

MOSAIX as a Case Study: Overview of the Stitched Chip for ITS3

Gianluca Aglieri Rinella, CERN EP-ESE

Acknowledgements: results and material from the work of a very large group of persons of the ALICE ITS3 project, EP R&D WP1.2 and beyond

Outline




Introduction and brief overview of ALICE ITS3 Upgrade

Technical overview of the MOSAIX sensor design


Selection of lessons learnt

ALICE ITS3 Upgrade


ALICE ITS2 (TAKING DATA)



~12.5 Gpixels, 10 m² sensitive area
24120 **ALPIDE** Monolithic Pixel Sensors (CMOS 180 nm)




ITS2 Inner Layers



ITS3



ALICE ITS3 Upgrade



Replace the ITS2 inner barrels by real half-cylinders (of **bent, thin** silicon)




Employ wafer-scale MAPS sensors (1 sensor per half-layer) in **65 nm** technology

Minimize material budget and distance to interaction point ($0.07\% X_0$ / layer, 19 mm)

Large improvement of vertexing precision and physics yield

ALICE ITS3 Upgrade

ITS3 Engineering Model 2
Wafer size blank silicon dies, 40 μm thin, 280 mm long



ALICE ITS3 TDR Requirements

3 Cylindrical Layers

Made with **6 curved wafer-scale single-die**
Monolithic Active Pixel Sensors

Thinned down to **<50 μm (0.07% X_0)**

Position resolution $\sim 5 \mu\text{m}$

-> Pixel pitch **20-25 μm**


Electro-mechanical integration

No flexible circuits in the active area

-> Distribute supply and transfer data *on chip*,
interconnects only on short edges

Cooling by air flow

-> Dissipate less than **40 mW/cm²** (in sensitive
area)




	Requirement
Pb-Pb Interaction Rate	50 kHz
Particle Flux	5.75 MHz/cm ²
Integration time	< 10 μs
TID	<10 kGy
NIEL	1×10^{13} 1 MeV n_{eq} cm ⁻²


ALICE ITS3 TDR <https://cds.cern.ch/record/2890181>

Wafer Scale Sensors with Stitching

Design Reticle (typ. 2x3 cm)



Circuits on wafer



ER1 Submission

Learning design with **stitching** and handling procedures

Submitted in December 2022, 65 nm CMOS Imaging Technology


Two wafer scale stitched sensor chips

MOSS: 14 mm × 259 mm, 6.72 Mpixels, (22.5 × 22.5 and 18 × 18 μm²), conservative design, different layout densities


MOST: 2.5 mm × 259 mm, 0.9 Mpixels (18 × 18 μm²), full density design

Single stitch devices

Several small test chips (1.5 mm × 1.5 mm)



