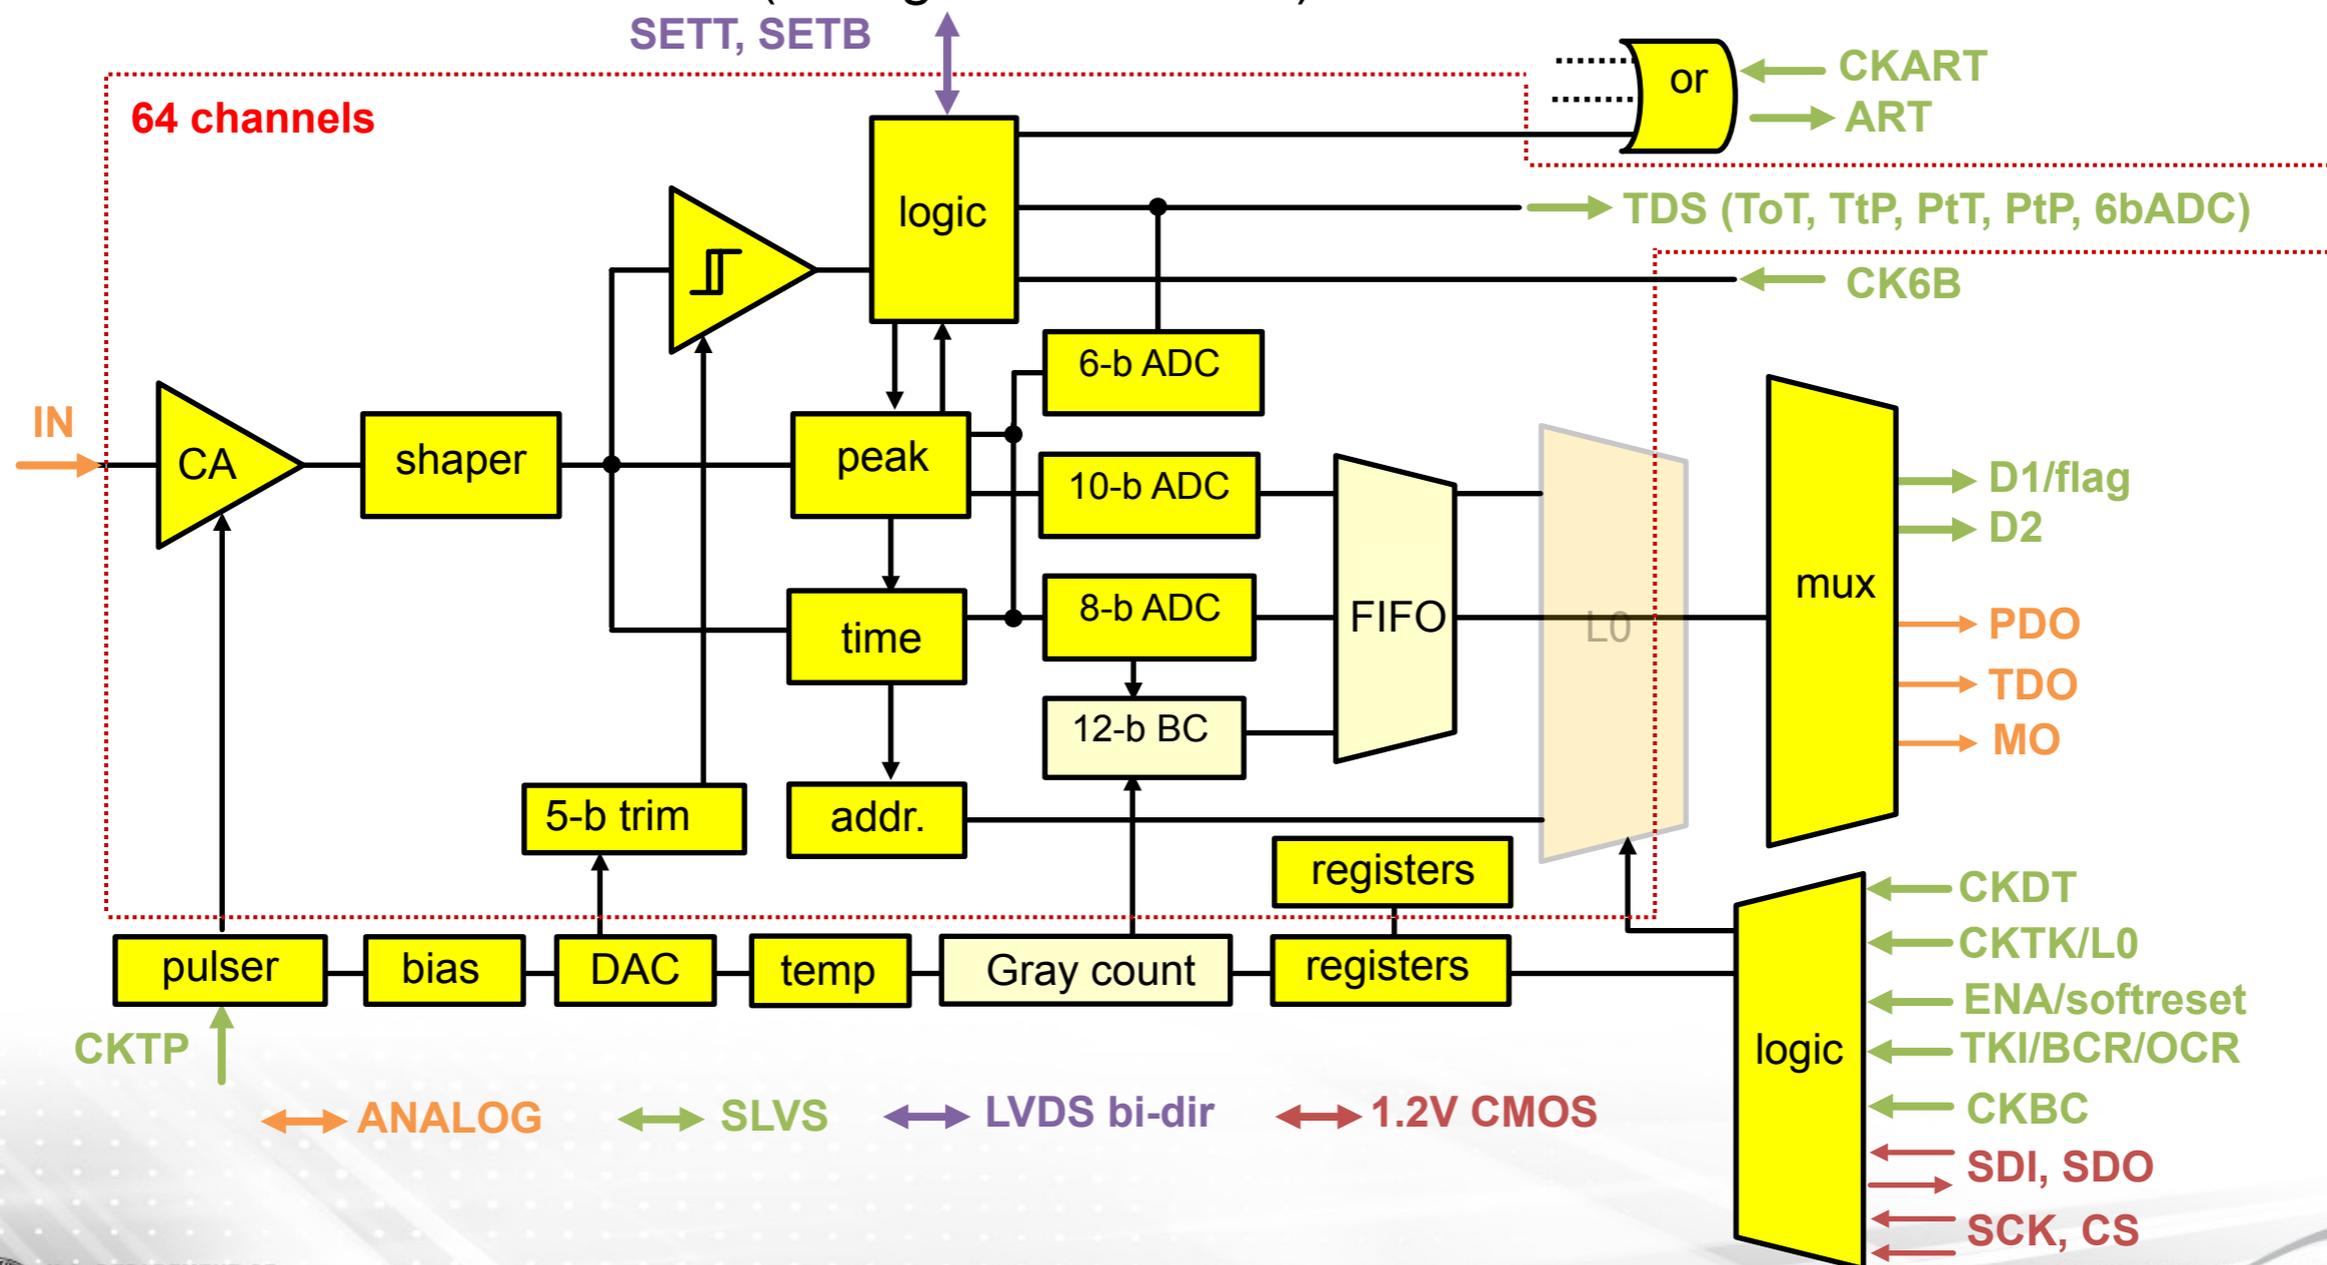


Three readout modes

- **Mixed** mode with peak & time analog output + address (**external ADCs**)
- **Digital continuous** with **internal ADCs** and 38-bit data at 2 outputs with 200MHz DDR, **triggerless** or with **external trigger** and auto reset
- **Level-0** processor **external trigger mode** with 64-deep latency FIFO programmable acceptance windows with 8b/10b encoding. Each channel could achieve 4MHz, total latency 16us

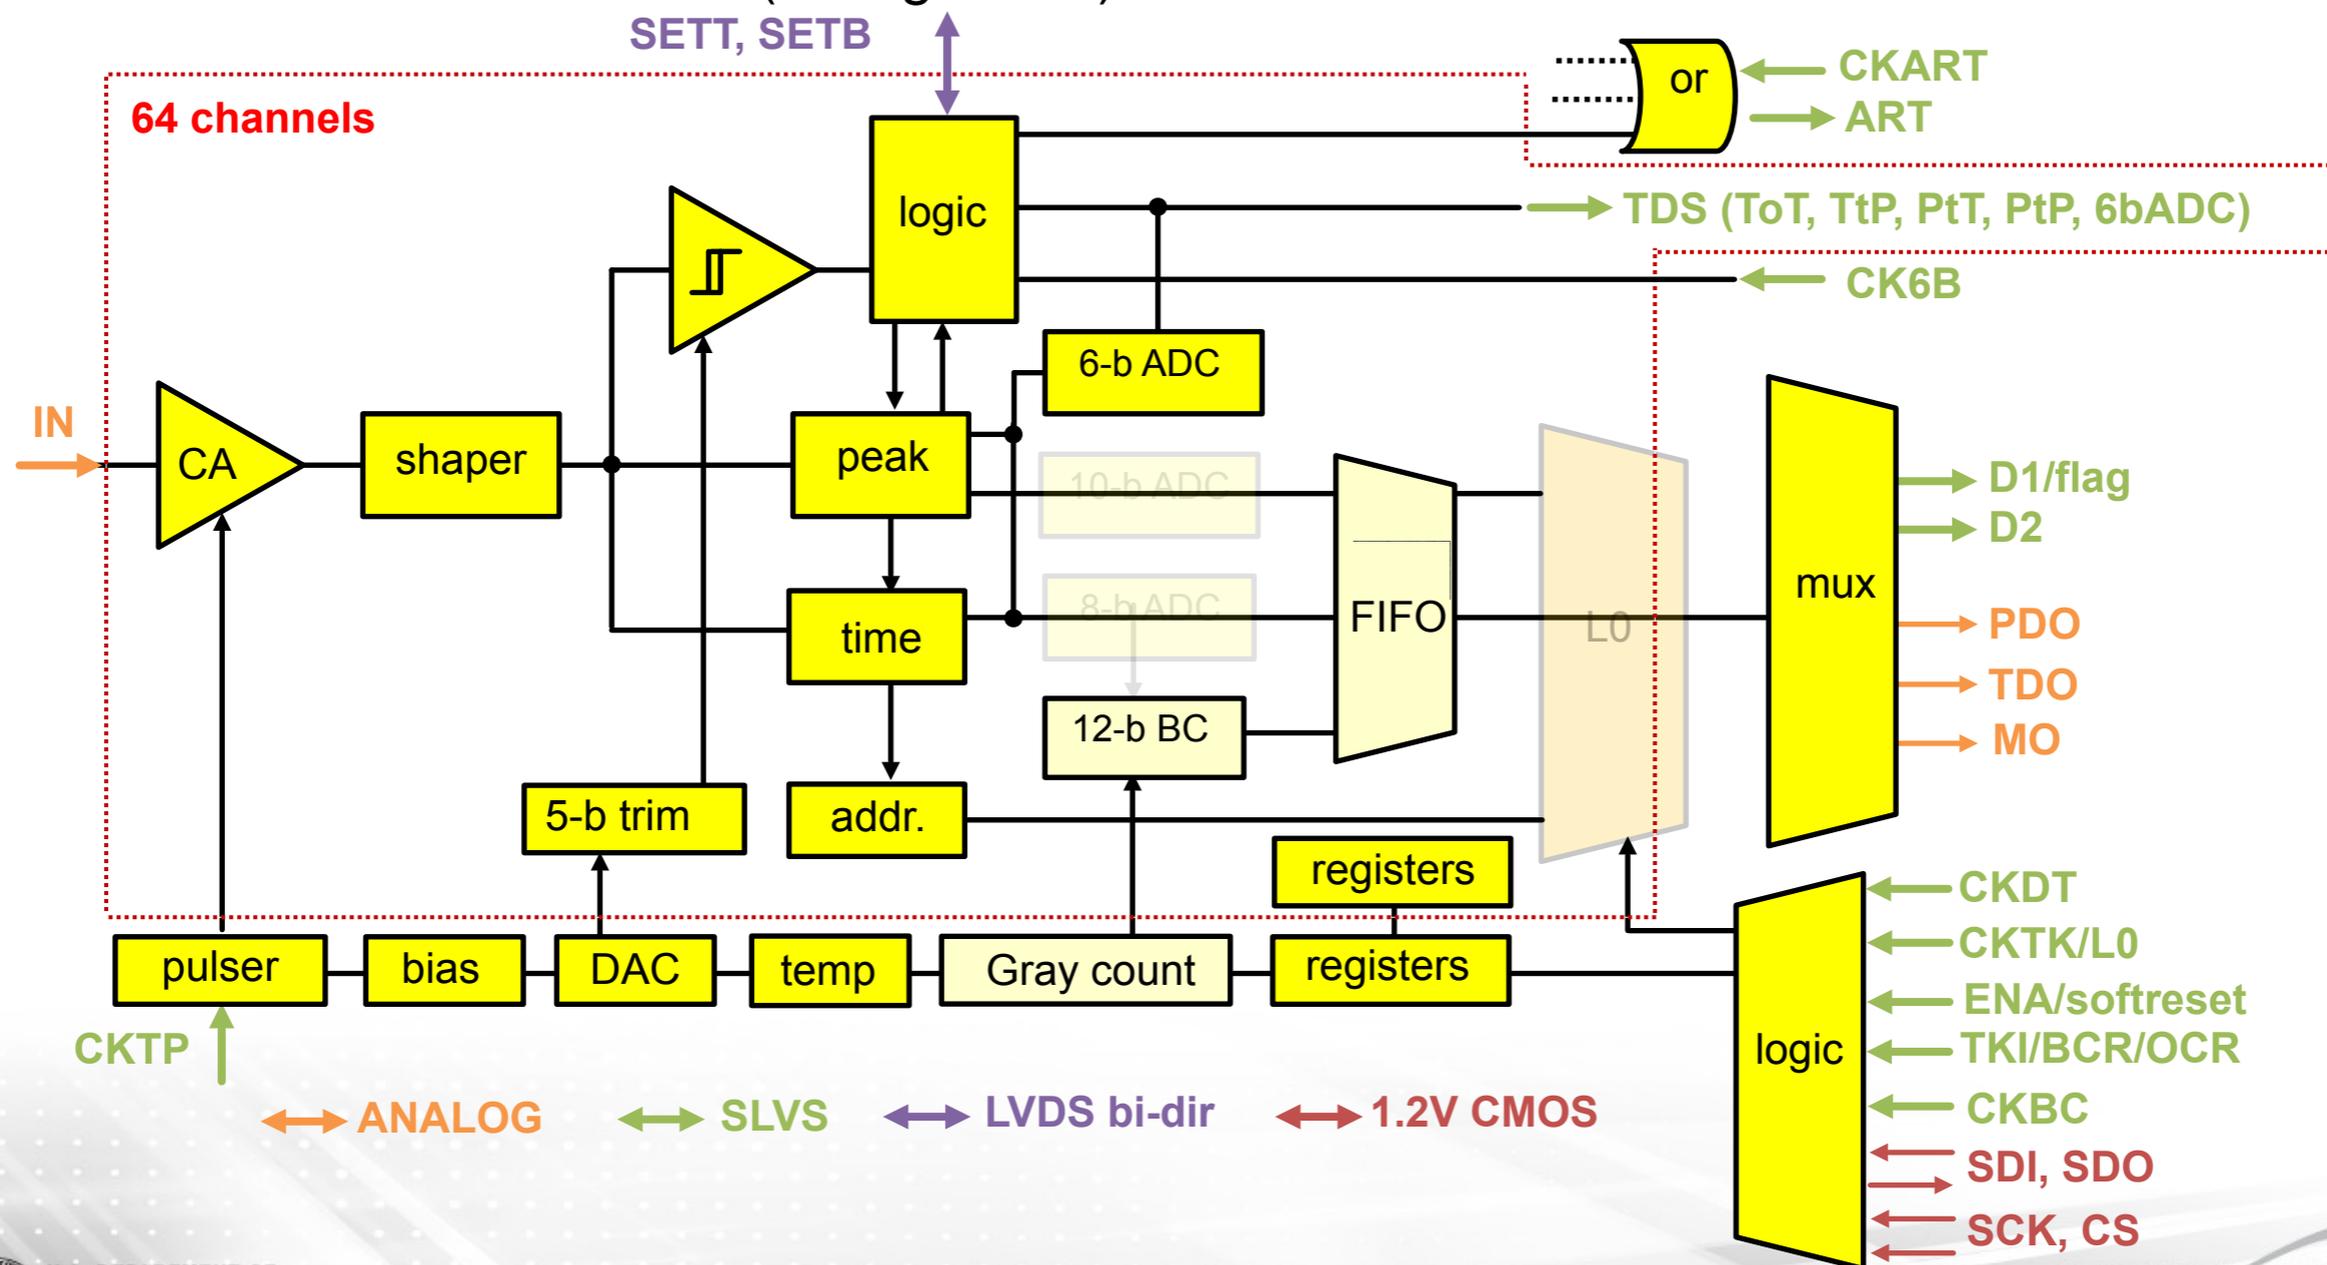
Modes of operation - Continuous

- This mode is **continuous trigger-less readout** (can be trigger driven as well see backup)
- All the outputs and inputs are active and **independent**
- 6bit ADC conversion within 25ns (configurable + reset)
- 10bit ADC conversion at 200ns (configurable + reset) ← Leading dead-time per channel
- 8bit ADC conversion at 100ns (configurable + reset)