MOSS Tests



MOSS design **fully functional**

Design concepts and methodology **validated**


Much learning on yield, handling and performance of a full-scale device

Paved the way for the design of the engineering prototype MOSAIX

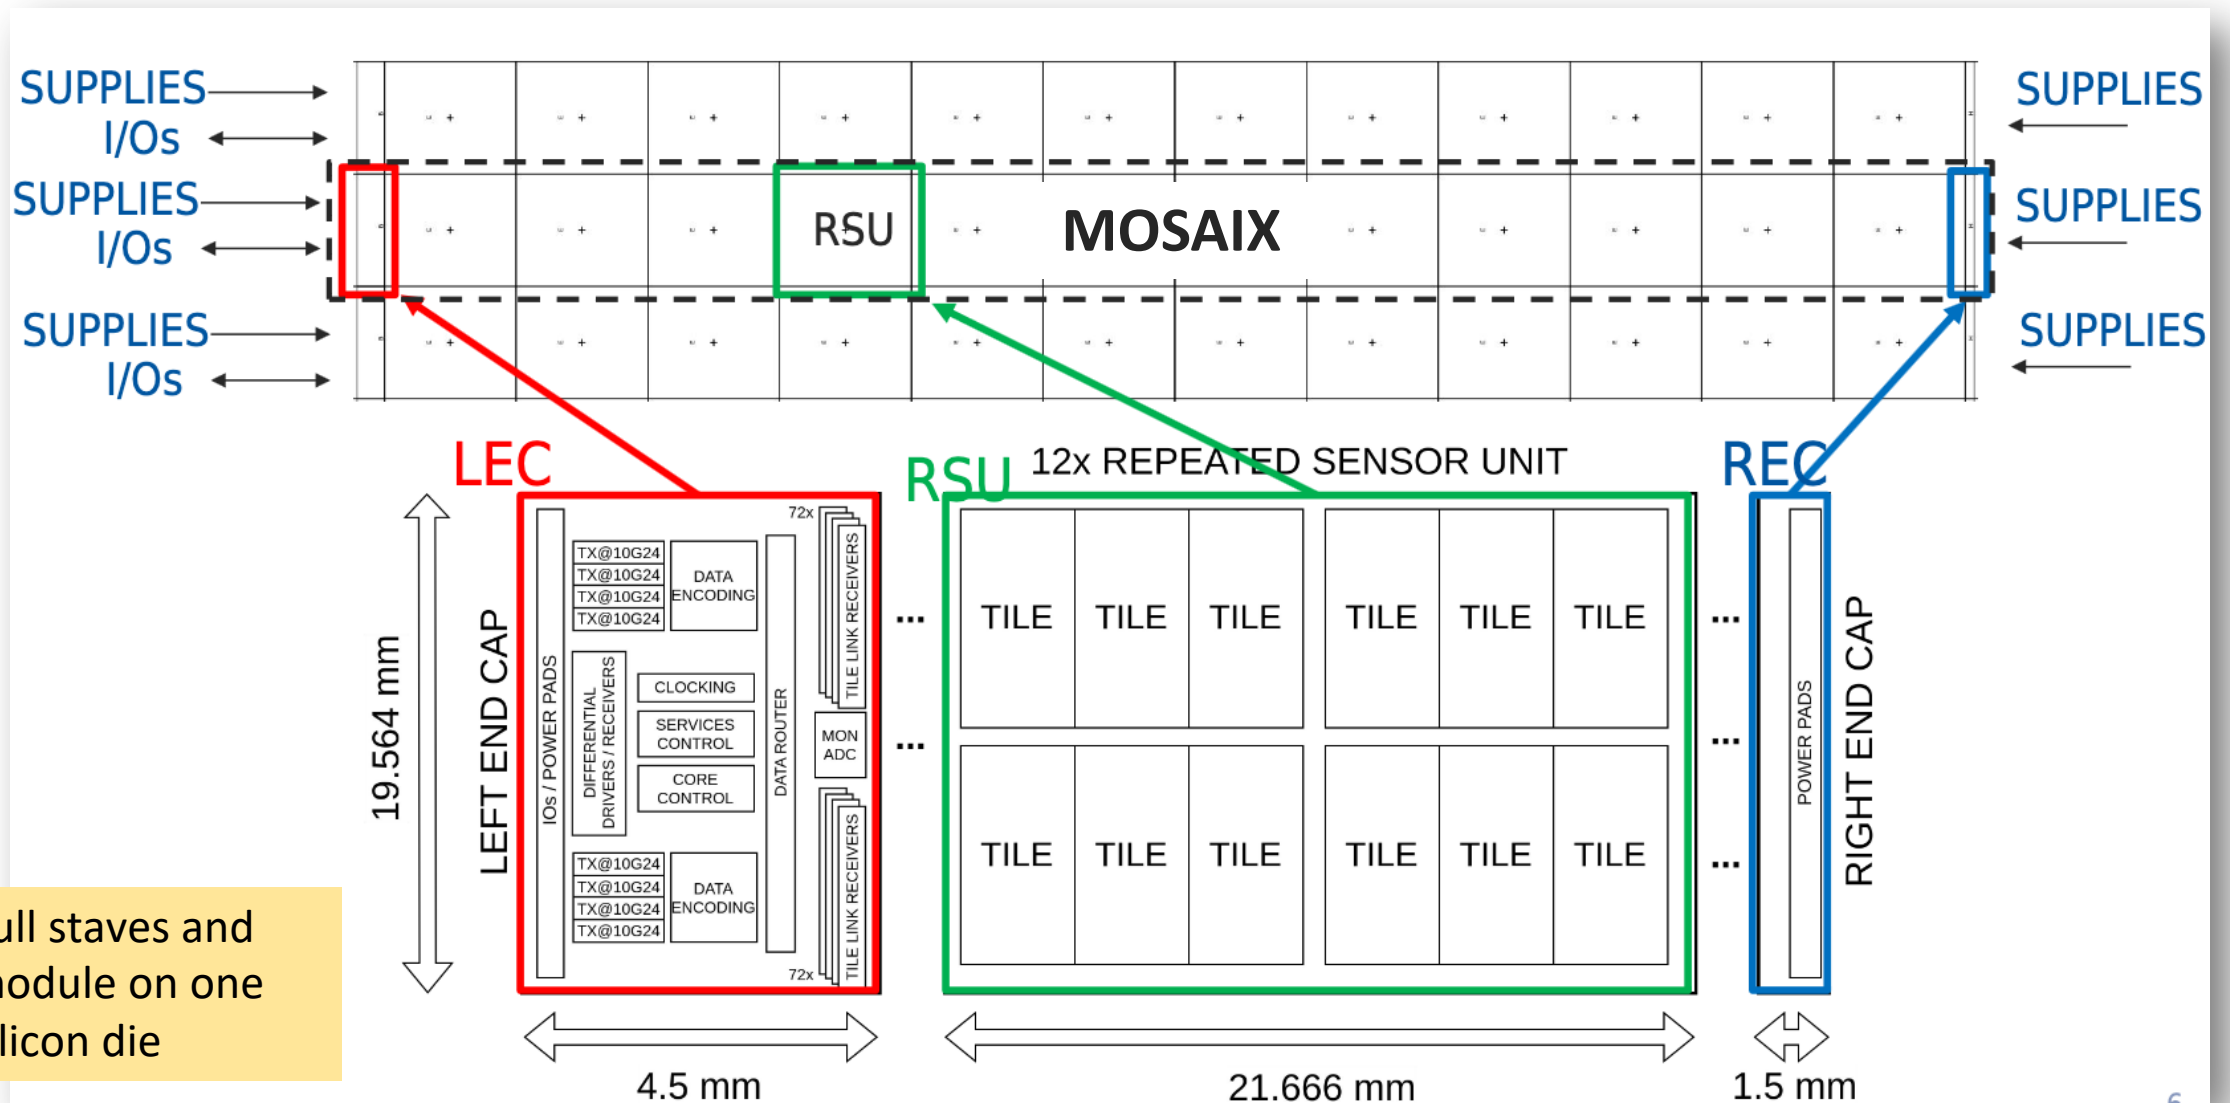
ALICE ITS3 and MOSAIX



MOSAIX is a full feature prototype of the sensor for ALICE ITS3

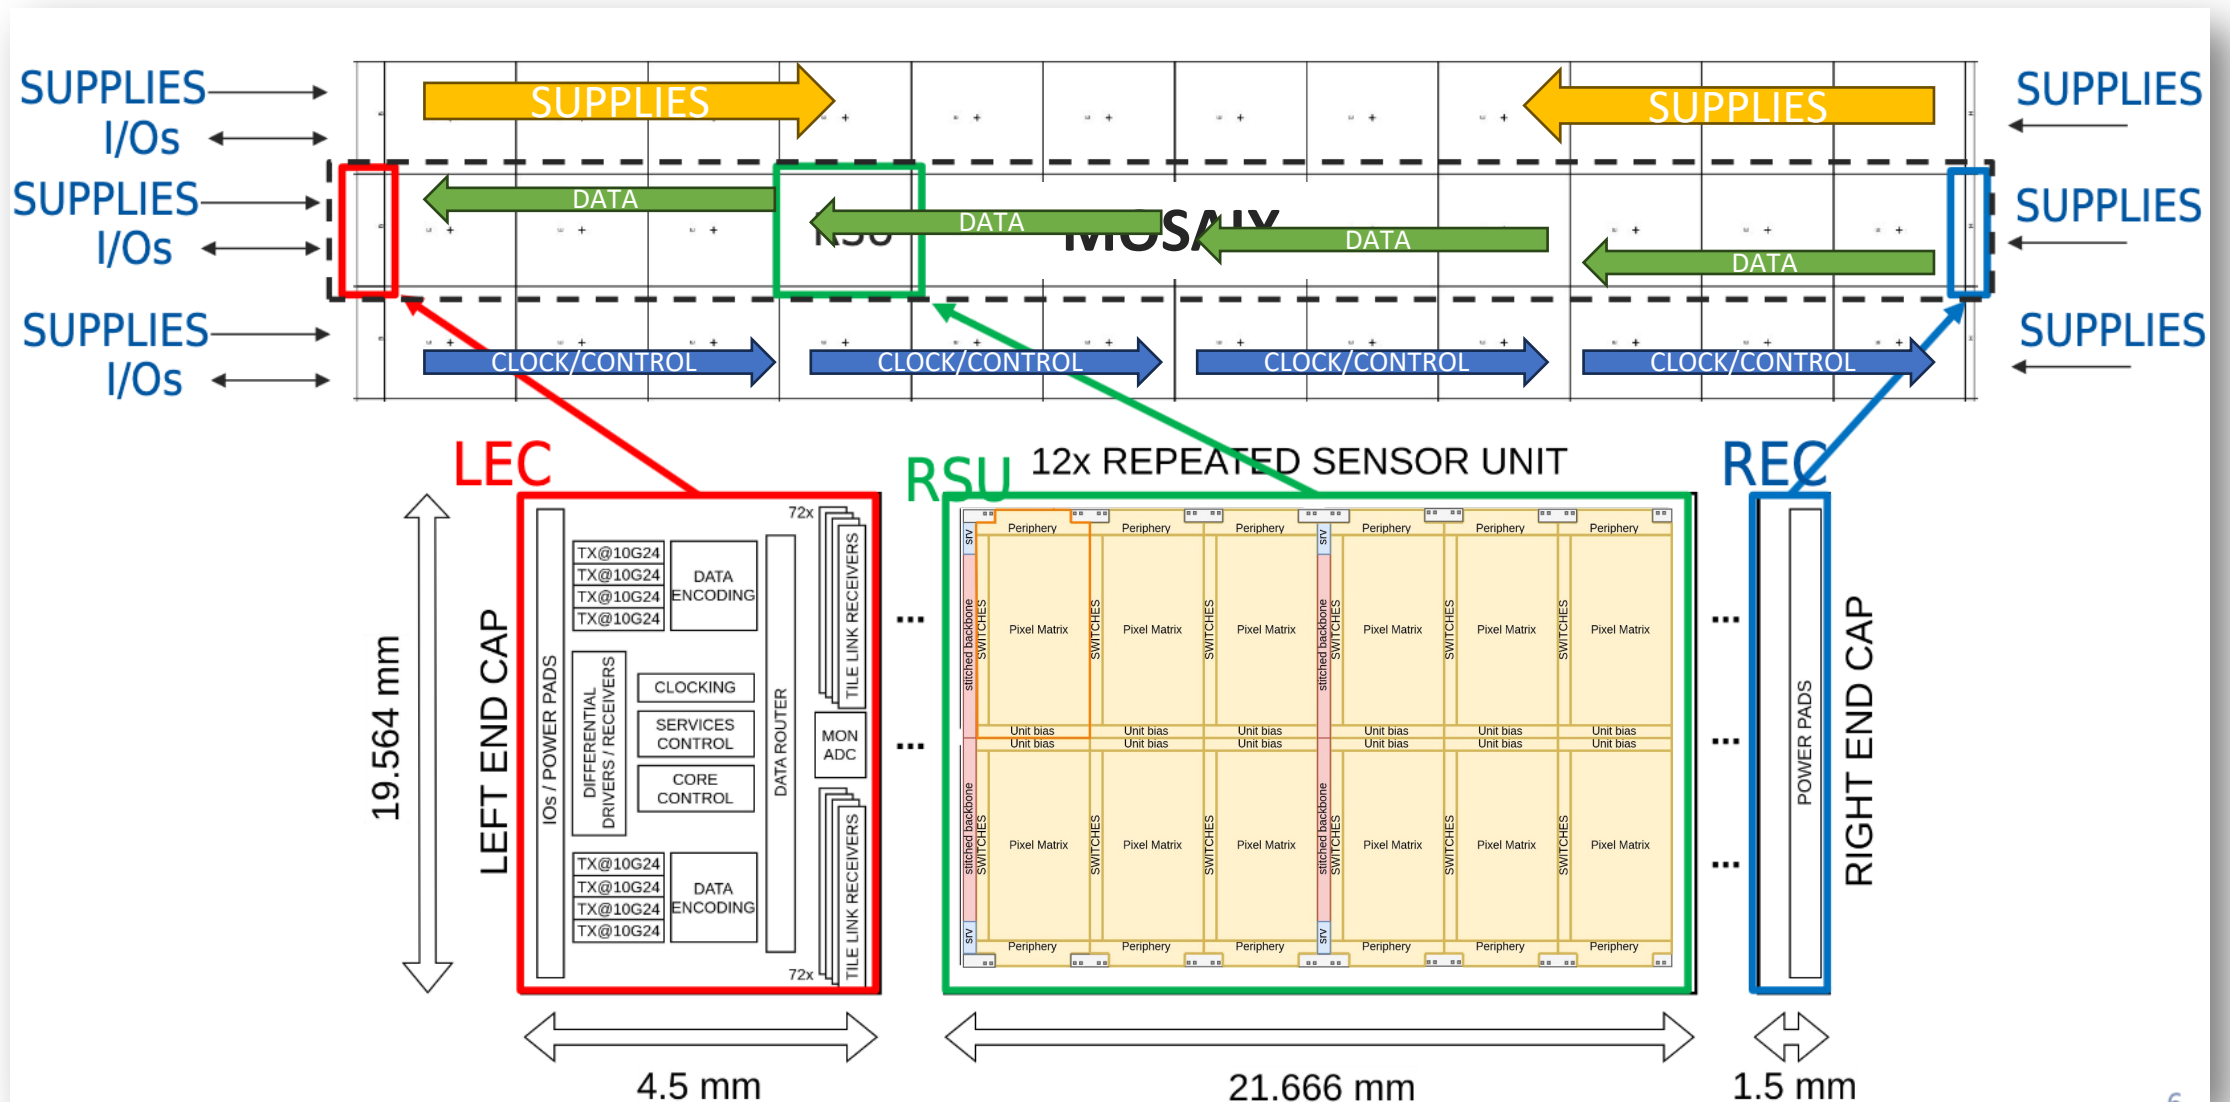


MOSAIX Architecture




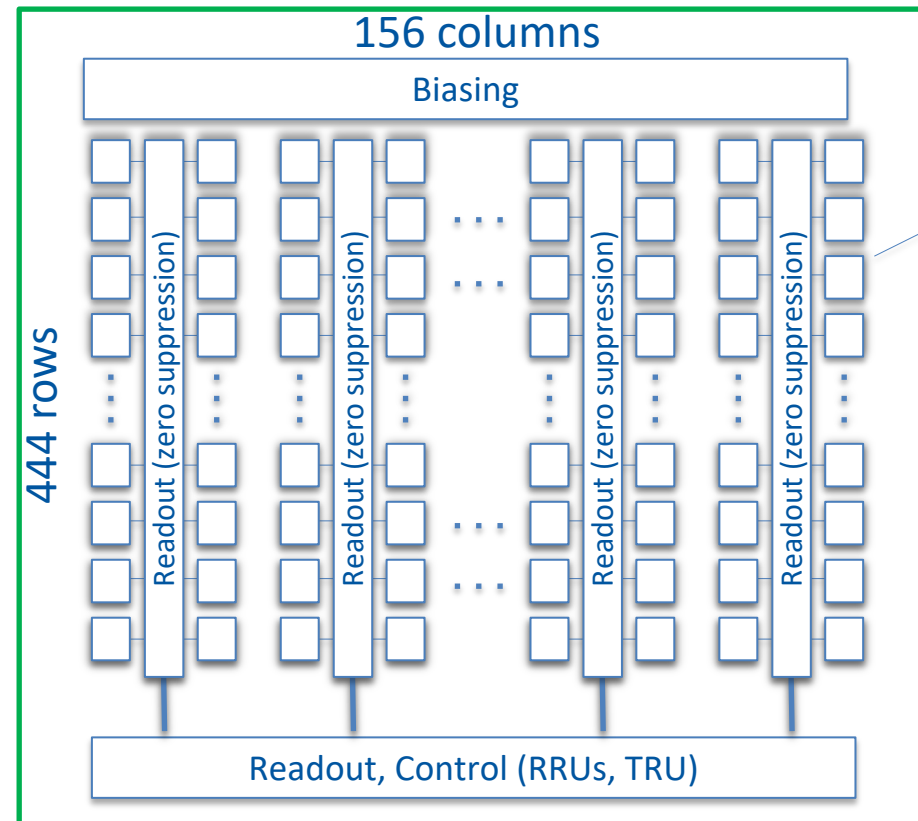
Full staves and module on one silicon die