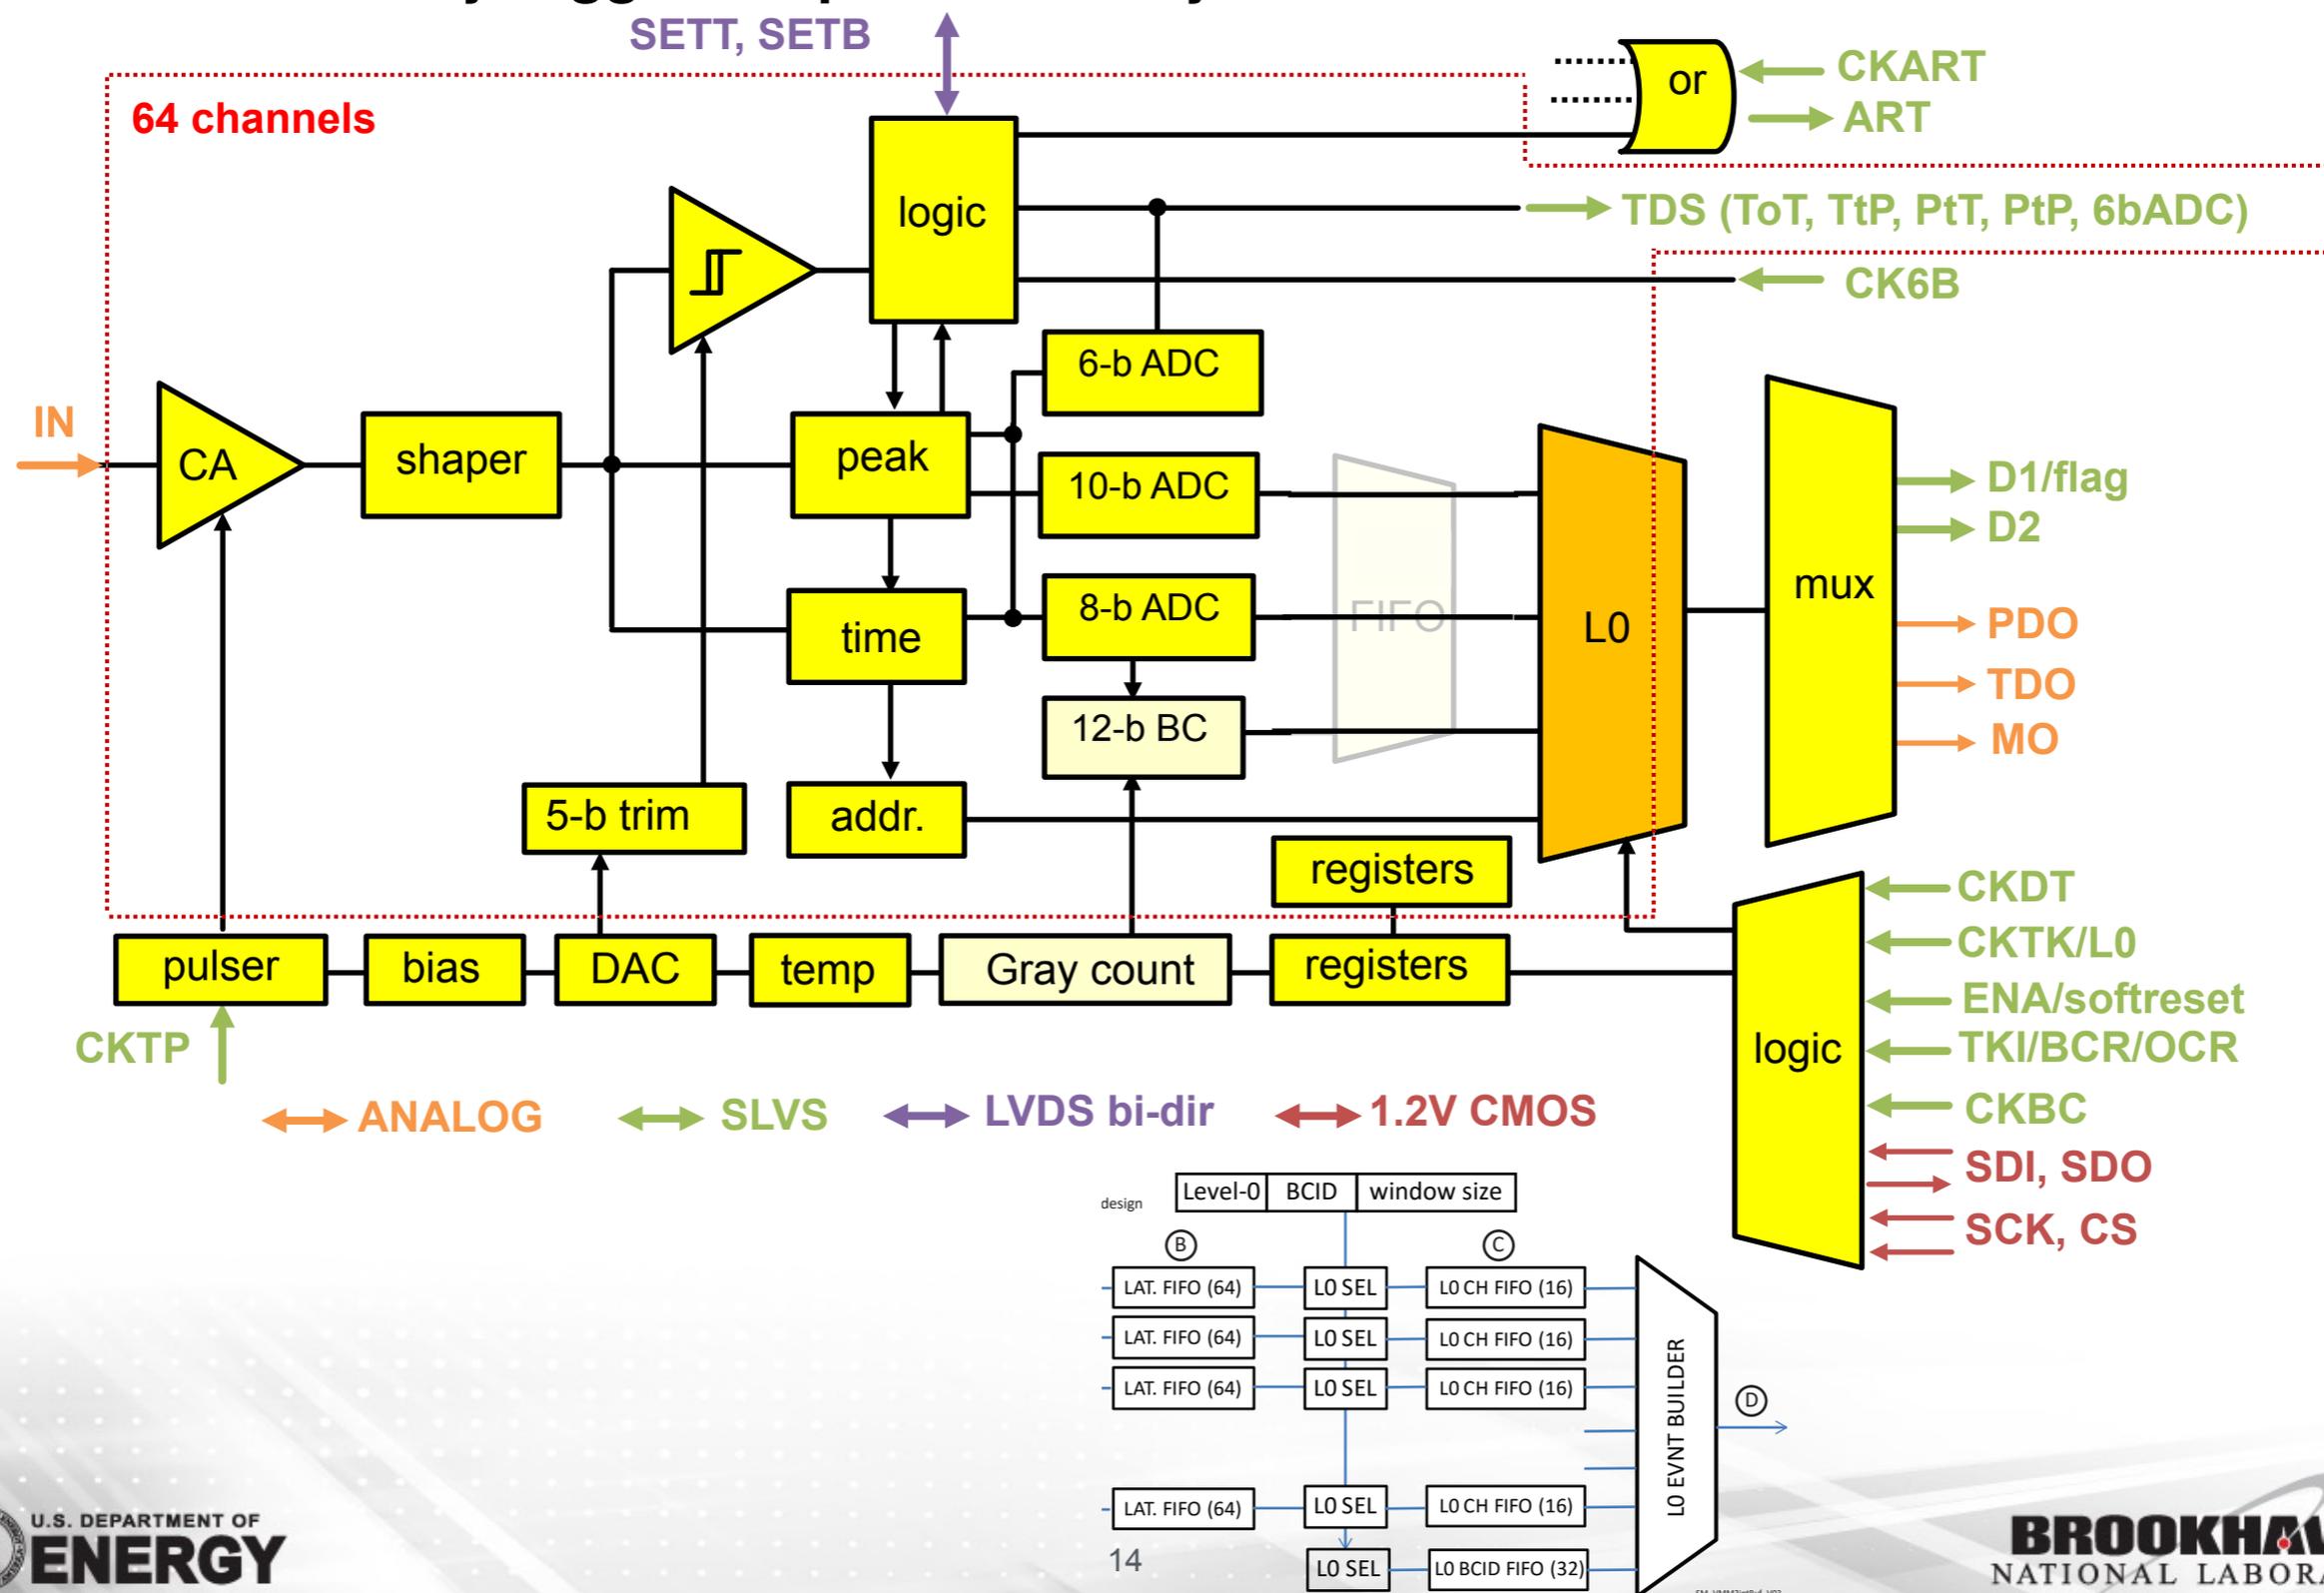
Modes of operation - Direct output

- This mode provides **continuous trigger-less readout**
- All the outputs and inputs are active and **independent**
- 6bit ADC conversion within 25ns (configurable) ← Enabling sfrst, channel resets after 6bit conversion ~ 50ns dead-time
- 10bit ADC conversion at 200ns (configurable) ← ~~Leading dead-time per channel~~
- 8bit ADC conversion at 100ns (configurable)



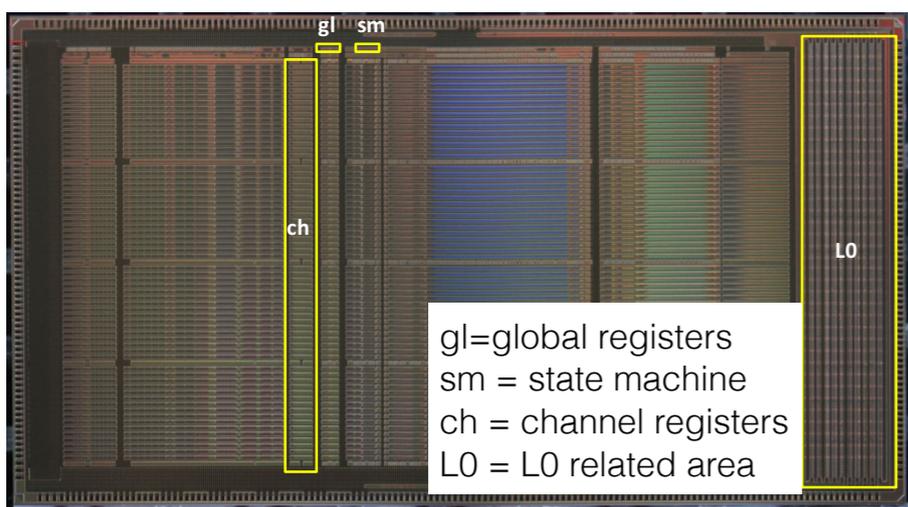
Modes of operation - L0

- The **signal processing** is done in the **same way** but the **readout is different**.
- This is an **externally triggered operation** for **synchronous machines**



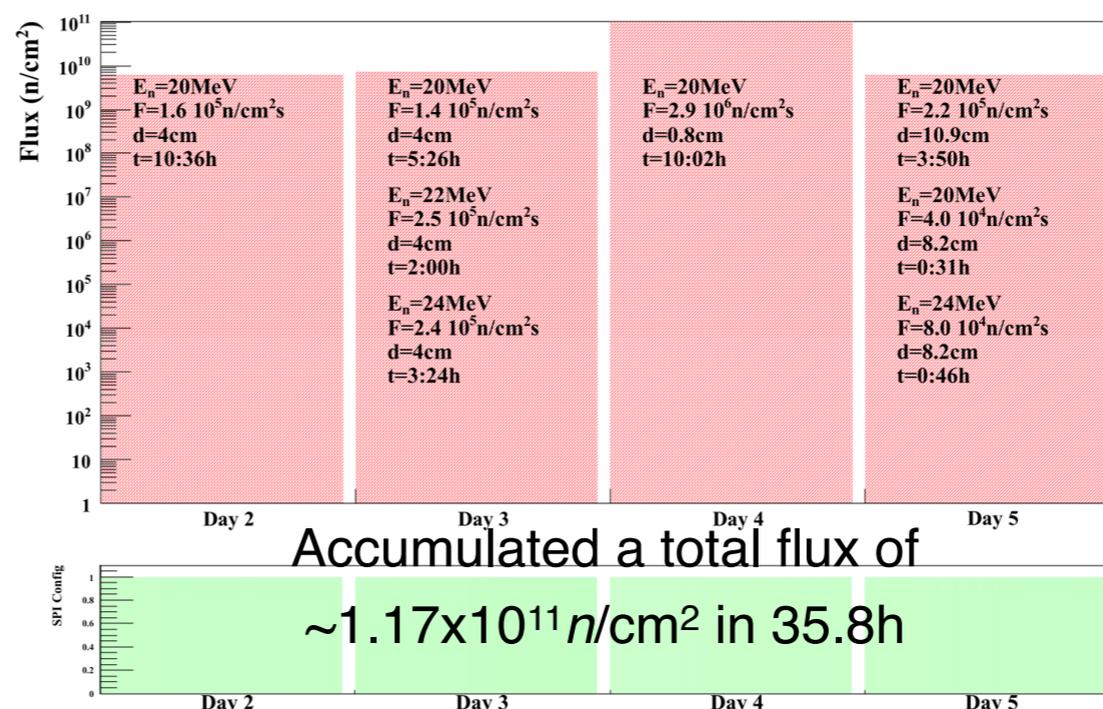
Single Event Upset & Total Ionisation Dose

- In the VMM3a there are **three types of storage elements** that require SEU protection, the **configuration registers**, the **state machine** control logic and the **L0 logic**
- To mitigate for SEU two techniques are used:
 - **Dual Interlocked Cells (DICE)** for the protection of the configuration registers
 - **Triple Modular Redundancy (TMR)** for the state machines and the L0 Logic blocks
- L0 Data
 - Single-bit faults on data are flagged by a **parity bit**
 - The **parity** is registered in the FIFOs and **transmitted** outside

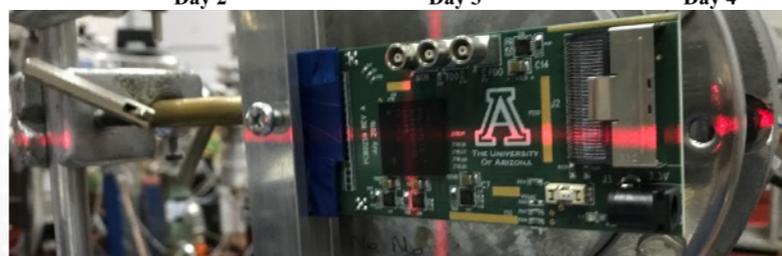


Block	Method
BC counter	TMR
Latency FIFO CTRL	Parity on FIFO pointer, FIFO resets if parity err
L0 FIFOs Control	TMR
Event Builder	TMR
L0A register/Nskip circuit	TMR

L0 block protection

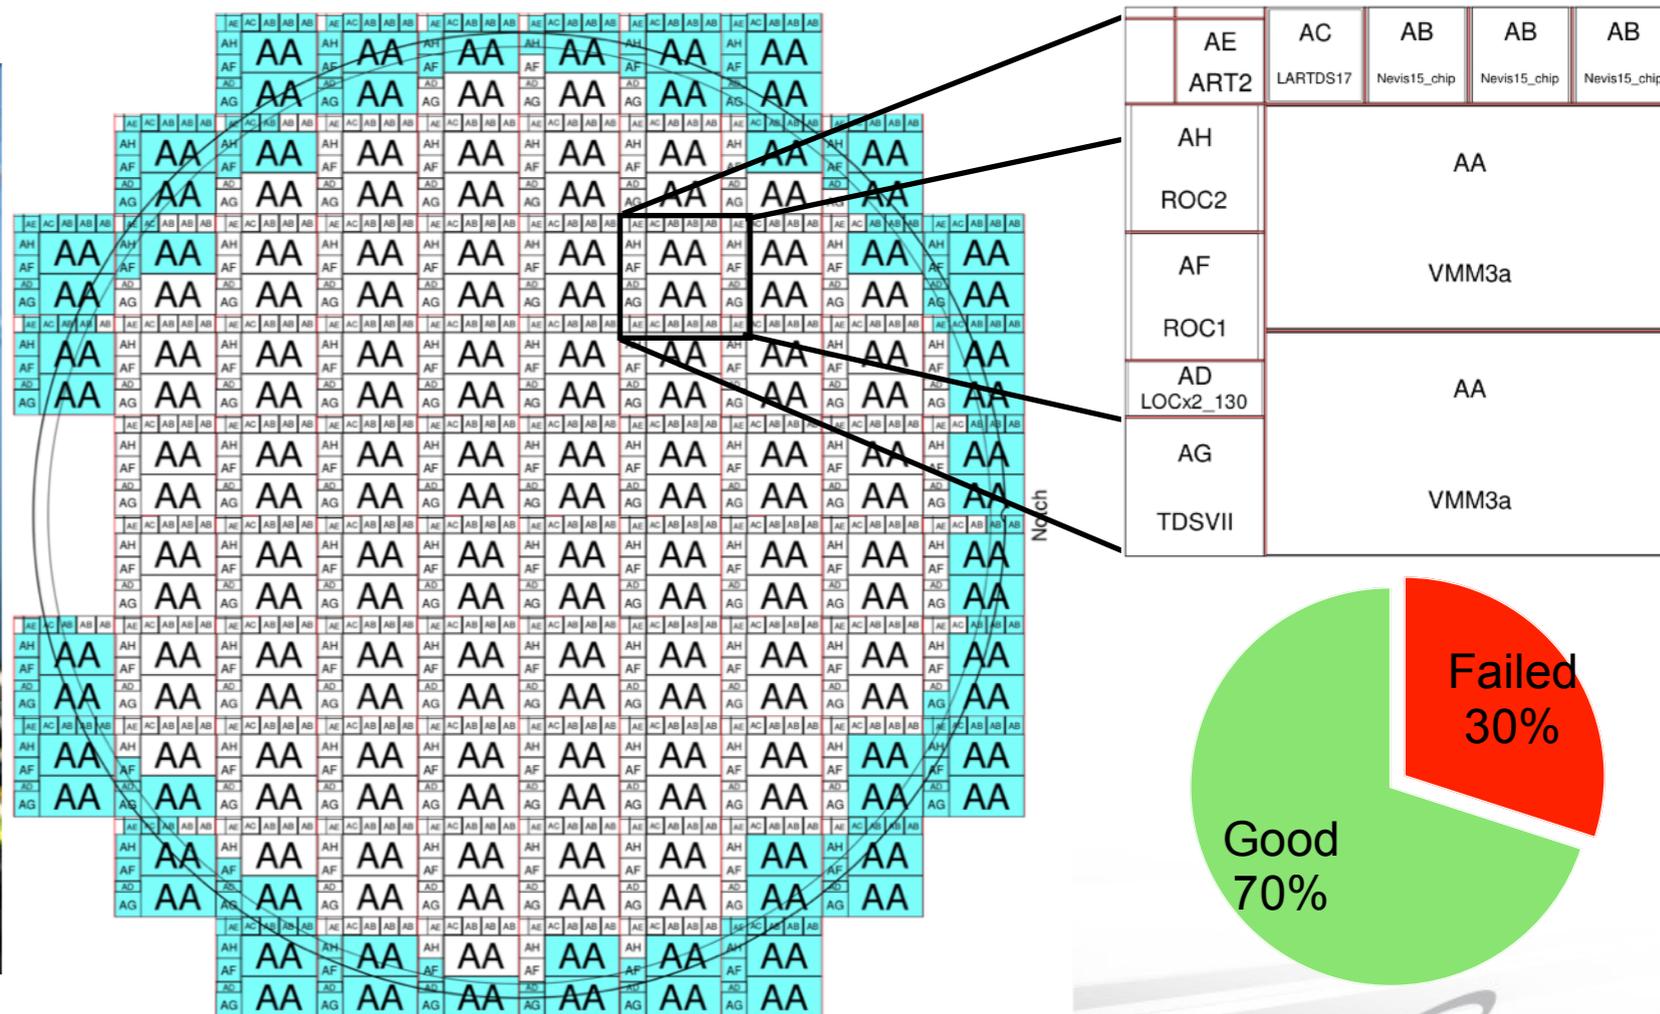
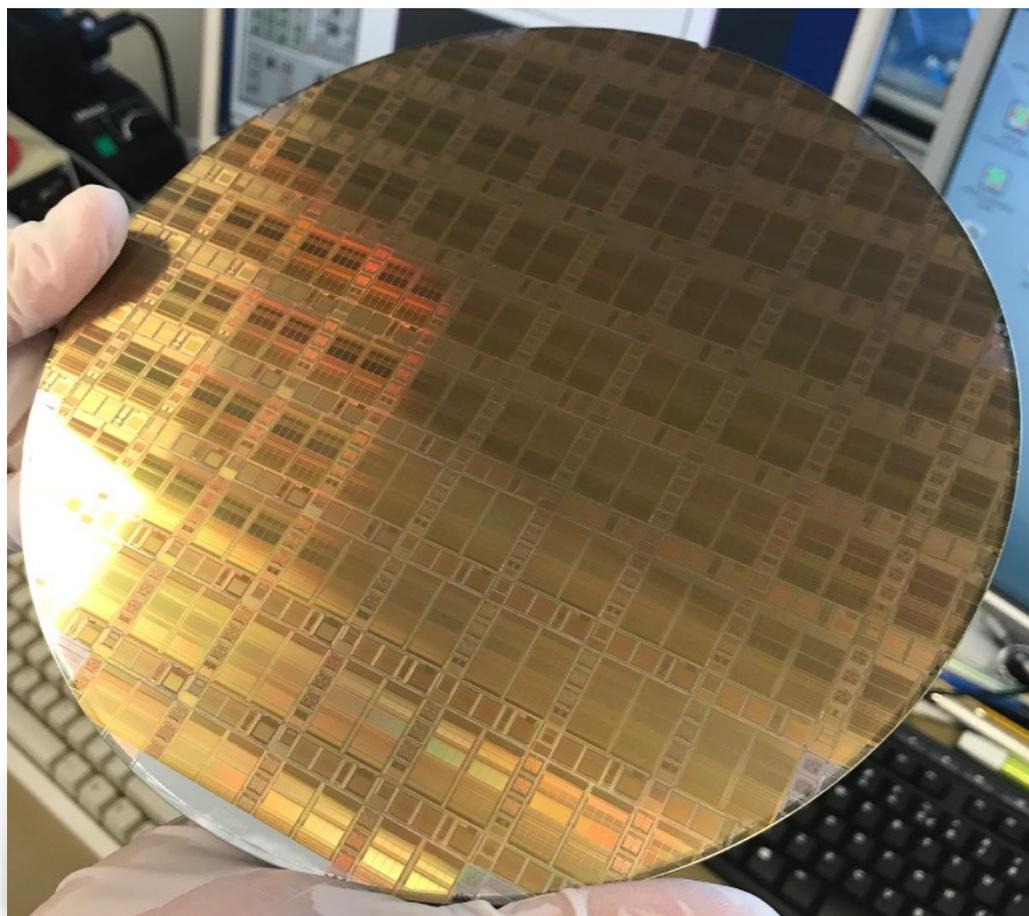


4 VMM3a were irradiated at the ^{60}Co



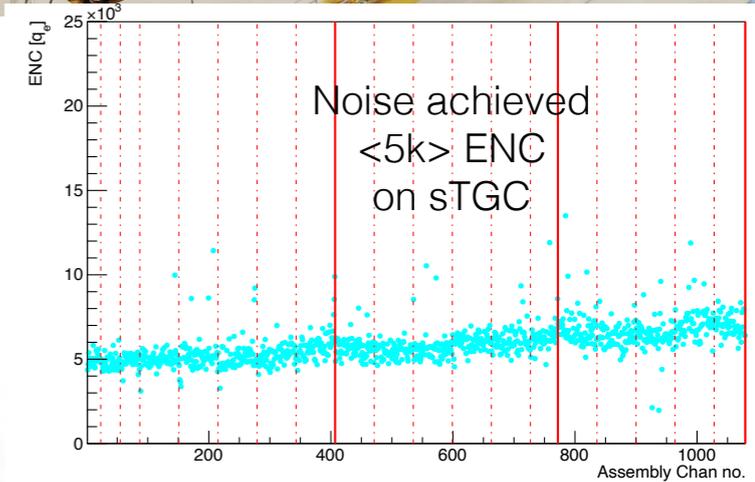
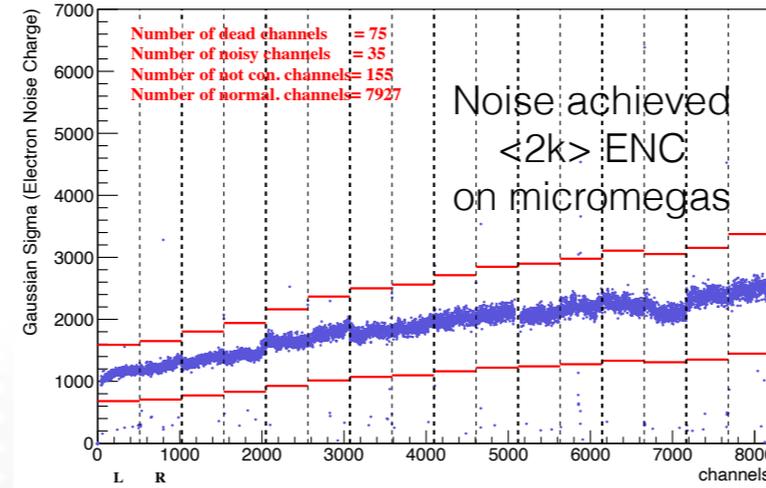
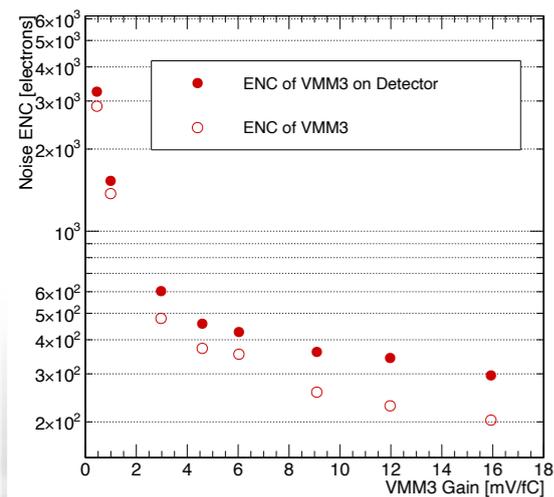
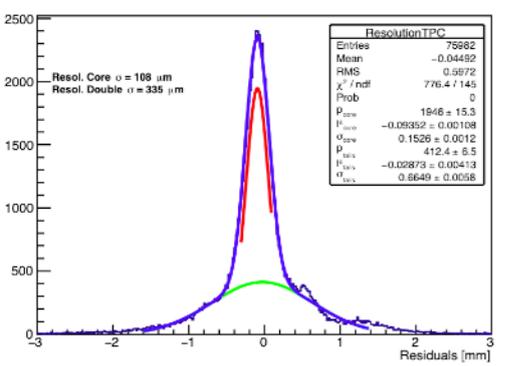
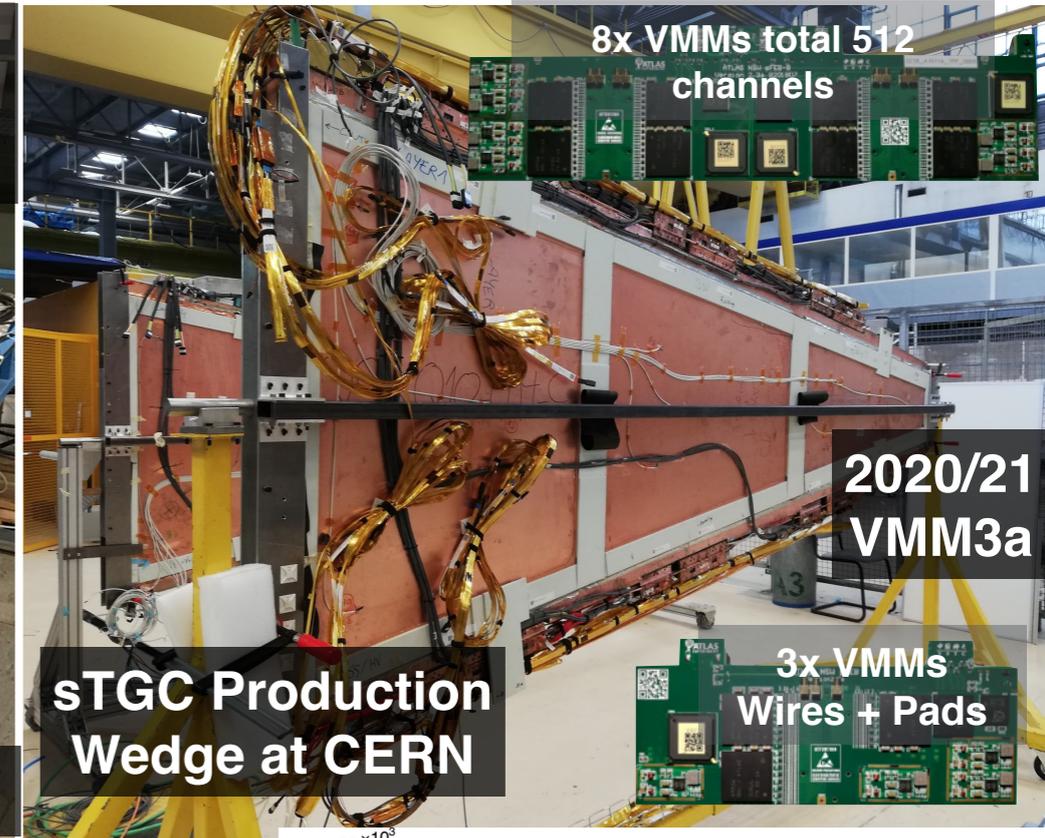
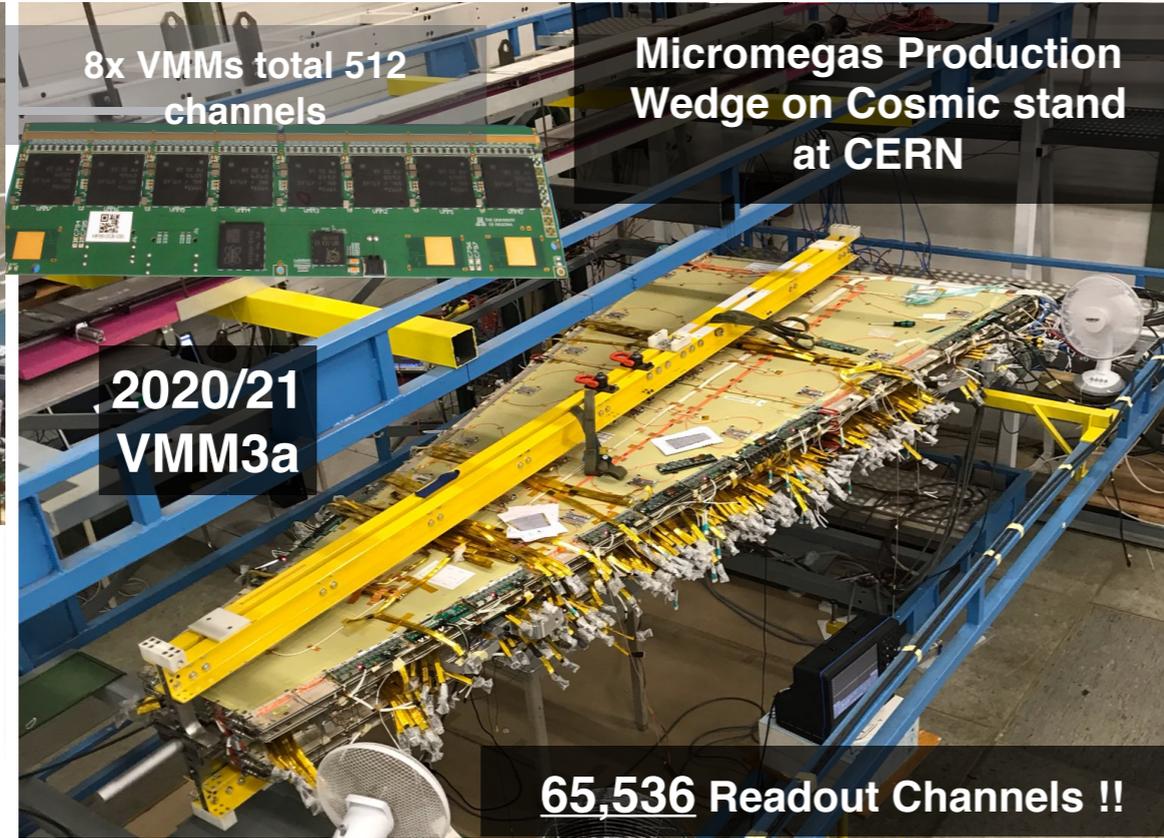
VMM Production