MOSAIX Architecture



Pixel Matrix Architecture

20.8 μm \times 22.8 μm pixel pitch
444 \times 156 pixels / matrix
144 matrices / segment
10.73 Mpixels / segment



In pixel:

- Amplification
- Discrimination
- Hit integration register and readout register
- Test charge injection
- Masking

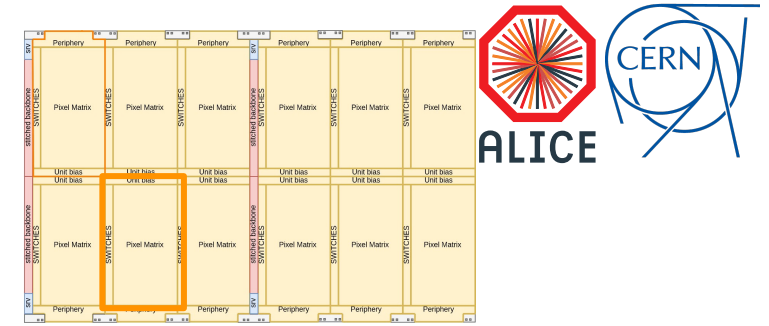
Always active front-end (40 nW typ.)

Global shutter

Zero-suppressed matrix readout

Time-framed continuous readout

Pixel Array Readout



Zero suppressed readout with Priority Encoders

Priority Encoder (PE) encodes the address of the first pixel in the column STATE vector with a hit

PE is fully combinatorial circuit steered by peripheral sequential circuits during the readout of a frame

No free clock over matrix. **No activity where there are no hits**

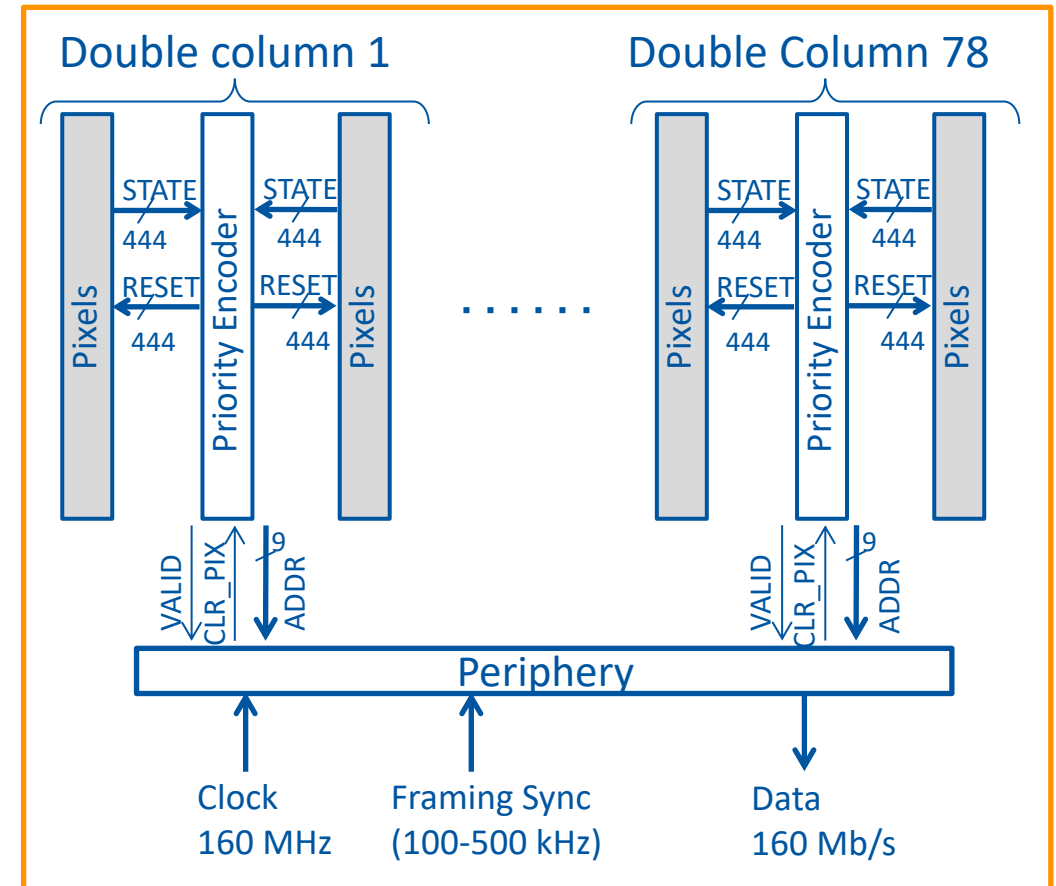
Dynamic energy per hit encoding and transfer over ~10 mm
 $E_h \approx 30 \text{ pJ}$

Leakage power is significantly larger

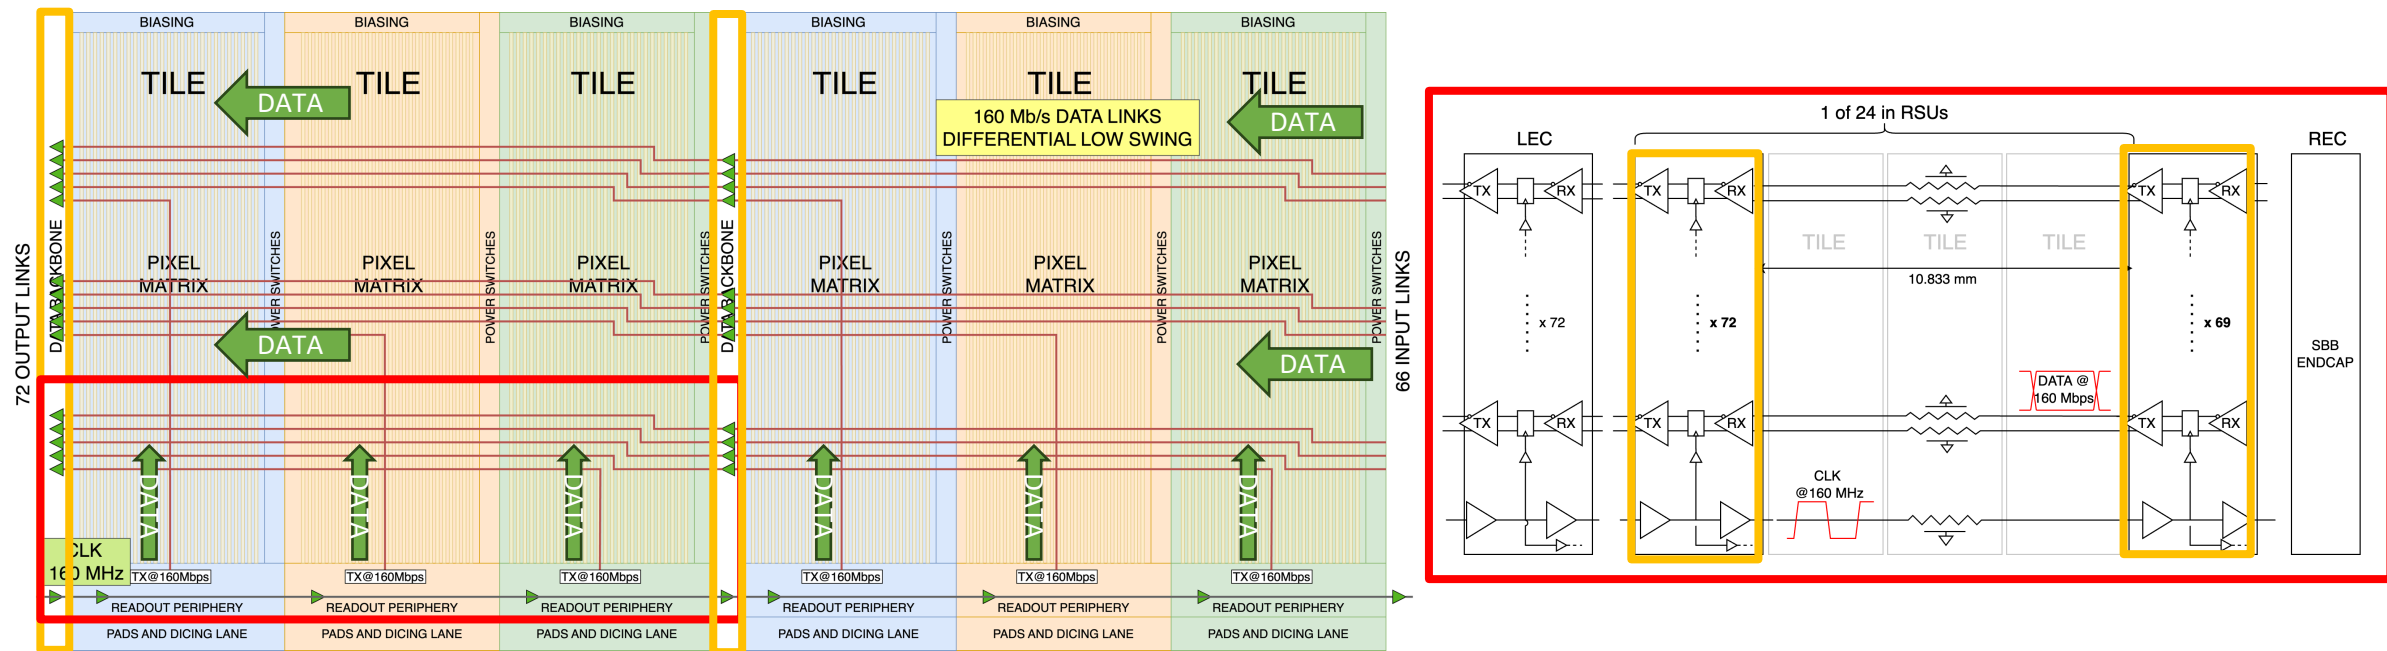
Serial transmission of tile packet to LEC