- The **VMM** is produced in a 8" wafer with 2 copies of the chip in a reticle, **total 113 chips / wafer**
- During the production we faced several issues due to GF processing affecting the yield
- Many iterations with experts from Global Foundries to improve the yield and understand the issue
- Investigation concluded (HPT process maintained throughout the production for high density metal layers)
- In ATLAS we had no time to further investigate the issue and moved forward to production
- ATLAS has already produced and package **73k** VMMs (incl. prototyping)
- Many other experiments and applications submitted proposals to use it (already in use by others)



Integration with Micromegas & sTGC Production modules

- **Micromegas & sTGC Production modules at CERN integrating the VMMs**
- All the FE boards are readout through custom made 4.8Gbit serialiser boards with fibres
- Water Cooling is a must in these applications



Conclusions

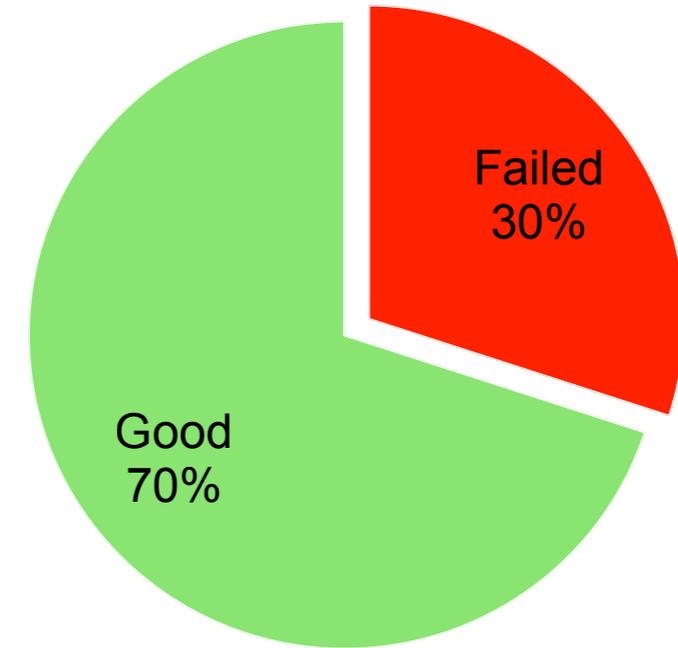
- VMM project was recently concluded successfully for ATLAS New Small Wheels
 - VMM is used for other R&D applications and proposed for several other experiments
- It has been a ~10y project from R&D to production and integration
- It is a state of the art ASIC for physics applications low power, low noise, fully adjustable providing many features and readout techniques
- We have faced many issues along the road:
 - Logic issues on the circuit (reset conflicts, logic conflicts, readout issues with token passing technique)
 - ADCs have been improved significantly but it took two iterations to achieve it
 - Baseline stabilisation issues (design issues and manufacturing problems, still the majority of failures)
- Eventually most of the problems have been addressed and VMMs are already integrated in high density low noise front-end boards on detectors showing the expected performance

backup

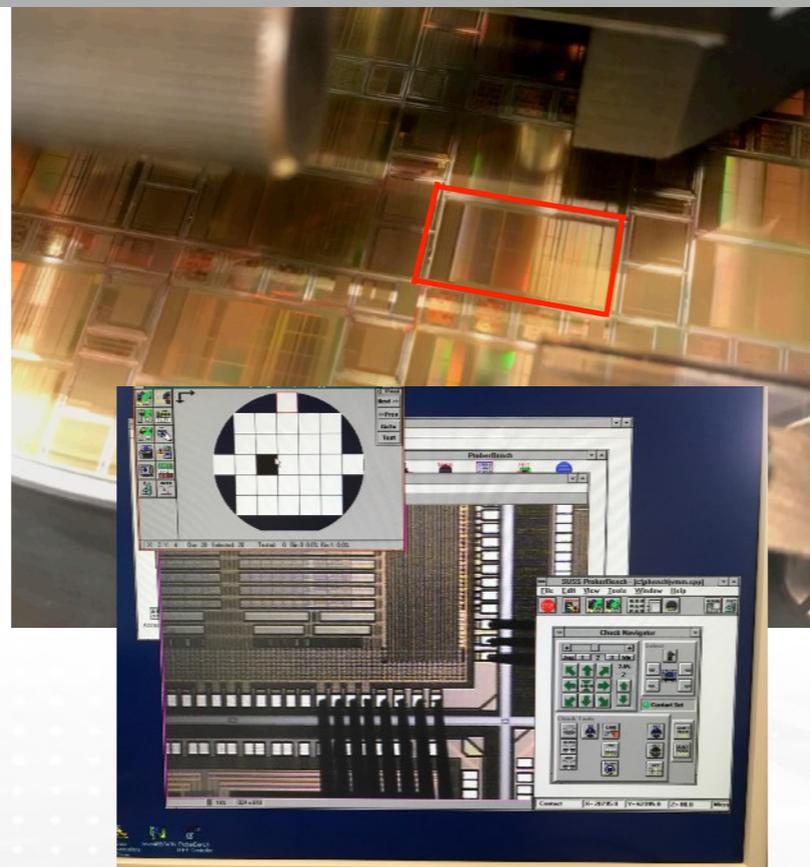
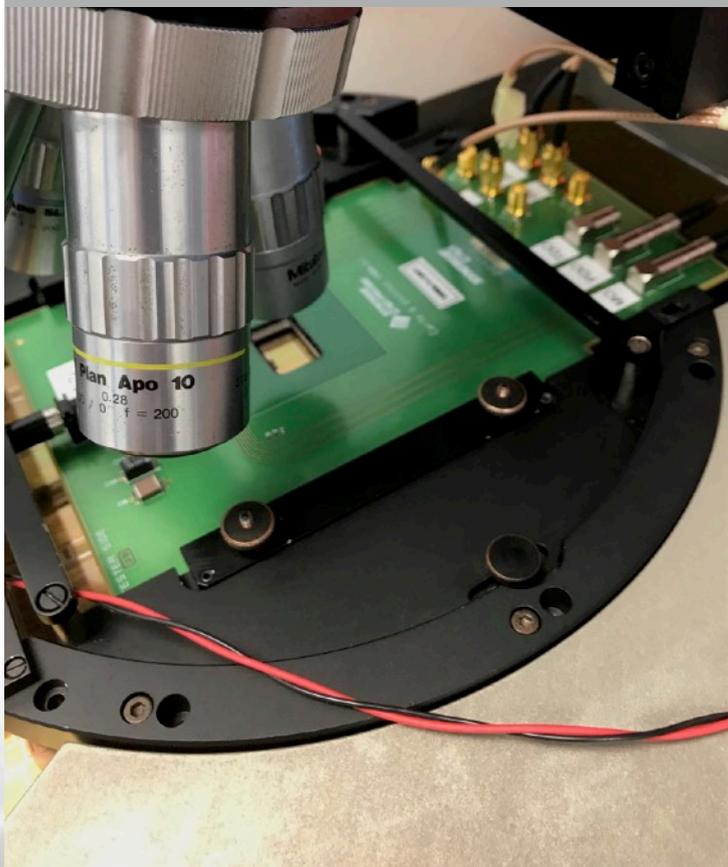
VMM Production testing

- Many testing protocols have been produced for testing the devices
- Direct wafer probing was developed to allow initial screening of the production batch (no detailed testing, only the frontend)
- Half of the production was tested by manual operators (lengthy process)
- Due to constraints in time, we developed automated testing (30sec/device) which accelerated dramatically the process
- Throughout the process, 70% yield was achieved, 30% mainly due to ESD damage on input transistor or baseline stabilising circuit