Tile periphery builds and transmits one data packet for each framing interval (2-10 μs)

Global SYNC signal **aligns in time** the integration intervals across all the tiles



Stitched Data Backbone (SBB)



Distribute 160 MHz clock from LEC to TILES

Long range (25 cm) on-chip 160 Mb/s point to point serial data links, clock synchronous

From 144 tile peripheries to the data aggregation in the Left End Cap

Differential transmission for data links with low voltage swing, and periodic retiming and regeneration

Power efficiency (~ 0.7 pJ/bit/cm), immunity to supply noise, reduction of noise injection into sensing nodes

LEC Left End Cap Architecture



Interfaces and peripheral data hub

Input capacity $144 \times 160 \text{ Mb/s} = 23 \text{ Gb/s}$

No memory and no data processing in LEC

Data Router allows to reroute data from the 144 Tile Links to different serializers

8 High Speed Serializers

Redundancy to mitigate the risk of failures of off-chip optical link components

Two operating modes


10.24 Gb/s line rate: 3 serializers are used

5.12 Gb/s line rate: 6 serializers are used

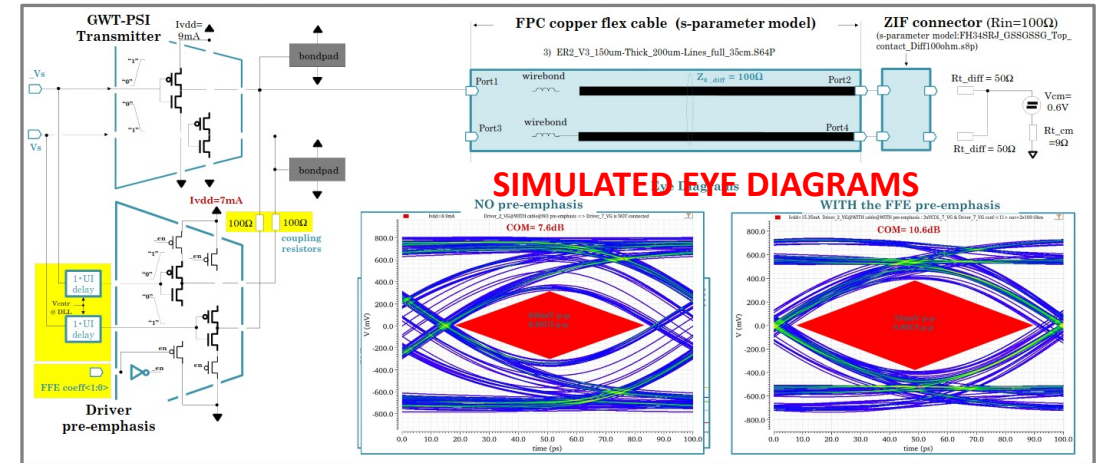
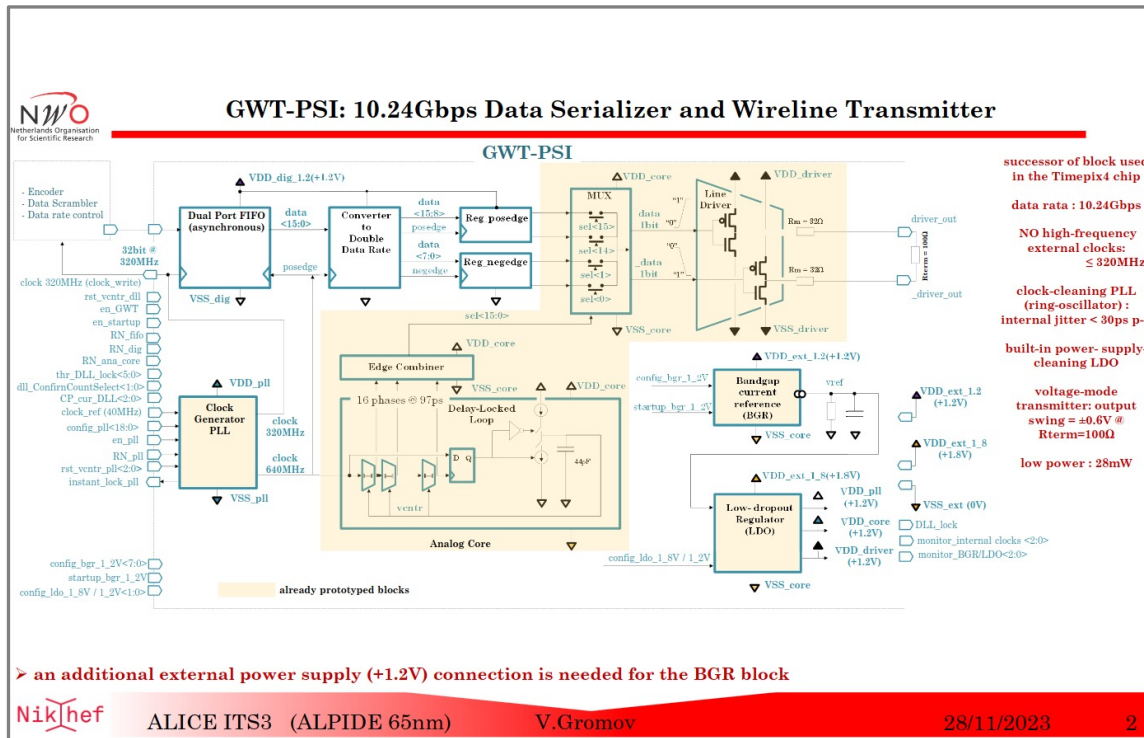
Unused serializers are switched off

MOSAIX can drive electro-optical transceivers directly

Data Encoding block ported from the IpGBT chip



High speed serializers



10.24 Gb/s serializers

Designed for transmission over ~ 30 cm wirelines on flex PCBs
 Driver with pre-emphasis
 Core consumption ~ 30 mW, 45 mW including LDOs (3-5 pJ/bit)
 On-chip linear supply regulation for immunity to supply noise

Power

RSU (sensing area) power consumption

Expected typical value **120 mW**, density **~30 mW/cm²**

Large uncertainty due to process variations and operating settings

Biasing point, reverse bias, temperature

Max estimate 190 mW (44 mW/cm² @ 25 °C)

Data readout and transfer larger than analog sensing


Undesired significant leakage component

Digital circuits consume even idling

Strong variability with process and temperature (doubles every 20 °C)

Mitigated designing a Custom Low Leakage Standard Cell Library

LEC consumption **700 mW** (~ 800 mW/cm²)

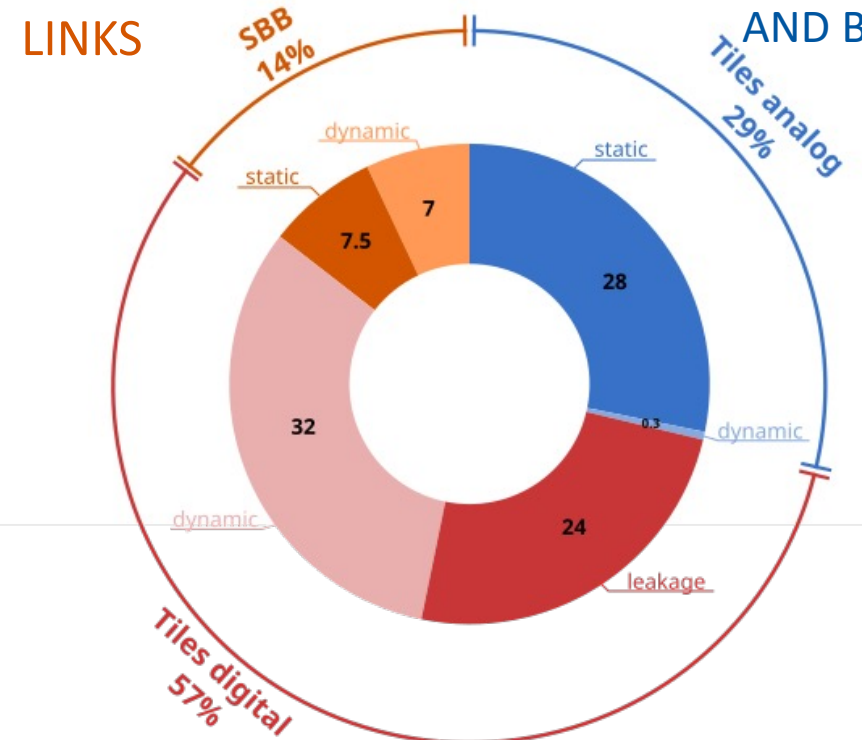


RSU Current Consumption

Typical

ON-CHIP DATA LINKS


ANALOG PIXEL AND BIASING



READOUT

Total Typ. Current 100 mA
Power 120 mW (@1.2 V)