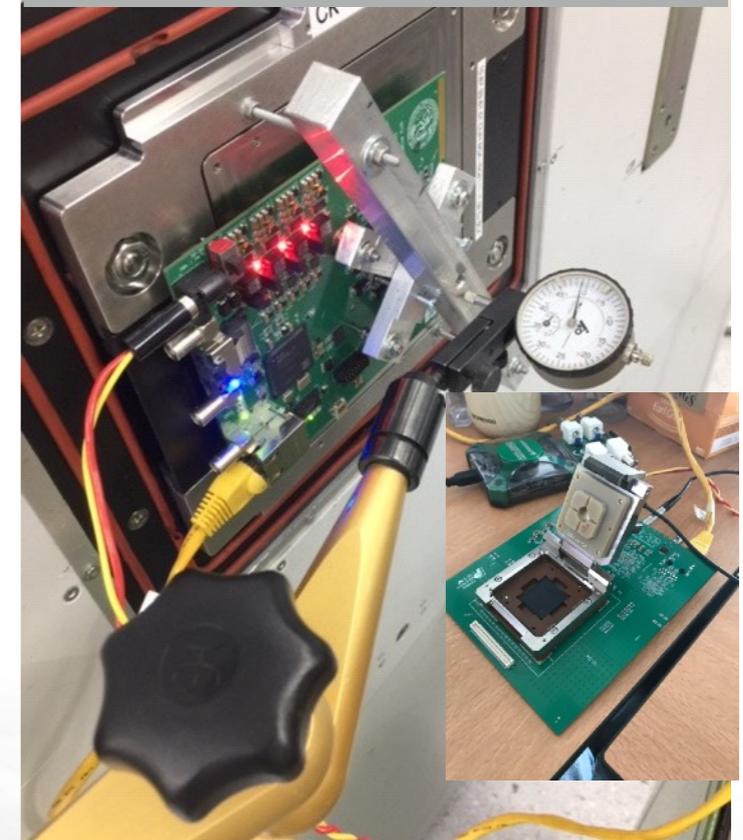
VMM Yield



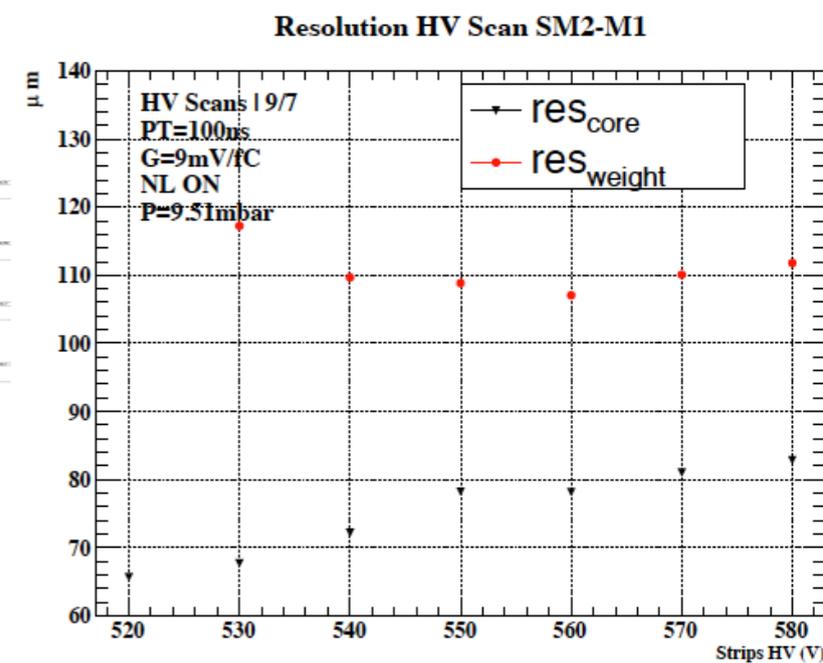
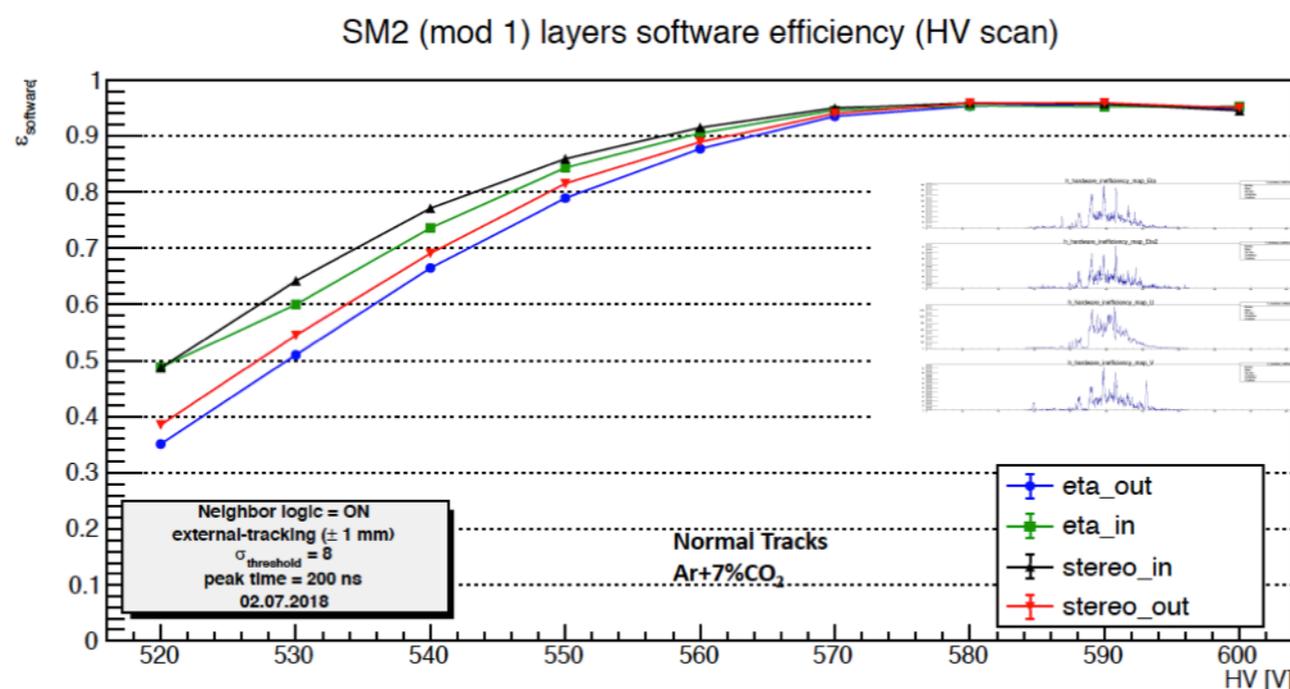
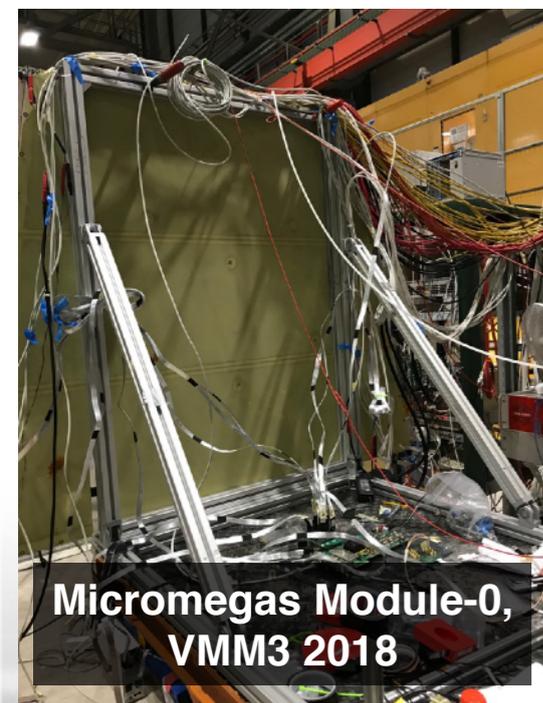
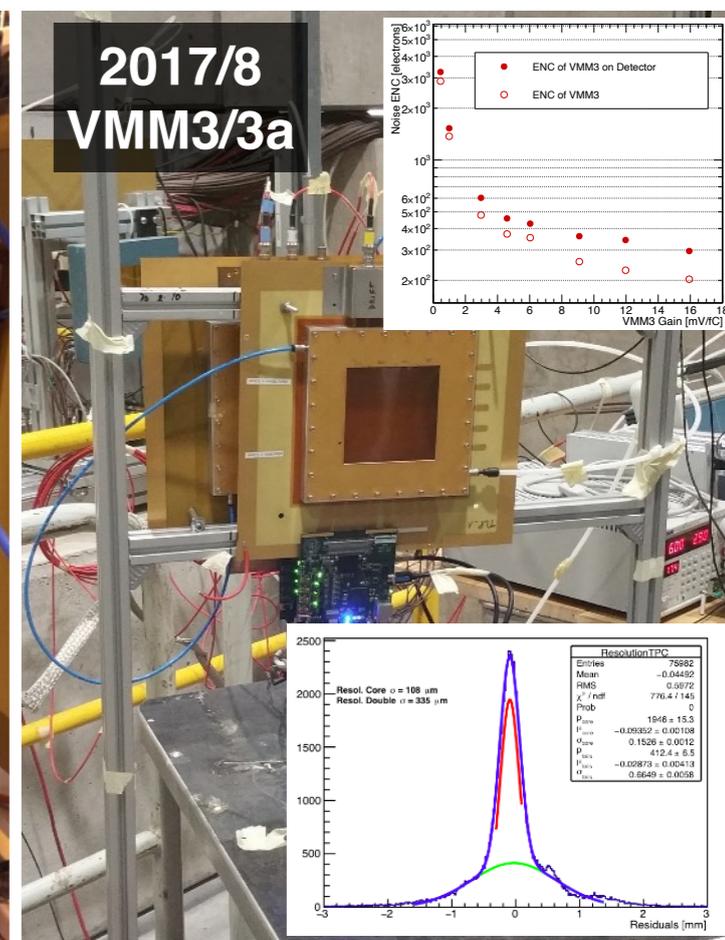
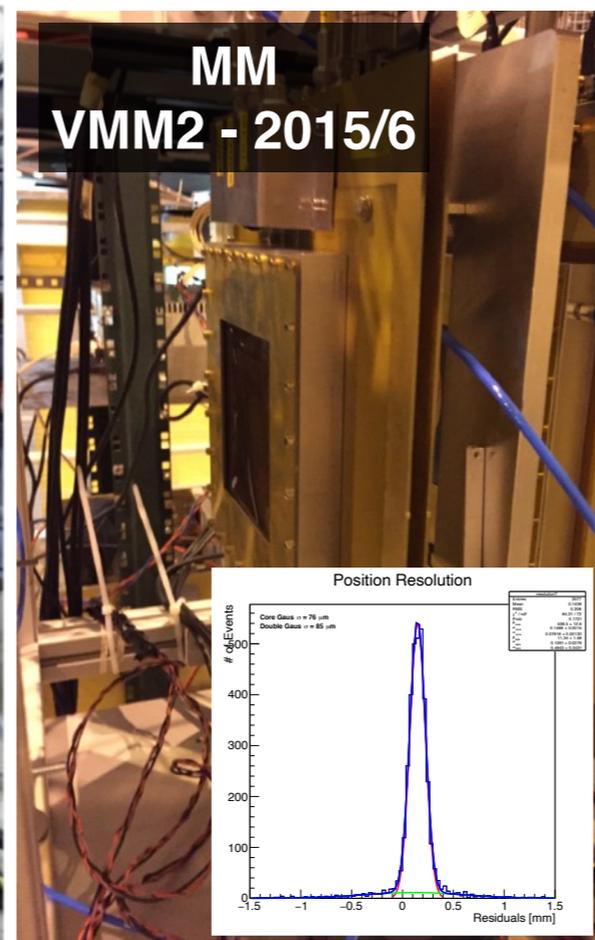
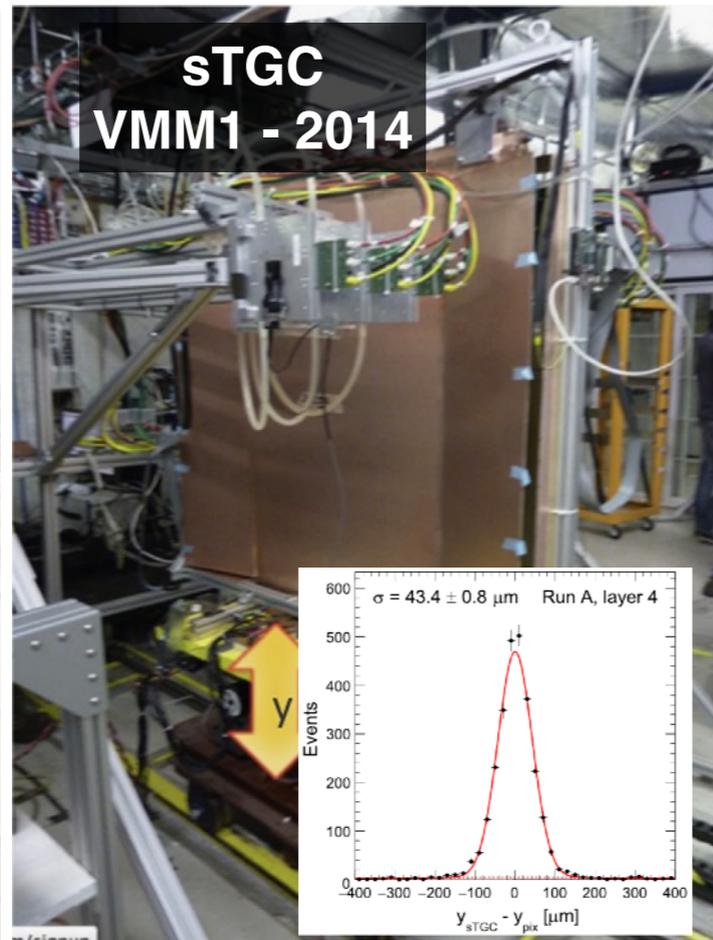
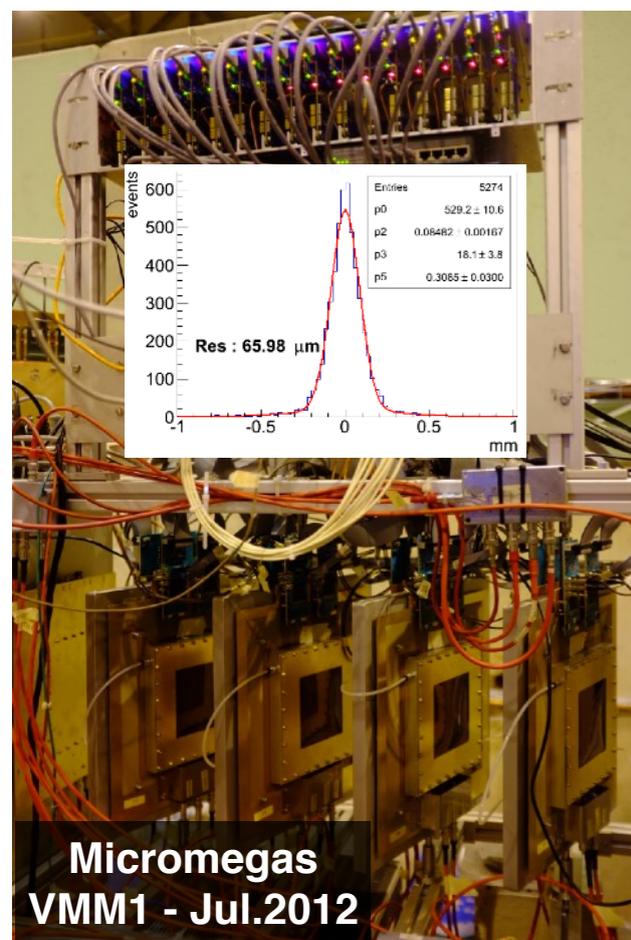
Wafer probing