On-chip IR Drops



Significant supply and ground drops


Power connections only from the two short edges

Large size of the device

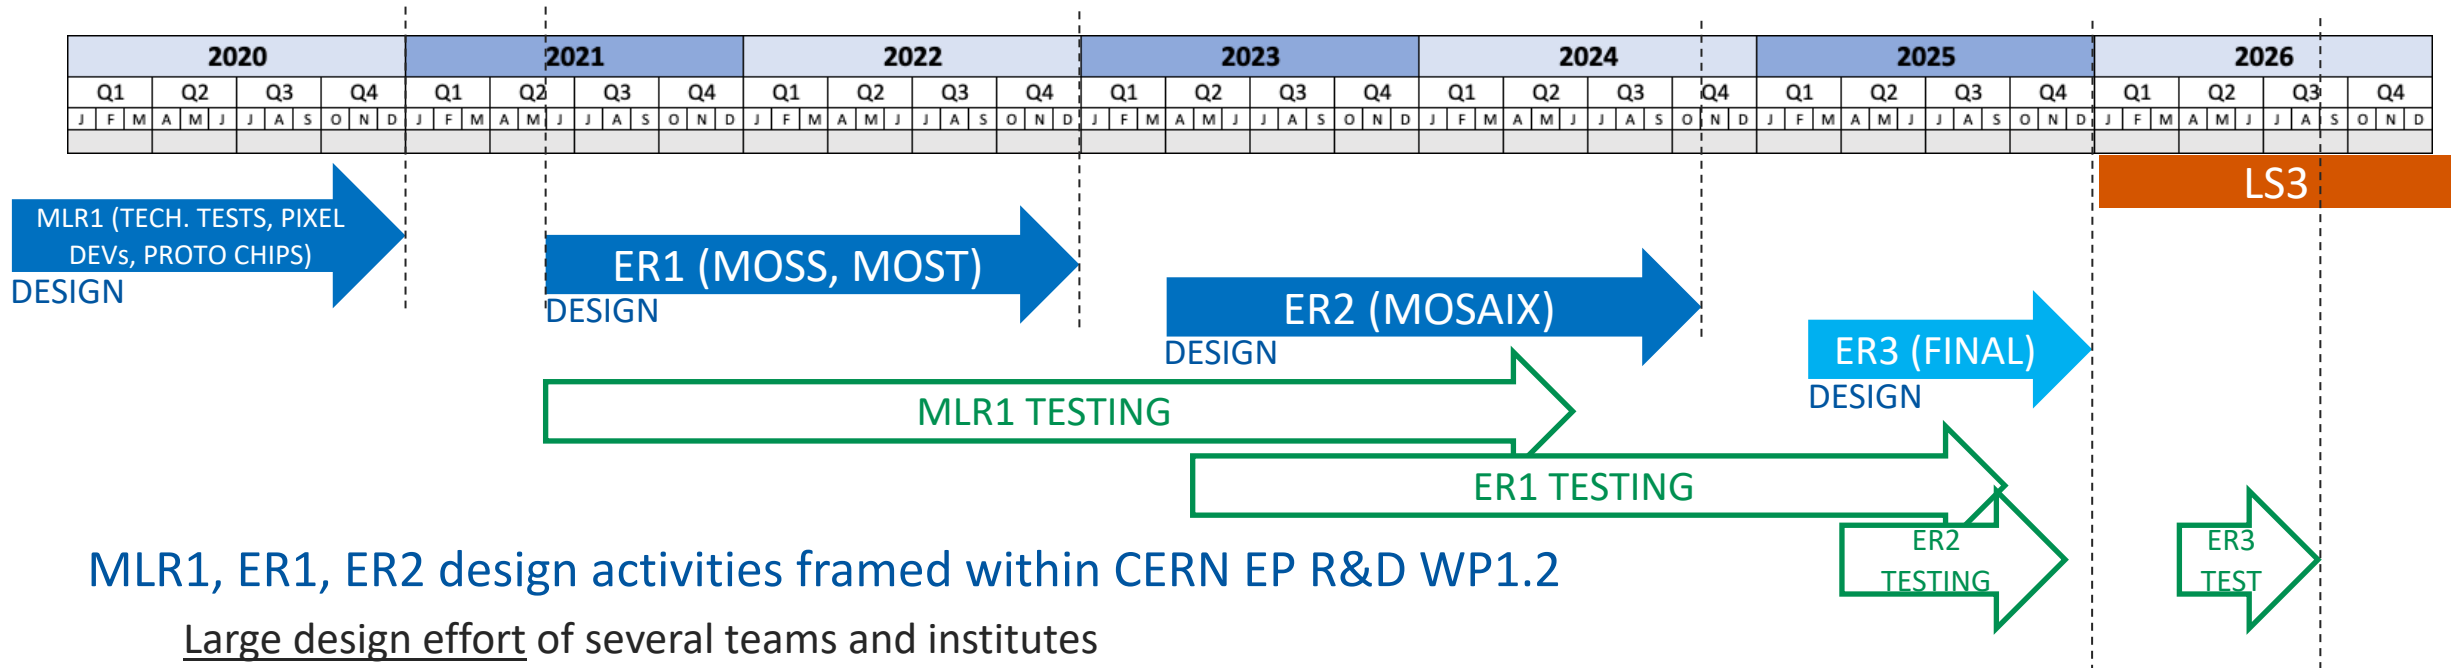
Limited conductivity of the metal layers

Needed new metal stack with thicker top metal

Distributed analog biasing and monitoring circuits



Timeline



MLR1, ER1, ER2 design activities framed within CERN EP R&D WP1.2

Large design effort of several teams and institutes

Design cycles of large complex chips > 1 year

Design teams: ~14-18 persons, effort: ~18-22 person-year / cycle

Testing and characterization effort even larger

$\Theta(40)$ persons, all test chips included. E.g.: ~10 persons *only* for testing MOSS

Development of dedicated hardware, setups and handling shall not be underestimated

Challenges and Learning



Interdependencies and Integration

Design for Manufacturability

Constraints density of features. Widening spacing and widths everywhere not feasible. Custom library of standard cells for DFM
Pixel pitch $O(20 \mu\text{m})$

Fill factor above 94%-95% difficult

Readout of data needs peripheral circuits, whose area and complexity increase with amount of data to move

Power Distribution

IR drops on the metals of the CMOS stack significant even with very low power

Complex segmentation in many independent domains that can be maintained off in case of short circuits

Switches and cross-domain signaling and protections

Significant leakage

Large variations with process and temperature

Needed to devise mitigation techniques, e.g., library of low leakage standard cells

Data transmission

Integrate 144 on-die transmission lines of 25 cm working at 160 Mb/s

High speed (10 Gb/s) wireline drivers for off-chip transmission

Summary



MOSAIX

Wafer scale MAPS sensor made with stitching using TPSCo 65 nm CMOS Imaging

Full feature prototype for the ALICE ITS3 Upgrade (LS3) of the three innermost ALICE tracking and vertexing layers (0.07% X_0 /layer)

Complex integrated system, full staves and modules on one wafer

10.73 Mpixels/segment, $20.8 \mu\text{m} \times 22.8 \mu\text{m}$, 30 mW/cm²

Design ongoing, submission target October 2024

Building on much knowledge acquired from MLR1 and ER1


Test chips and MOSS and MOST stitched sensors

Developments made possible by

- A large design effort of many persons and institutes working in synergy within EP R&D WP1.2
- A very large development, testing, and characterization effort within the ALICE ITS3 project


REFERENCE MATERIAL

DPTS Beam Test Results – Det. Efficiency vs FHR



<https://doi.org/10.1016/j.nima.2023.168589>

DPTS Beam Test Results – Spatial Resolution



(b) Sensors irradiated to different levels.

<https://doi.org/10.1016/j.nima.2023.168589>

MOSS Monolithic Stitched Sensor Prototype




Goals


- Learn techniques to make wafer scale sensors
- Study yield, constraints, powering, spread of parameters

Repeated units abutting on short edges

- 10 Repeated Sensor Units, 1 Endcap Left, 1 Endcap Right
- Metal traces cross stitching boundaries for power distribution and long range on-chip control and data transfer
- Module integration on wafer scale die for the first time
- 1/20 power segmentation

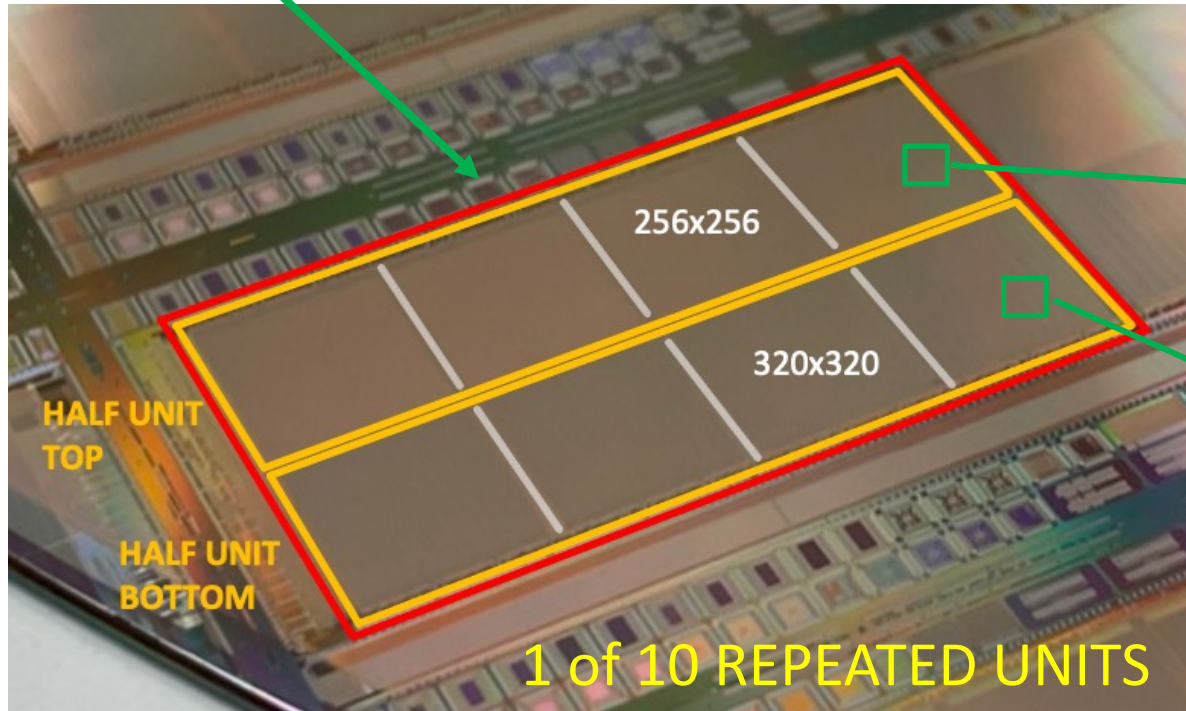



Reticle





259 mm

MOSS Layout



Large Pixel Pitch
22.5 μm

Fine Pixel Pitch
18 μm



Investigate effects
of layout density

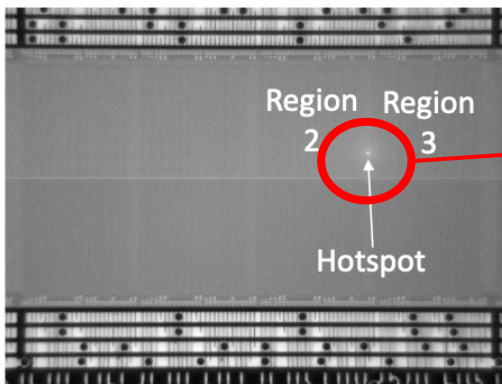
MOSS Testing – Powering Yield




- Dominant failure mode: short circuits between power nets
- Long and intense investigations. Finding: unexpected inter-metal vertical shorts
 - Related to manufacturing.
 - Wafer to wafer variations.
 - Followed-up with foundry. Expected to disappear or reduce with new metal stack and mitigation by layout

- Results before and after thinning and dicing consistent


Thermal Imaging reveals shorts



Optical Microscope on due to wafer post processing



SEM image of metal-to-metal short



CERN-LHCC-2024-003 / ALICE-TDR-021

Powering tests from chips of the first three wafers tested. The chips were thinned, diced, glued and bonded before testing.

Wafer	1-TOP	1-BOT	2-TOP	2-BOT	3-TOP	3-BOT	4-TOP	4-BOT	5-TOP	5-BOT	6-TOP	6-BOT	RSU1	RSU2	RSU3	RSU4	RSU5	RSU6	RSU7	RSU8	RSU9	RSU10	
Wafer 17	LIMIT	OK - I	OK - I	LIMIT	LIMIT	LIMIT	LIMIT	OK - I	OK - I	OK - I	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	
	OK - II	OK - II	OK - II	OK - II	LIMIT	LIMIT	LIMIT	LIMIT	OK - II	OK - I	OK - I	OK - I	mechanically broken										
	LIMIT	OK - I	OK - II	LIMIT	LIMIT	LIMIT	LIMIT	LIMIT	OK - I	OK - II	LIMIT	LIMIT	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	
	OK - I	OK - I	OK - I	LIMIT	LIMIT	LIMIT	LIMIT	LIMIT	OK - I	OK - II	OK - I	OK - I	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	
	OK - I	LIMIT	OK - I	LIMIT	LIMIT	LIMIT	LIMIT	LIMIT	LIMIT	OK - II	LIMIT	LIMIT	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	
	OK - II	OK - I	OK - II	LIMIT	LIMIT	LIMIT	LIMIT	LIMIT	LIMIT	OK - I	OK - I	OK - I	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	
Wafer 23	OK - II	OK - I	OK - II	LIMIT	OK - II	OK - II	OK - I	LIMIT	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	
	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	
	OK - II	OK - I	LIMIT	LIMIT	LIMIT	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	OK - II	
	OK - II	OK - I	LIMIT	LIMIT	LIMIT	LIMIT	LIMIT	LIMIT	OK - I	OK - II	LIMIT	LIMIT	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I
	OK - II	OK - I	OK - I	OK - II	OK - II	OK - I	OK - II	OK - II	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I
	OK - I	LIMIT	LIMIT	LIMIT	OK - I	OK - II	OK - II	OK - II	OK - II	OK - II	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I
Wafer 24	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	
	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	
	OK - II	OK - I	OK - I	OK - II	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	
	LIMIT	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	
	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I
	OK - I	LIMIT	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I	OK - I

MOSS Characterization with Beams





- Sensor performance characterization ongoing

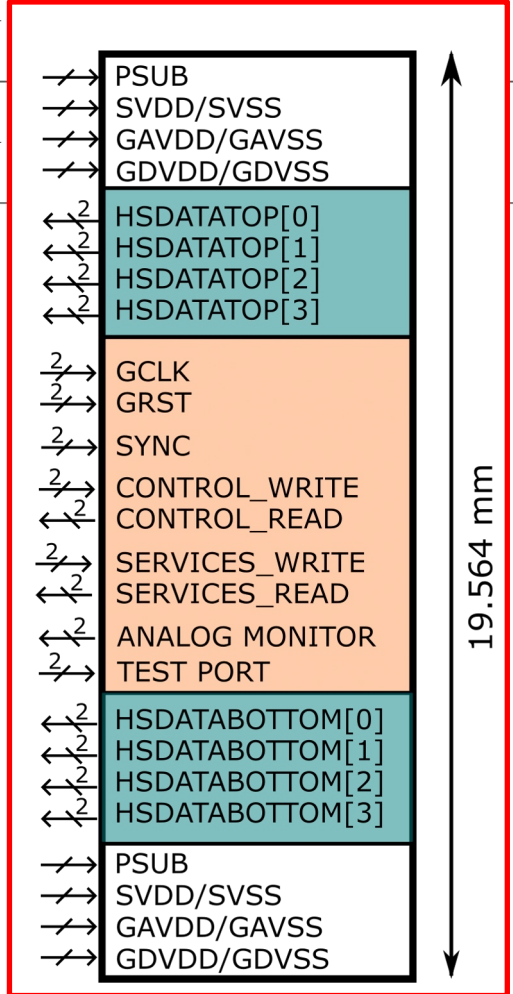

- Based on laboratory measurements and multiple **beam tests** with Full MOSS and Single-Stitch MOSS
- Studying detection efficiency, FHR, position resolution and tuning operating settings
- Compare 6 variants of pixels × 2 process splits × Non-irradiated and NIEL Irradiated samples

- Cross sections of SEU and SEL events

- Beam tests with Single Stitch MOSS
- SEUs as expected
- Indications of sensitivity to SEL, will investigate to localize and mitigate



MOSAIX Supplies and I/Os



10 Gb/s DATA OUTPUTS

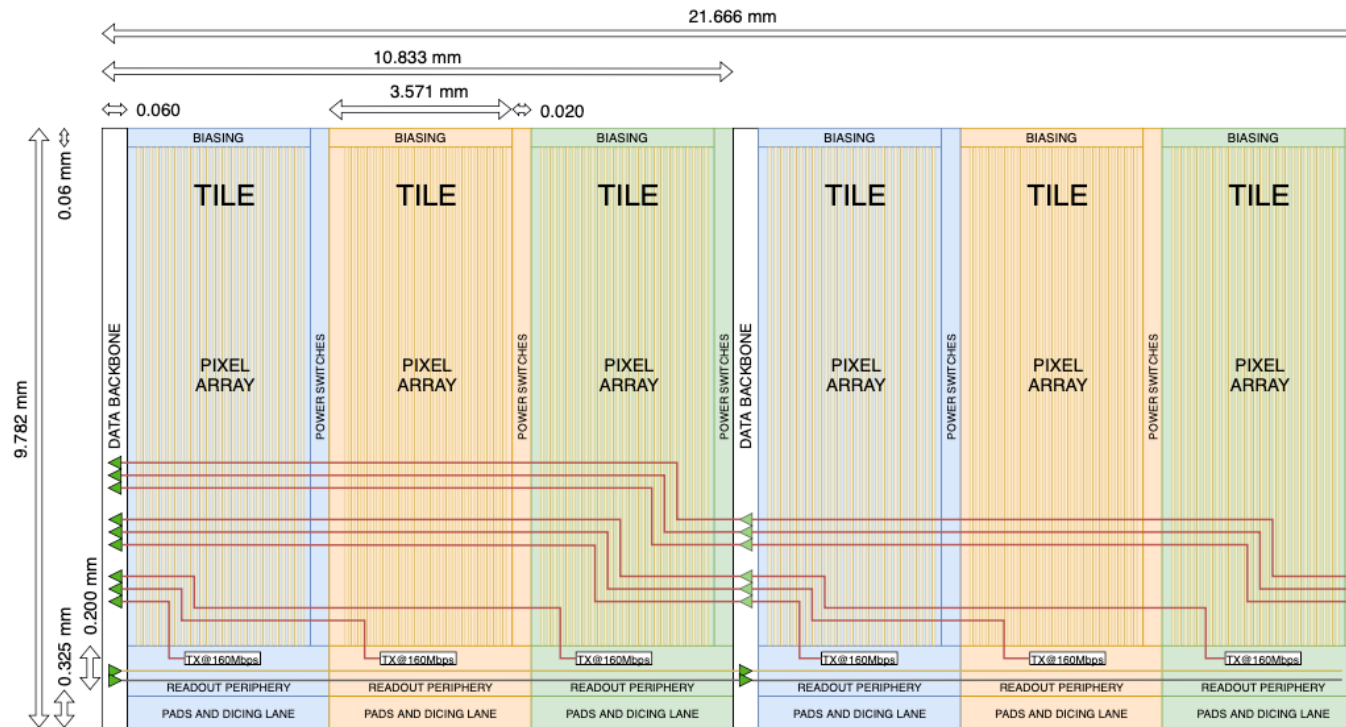
160 MHz Clock

5 Mb/s Slow Control Ports

10 Gb/s DATA OUTPUTS

PSUB Substrate reverse biasing (full die)
 Services, Analog, Digital supplies (per segment)

RSU – Floorplan



Current design fill factor ~93%

Depending on ER2 results, optional removal of test pads could reach 95.5%

Block	Width [mm]	Height [mm]	Block Area [mm ²]	Instances	Percent area
RSU	21.666	19.564	423.873	1	100 %
Pixel Matrix	3.571	9.197	32.843	12	92.98 %
Biasing	3.571	0.060	0.214	12	0.61 %
Power switches	0.020	9.257	0.185	12	0.52 %
Data Backbone	0.060	9.257	0.555	4	0.52 %
Readout periphery	3.591	0.200	0.718	12	2.03 %
Test pads	21.666	0.250	5.416	2	2.56 %
Seal ring and dicing lane	21.666	0.075	1.625	2	0.77 %

Table 3.4: Plan of dimensions of the blocks composing one Repeated Sensor Unit and percentage of the RSU area occupied by the instances of the block.