Testing in industry

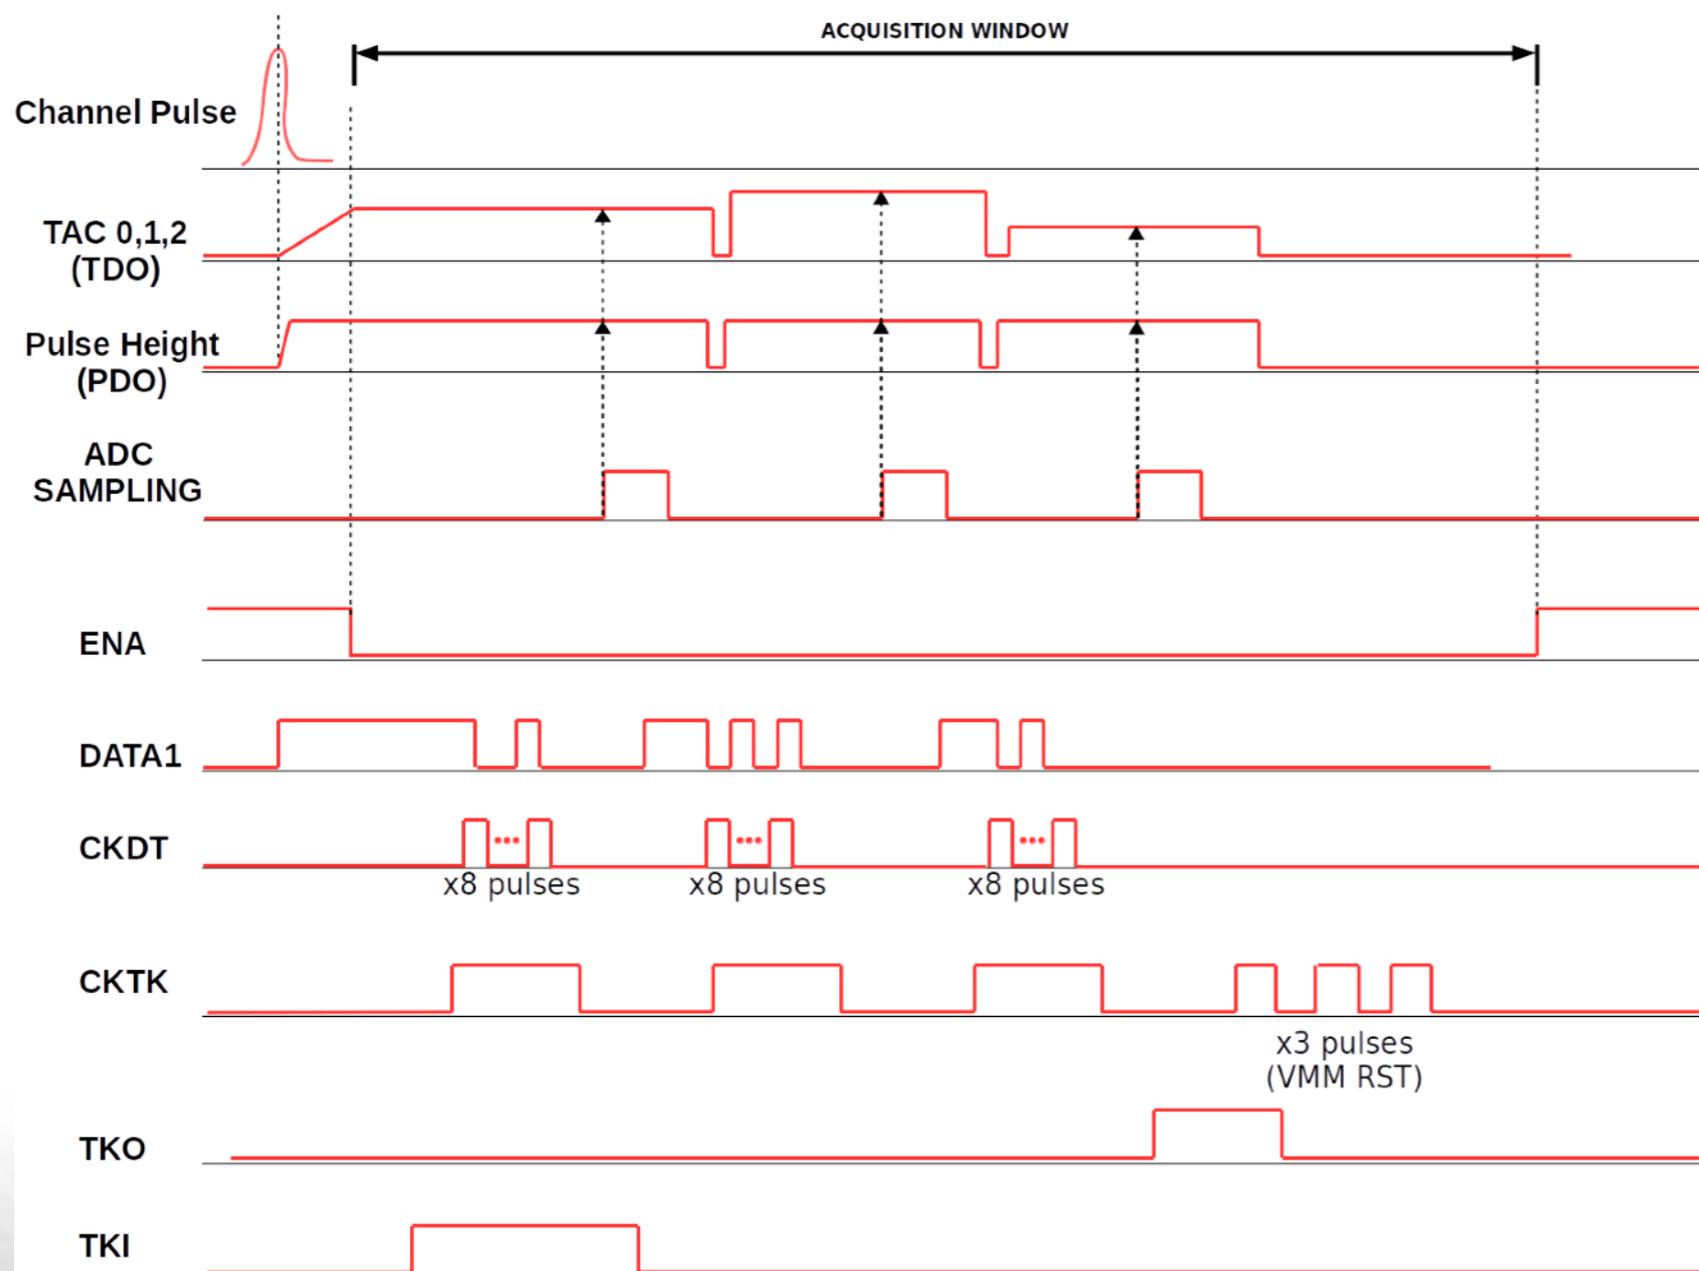


Test Beams from 2012 to 2018



Modes of operation - Analog

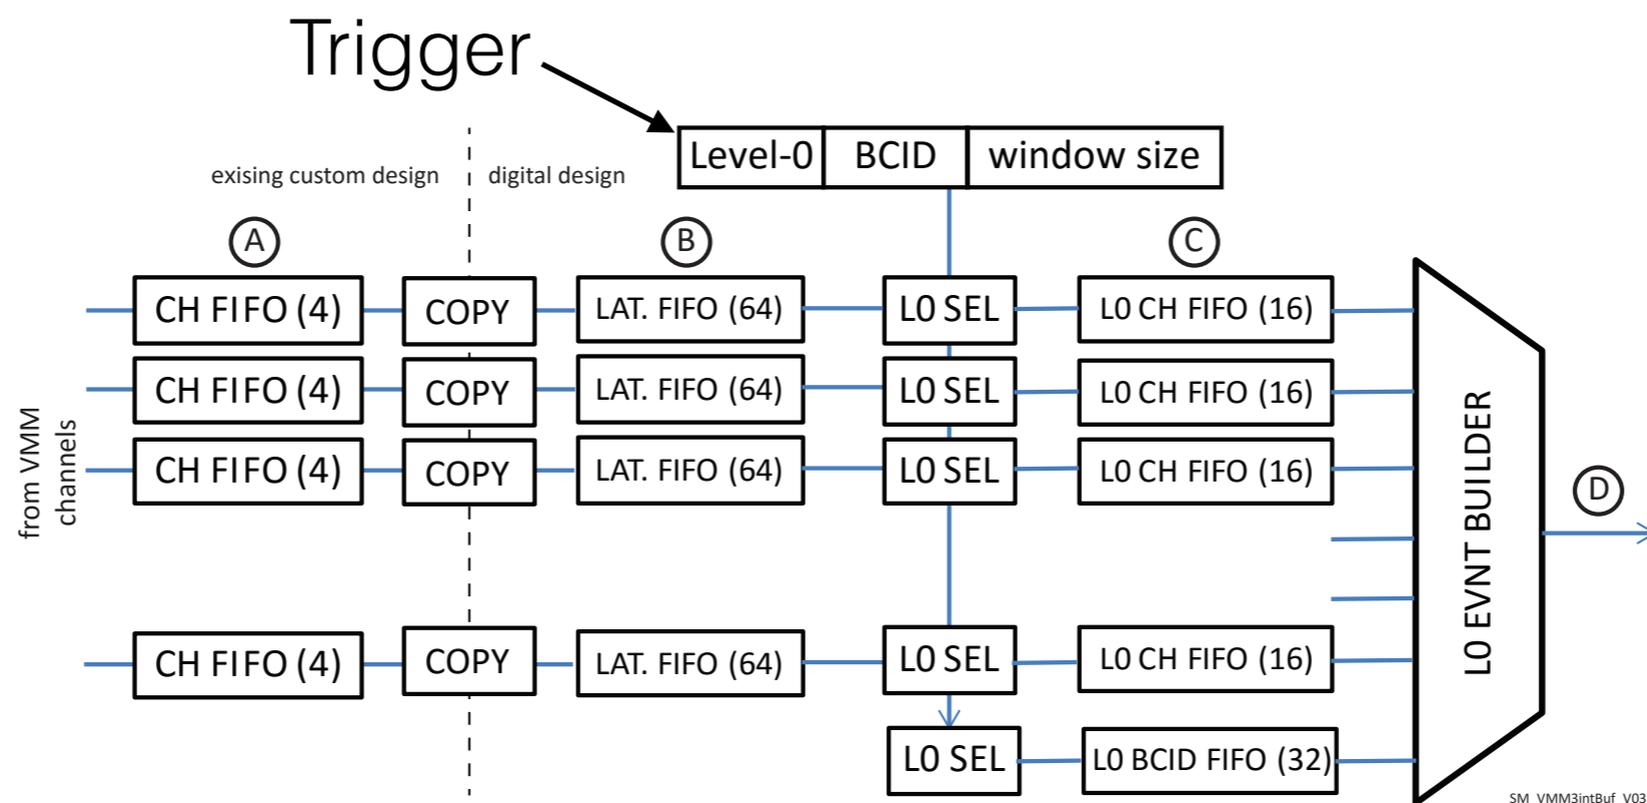
- In two-phase (analog) mode which is the mode originally implemented in the VMM1, the ASIC operates in two separate phases: **acquisition** and **readout** - During the acquisition phase the events are processed and stored in the **analog memories** of the **peak and time detectors**. As soon as a first event is processed, a flag is raised at the digital output.