Pixel Diagram

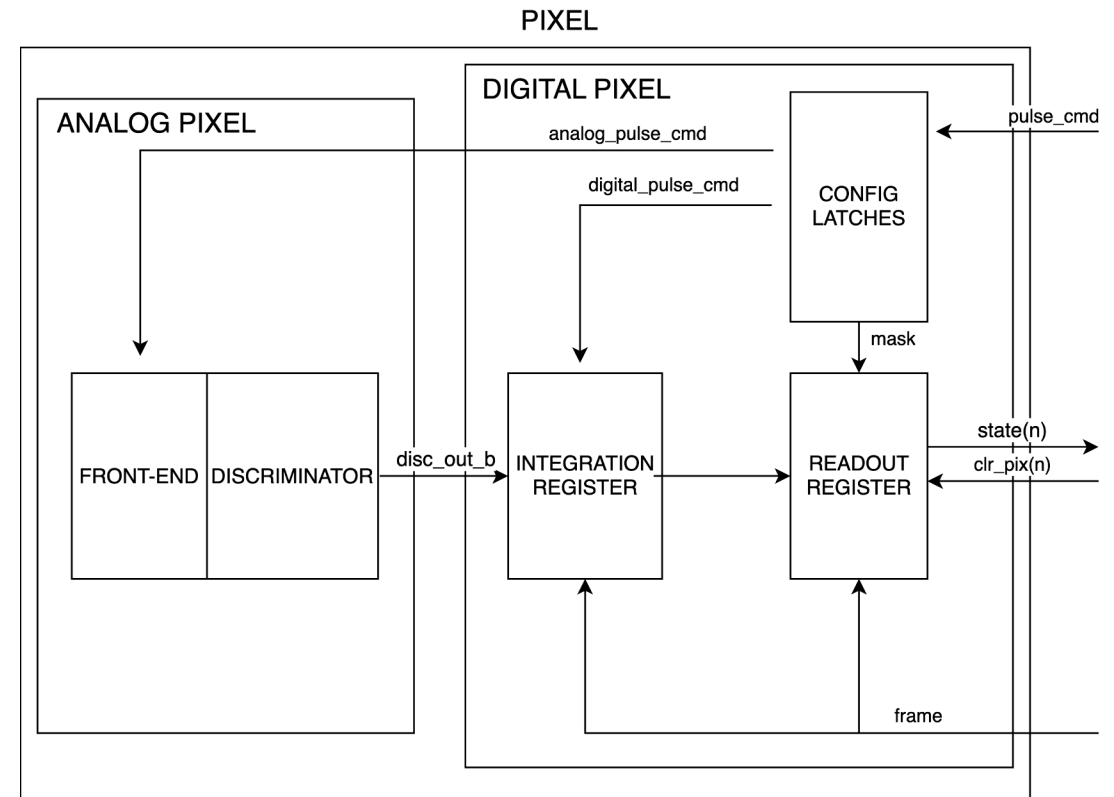


Analog front-end and discriminator continuously active

- Test pulse charge injection circuitry
- Global threshold and analog settings per TILE

Digital pixel

- Edge sensitive recording of FE discriminator output edges
- 2-stage hit data buffer
- In-pixel **masking** and **pulsing** configuration



Analog Front-End



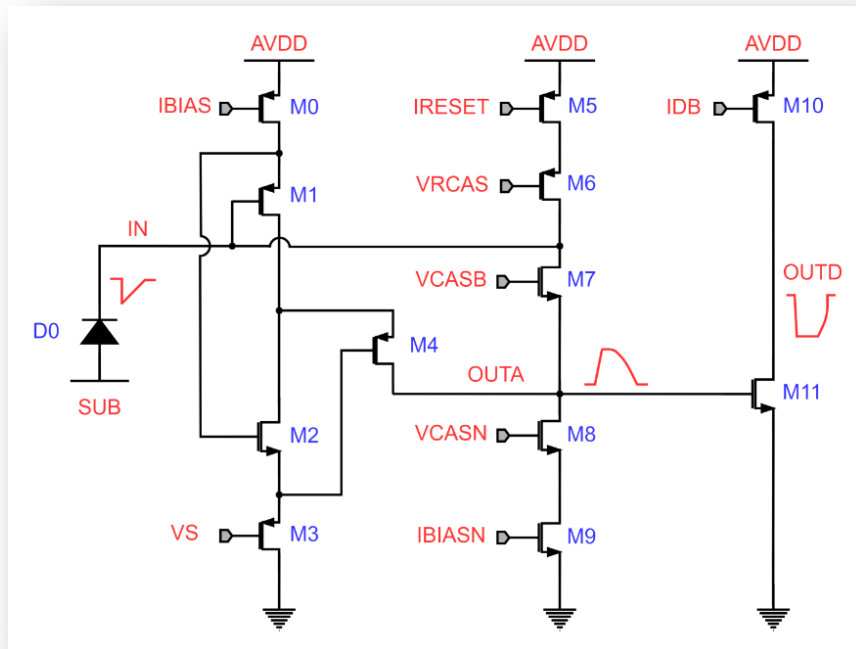
Evolution path

DPTS -> MOSS -> ER2

In-depth revision of transistor sizing

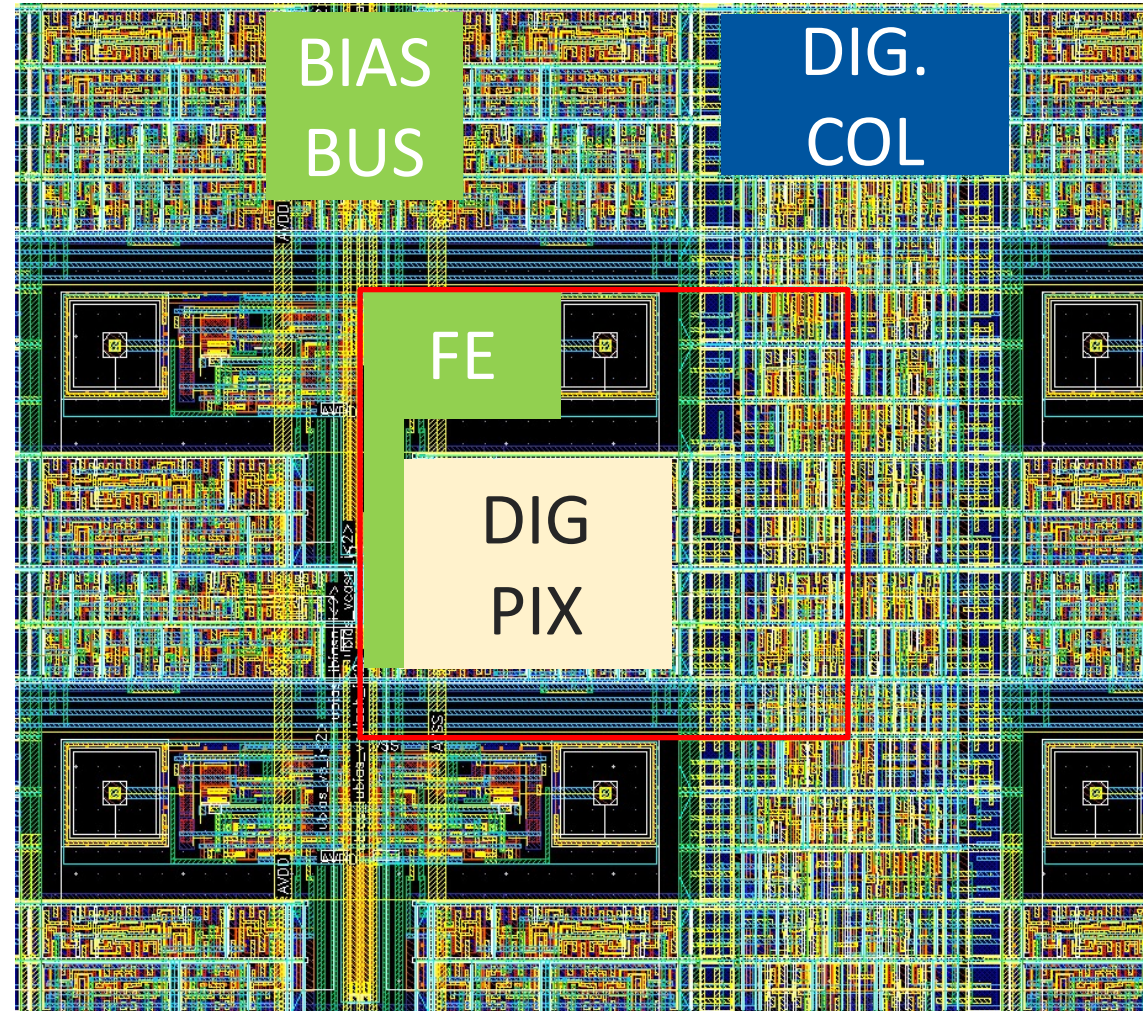

Reduction of spreads, reduce sensitivity

New layout, area: 40 μm^2



Specification		Value	Note
Power consumption	Static current consumption	< 30 nA	IBIAS + discriminator standby current
	Dynamic energy (@600e-)	< ~10 pJ	Not an hard requirement, it seems negligible wrt the total power budget
Nominal threshold		~ 150 e-	¼ of MIP
ENC		< 18 e-	
Threshold mismatch		< 18 e-	
Gain (@TH)		> 400 $\mu\text{V}/\text{e}^-$	Simulated avoiding discriminator kick-back
Phase margin		> 45°	
Timing	TOA (@150 e-)	< 1 μs	To cope with 100 kHz strobe
	TOT_a (@600 e-)	\ll 1 ms	To have a lost hit probability < 1%
Threshold sensitivity vs supply levels drop		< 2 e-/mV	Supply drop on both AVDD and AVSS

Pixel Matrix Layout Snapshot



Data Readout Processes in the Periphery



Framing time base re-generated in each TILE periphery

FRAME local signal synchronizes pixels, Region Readout and Top Readout

Global SYNC input signal *aligns in time* the integration intervals across tiles

Four parallel readout processes in each tile

Regions have 38 or 40 columns

Double columns in one region are sequentially read out

Region data packet is stored in FIFOs

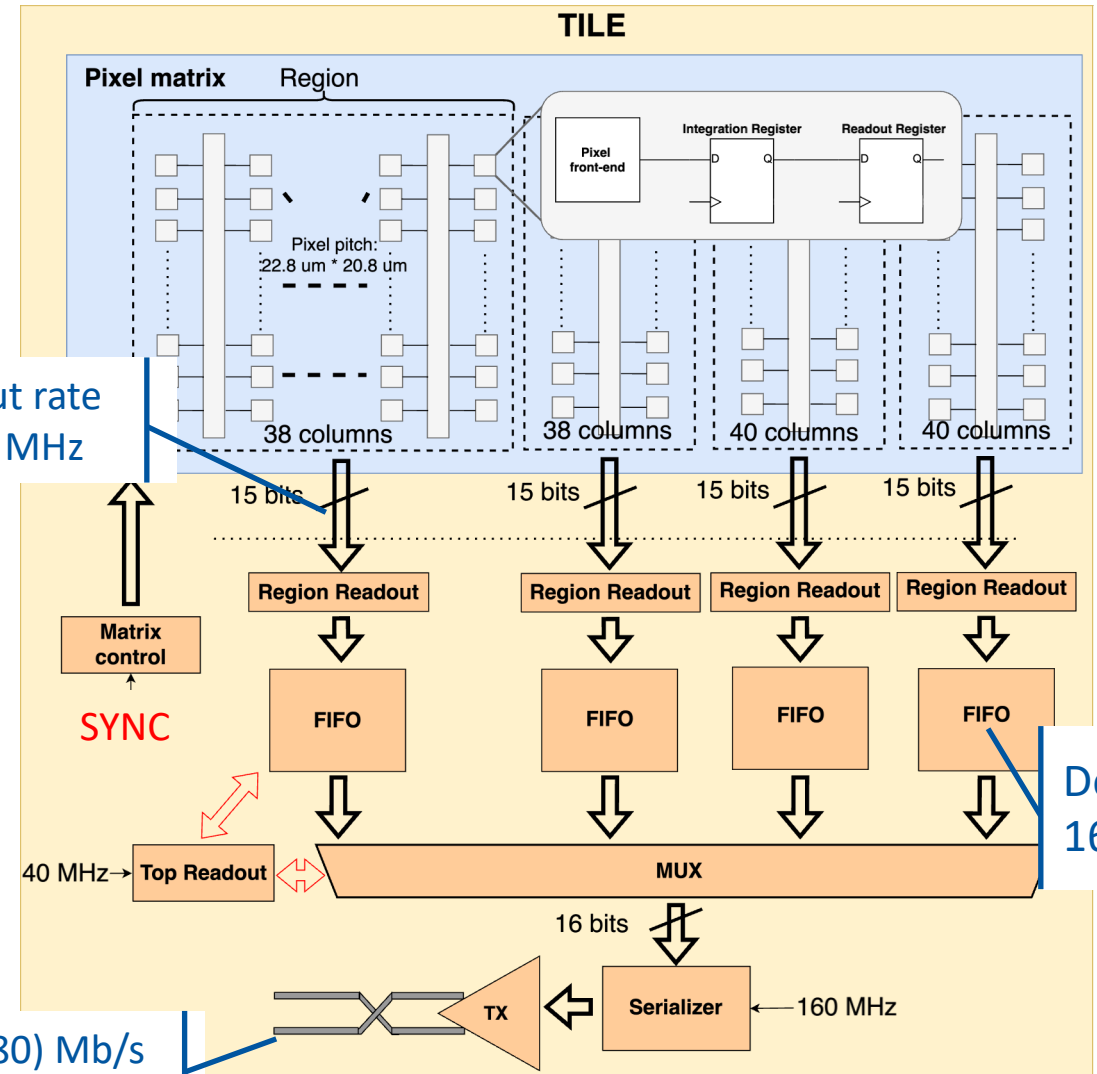
Double columns and full regions can be masked

Serial transmission of tile packet to LEC

Top readout aggregates region data packets for the same frame interval

Tile transmits one data packet for each frame interval, in order

Hit readout rate
40 (or 20) MHz



160 (or 80) Mb/s
144 links/segment

Readout Simulation Model



Simulation parameters	Value	Unit	Conditions
Particle Rates			
Average Pb-Pb Interaction Rate	164	kHz	
Particle flux (Hadronic)	2.55	MHz cm ⁻²	z=0 cm, all centralities.
Particle flux (QED)	3.20	MHz cm ⁻²	z=0 cm.
Total particle flux	5.75	MHz cm ⁻²	z=0 cm, all centralities.
Geometry, timing, encoding, data transfer capacity			
Pixel dimensions	20.8 × 22.8	μm × μm	
Tile pixel array size	442 × 156		
Pixels per Tile	68952		
Sensitive Area of the tile	0.327	cm ²	
Tiles per segment	144		
Readout regions per tile	3 or 4		
Frame Interval Duration (FD)	2 or 5	μs	
Minimum average cluster size	2.1		Δz = 0 cm, fig. 3.41.
Maximum average cluster size	6.3		Δz = 13.5 cm, fig. 3.41.
Pixel hit encoding time	25	ns	
Bits per pixel hit	16	bit	
Capacity of tile link	160	Mbit s ⁻¹	
Aggregated capacity (Segment)	23.04	Gbit s ⁻¹	

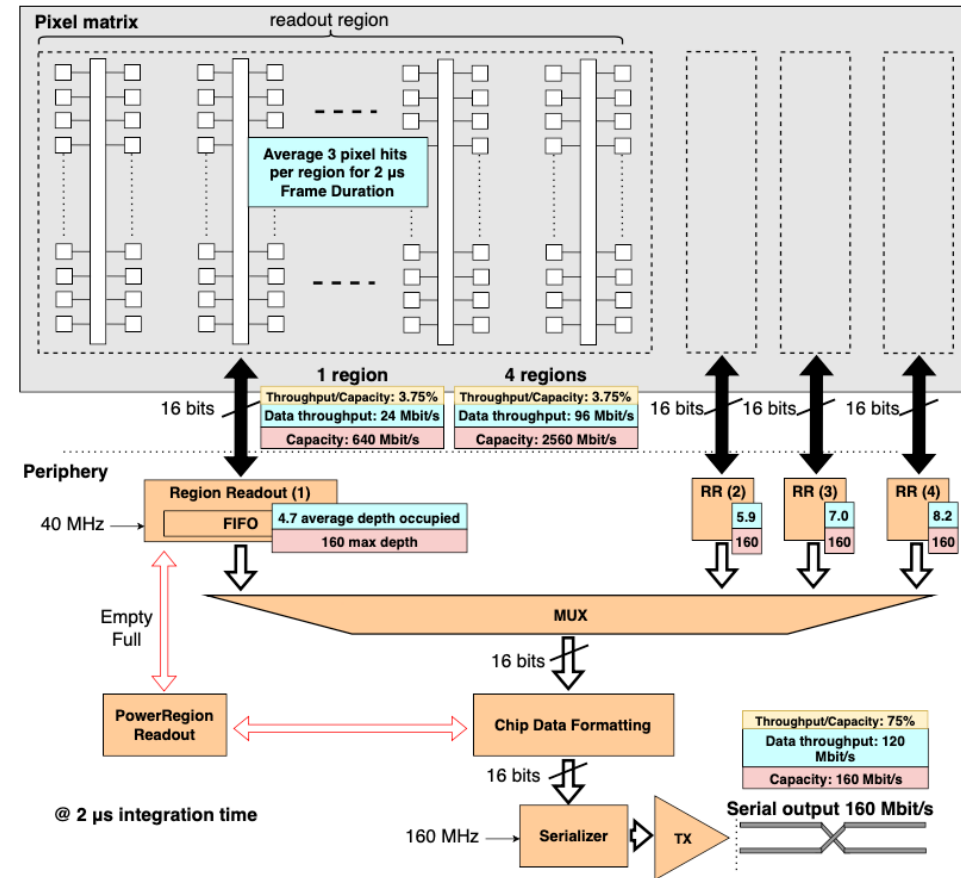



Figure 3.40: Block diagram of the peripheral readout of one tile. The illustration assumes four readout regions, FIFOs of 160 words depth and an integration time of 2 μs. The values of occupancy and data throughput of the internal busses reported on the diagram are results from the simulation runs.

Readout performance

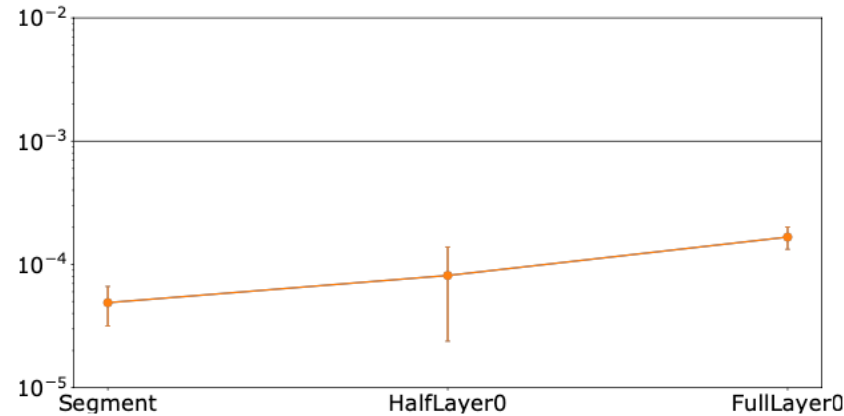


Figure 3.44: Readout performance simulation results. Fraction of collisions with missing data for one L0 segment, for three segments (half layer 0) and for six segments (full layer 0). The results are for 4 regions per tile, FIFO depth of 160 words and frame duration of 2 μ s.

Simulation results

Average pixel occupancy	$< 2.0 \times 10^{-4}$		$z=0$ cm.
Average pixel occupancy	$< 5.0 \times 10^{-4}$		$z=0$ cm, FD=5 μ s.
Data throughput	120	Mbit s ⁻¹ Tile ⁻¹	$z=0$ cm.
Data throughput	15.55	Gbit s ⁻¹ Segment ⁻¹	
Data throughput per unit area	365	Mbit s ⁻¹ cm ⁻²	$z=0$ cm.
Data throughput per unit area	329	Mbit s ⁻¹ cm ⁻²	Average over z .
Data throughput per link	2.58	Gbit s ⁻¹	
Incomplete event probability	$< 6 \times 10^{-5}$		Layer 0 segment.
Incomplete event probability	$< 2 \times 10^{-4}$		Full layer 0.