- Once the process is complete the **ASIC can be switched readout phase**. The first set of amplitude and time voltages is made available at the **analog outputs**. The **address of the channel is serialised** and made available at the digital output using **six data clocks**.

Modes of operation - ATLAS L0 readout

- Each channel has a **Level-0 Selector** circuit which is **connected to the output** of the channel's latency FIFO.
- The **selector finds events within the BCID window** (maximum size of 8 BC clocks) of a Level-0 Accept and **copies them to the L0 Ch FIFO**. The data are available in the output which is running on **IDLE K28.5** in two data lines and can be readout **DDR at a speed of 640Mbps** (160MHz clock tested, effective bandwidth 560Mbps due to 8b/10b encoding).



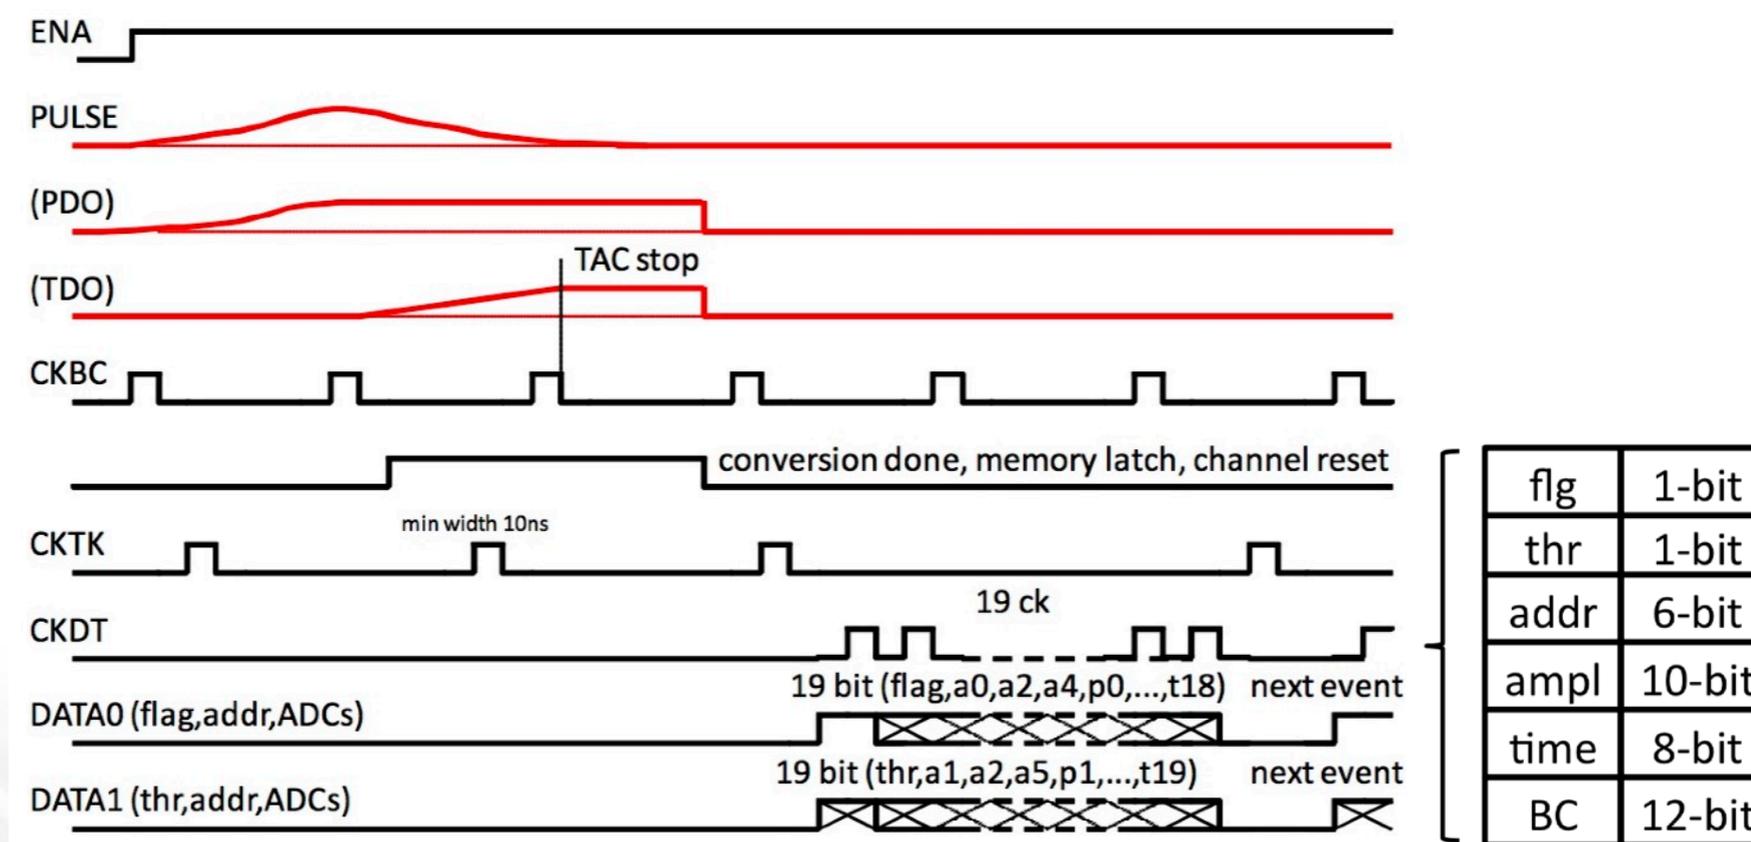
- VMM will build the event with common BCID +relative for each hit
- Header is sent out once no data found

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
header	V	P	orb		BCID (12)												1st word after comma															
hit data	1	P	R	T	Chan# (6)						ADC (10)						TDC (8)				N rel BCID											

LL_format_VMM3out_V04

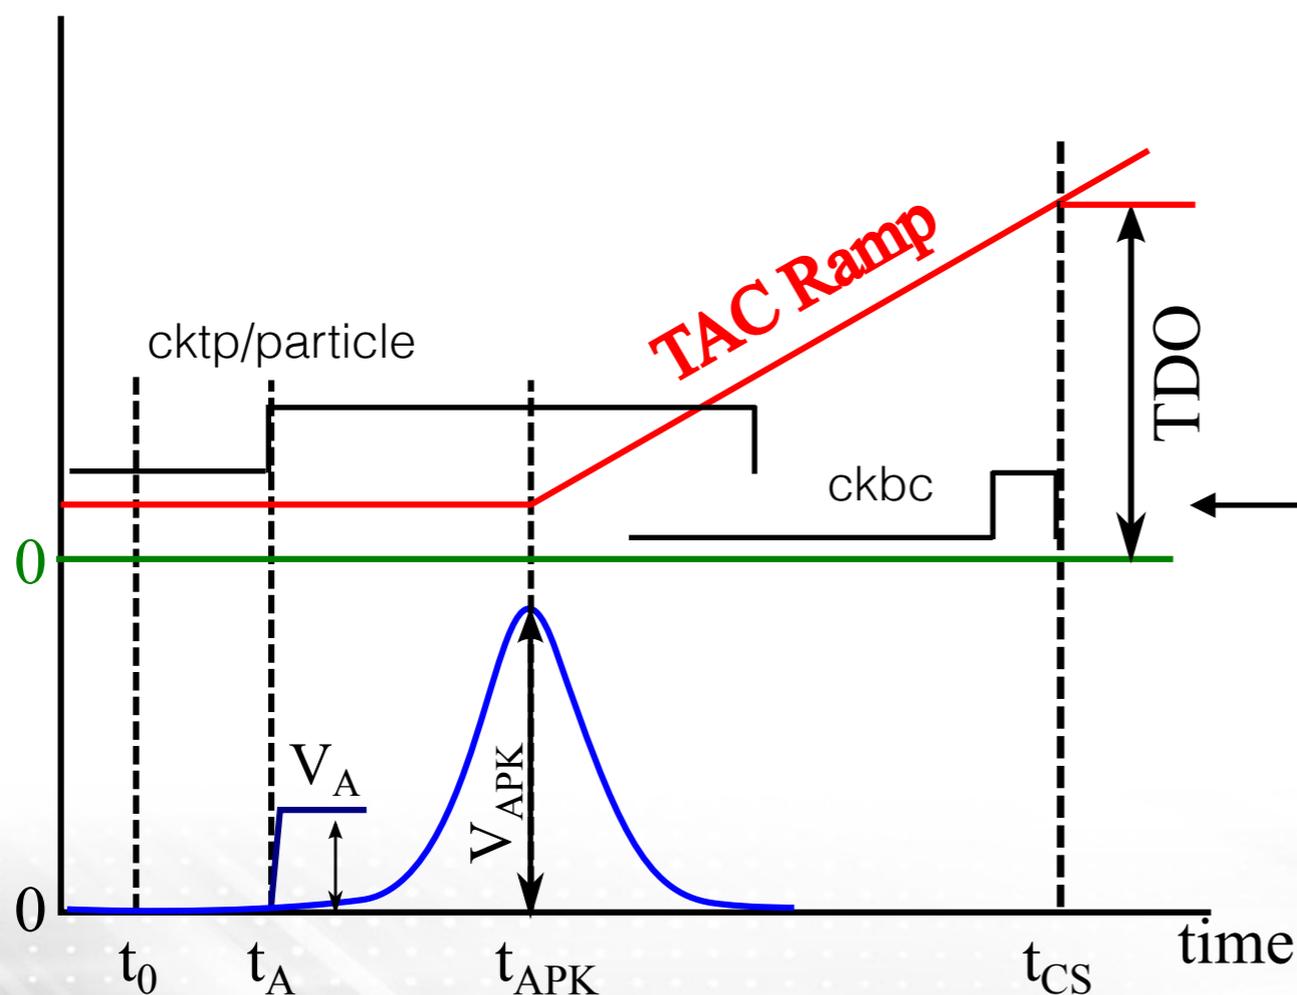
Modes of operation - Continuous

- Trigger-less mode with **peak and time detectors convert the voltages into currents** routed to the **6/10-bit ADC and 8-bit ADC** respectively.
- The 10-bit ADC provides a **high resolution A/D conversion** of the peak amplitude
- The 8-bit ADC provides the A/D conversion of the **timing (measured using the TAC peak or the threshold to a stop signal). Time associated with global 12-bit counter.**
- In the continuous mode the **64 channel direct outputs** are available as well providing time-over-threshold (ToT), threshold-to-peak (TtP), peak-to-threshold (PtT), or a 10 ns pulse occurring at peak (PtP) or the 6-bit ADC. The **channel self resets** at the end of the timing pulse, thus providing continuous and **independent operation of all 64-channels.**



Modes of operation - Continuous + ext trigger

- VMM design targets **synchronous** machines hence can be difficult to use in an environment like a **test beam where asynchronous operation** is needed but **precise timing** is needed to be measured (drift time)
- Most chips designed for synchronous machine **suffer from time jitter in such environment**
- On **VMM a mode was foreseen** to do such measurement where the **ckbc** can be used as a **strobe** and **not like a real clock**
- It can be **send as a trigger signal** with a **fixed latency** achieving precise time measurements



- Trigger signal from **external source**
- Can be combined with **register stcr** where **channel resets** if stop signal not occurs within the **TAC ramp**
- **Implies that trigger is propagated** within the TAC ramp up time (**60ns-650ns**)
- The **longer the TAC** though the **lower the resolution** on 8-bit information from the ADC
- Highly correlated trigger readout and noise subtraction