



Full Length Article

## Characterisation of the first wafer-scale prototype for the ALICE ITS3 upgrade: The monolithic stitched sensor (MOSS)

Omar Abdelrahman<sup>1</sup>, Gianluca Aglieri Rinella<sup>2</sup>, Luca Aglietta<sup>3,4</sup>, Giacomo Alocco<sup>3,4</sup>, Matias Antonelli<sup>5,6</sup>, Roberto Bacconi<sup>6</sup>, Francesco Barile<sup>7,8</sup>, Pascal Becht<sup>9</sup>, Franco Benotto<sup>4</sup>, Stefania Maria Beolè<sup>3,4</sup>, Marcello Borri<sup>10</sup>, Daniela Bortoletto<sup>11</sup>, Naseem Bouchhar<sup>12</sup>, Giuseppe Eugenio Bruno<sup>7,8</sup>, Matthew Daniel Buckland<sup>10</sup>, Szymon Bugiel<sup>2</sup>, Paolo Camerini<sup>5,6</sup>, Francesca Carnesecchi<sup>2</sup>, Marielle Chartier<sup>13</sup>, Domenico Colella<sup>7,8</sup>, Angelo Colelli<sup>7,8</sup>, Giacomo Contin<sup>5,6</sup>, Giuseppe De Robertis<sup>8</sup>, Wenjing Deng<sup>14</sup>, Antonello Di Mauro<sup>2</sup>, Vittorio Di Trapani<sup>5,6</sup>, Maurice Donner<sup>9</sup>, Ana Dorda Martin<sup>2</sup>, Piotr Dorosz<sup>2</sup>, Floarea Dumitrache<sup>3,4</sup>, Lars Döpper<sup>15</sup>, Gregor Hieronymus Eberwein<sup>11</sup>, Domenico Elia<sup>8</sup>, Simone Emiliani<sup>2</sup>, Laura Fabbietti<sup>1</sup>, Tommaso Fagotto<sup>5,6</sup>, Xiaochao Fang<sup>16</sup>, Henrik Friberth<sup>1</sup>, Roman Gernhäuser<sup>1</sup>, Piero Giubilato<sup>17,18</sup>, Laura Gonella<sup>5,6</sup>, Karl Gran Grodaas<sup>2,19</sup>, Ola Slettevoll Groettvik<sup>2</sup>, Vladimir Gromov<sup>20</sup>, Malte Grönbeck<sup>15</sup>, Philip Hauer<sup>15</sup>, Hartmut Hillemanns<sup>2,\*,\*</sup>, Guen Hee Hong<sup>21</sup>, Yu Hu<sup>22</sup>, Minyoung Chris Hwang<sup>23</sup>, Marc Alain Imhoff<sup>16</sup>, Barbara Jacak<sup>22,23</sup>, Daniel Matthew Jones<sup>13</sup>, Antoine Junique<sup>2</sup>, Filip Křížek<sup>24</sup>, Jetnipit Kaewjai<sup>25</sup>, Anouk Kaiser<sup>9</sup>, Jesper Karlsson Gumprecht<sup>26</sup>, Markus Keil<sup>2</sup>, Bernhard Ketzer<sup>15</sup>, Jiyoung Kim<sup>27</sup>, Lena Kirchner<sup>1</sup>, Kritsada Kittimanapun<sup>28</sup>, Alex Kluge<sup>2</sup>, Chinorat Kobdaj<sup>25</sup>, Artem Kotliarov<sup>24</sup>, Thanushan Kugathasan<sup>2</sup>, Marc König<sup>2,29</sup>, Paola La Rocca<sup>30,31</sup>, Natthawut Laojamnongwong<sup>25</sup>, Lukas Lautner<sup>1</sup>, Corentin Lemoine<sup>2</sup>, Long Li<sup>32</sup>, Beatrice Eva Liang-Gilman<sup>23</sup>, Francesco Licciulli<sup>8</sup>, Sanghoon Lim<sup>33</sup>, Bong-Hwi Lim<sup>34</sup>, Jian Liu<sup>13</sup>, Flavio Loddo<sup>8</sup>, Matteo Lupi<sup>2</sup>, Magnus Mager<sup>2</sup>, Philipp Mann<sup>1</sup>, Georgios Mantzaridis<sup>1</sup>, Davide Marras<sup>35,36</sup>, Paolo Martinengo<sup>2</sup>, Silvia Masciocchi<sup>9</sup>, Annalisa Mastroserio<sup>37,8</sup>, Soniya Mathew<sup>10</sup>, Serena Mattiazzo<sup>17,18</sup>, Marius Wilm Menzel<sup>2,9</sup>, Nicola Minafra<sup>38</sup>, Frederic Morel<sup>16</sup>, Alice Mulliri<sup>35,36</sup>, Luciano Musa<sup>2</sup>, Anjali Ila Nambrath<sup>23</sup>, Rajendra Nath Patra<sup>8</sup>, Iaroslav Panasenko<sup>26</sup>, Styliani Paniskaki<sup>26</sup>, Caterina Pantouvakis<sup>17,18</sup>, Cosimo Pastore<sup>8</sup>, Stefania Perciballi<sup>3,4</sup>, Francesco Piro<sup>2,39</sup>, Adithya Pulli<sup>2</sup>, Alexander Rachev<sup>15</sup>, Alexander Rachevski<sup>6</sup>, Ivan Ravasenga<sup>2</sup>, Felix Reidt<sup>2</sup>, Michele Rignanese<sup>17,18</sup>, Giacomo Ripamonti<sup>2</sup>, Isabella Sanna<sup>2</sup>, Valerio Sarritzu<sup>35,36</sup>, Umberto Savino<sup>3,4</sup>, Iain Sedgwick<sup>10</sup>, Serhiy Senyukov<sup>16</sup>, Danush Shekar<sup>40</sup>, Sabyasachi Siddhanta<sup>36</sup>, David Silvermyr<sup>26</sup>, Walter Snoeys<sup>2</sup>, Joey Staa<sup>26</sup>, Alessandro Sturniolo<sup>30,31</sup>, Miljenko Šuljić<sup>2</sup>, Timea Szollosova<sup>41</sup>, Daniel Tapia Takaki<sup>38</sup>, Livia Terlizzi<sup>2</sup>, Nicolas Tiltmann<sup>2,43</sup>, Antonio Trifirò<sup>30,31</sup>, Christina Tsolanta<sup>42</sup>, Rosario Turrisi<sup>18</sup>, Berkin Ulukutlu<sup>1</sup>, Gianluca Usai<sup>35,36</sup>, Isabelle Valin<sup>16</sup>, Giovanni Vecil<sup>5,6</sup>, Pedro Vicente Leitao<sup>2</sup>, Anna Villani<sup>5,6</sup>, Chunzheng Wang<sup>44</sup>, Zhenyu Ye<sup>22</sup>, Emma Rose Yeats<sup>23</sup>, Asli Yelkenci<sup>20</sup>, Zijun Zhao<sup>14</sup>, Alessandra Zingaretti<sup>17,18</sup>

<sup>1</sup> Technische Universität München, Munich, Germany

\* Corresponding author.

E-mail address: [hartmut.hillemanns@cern.ch](mailto:hartmut.hillemanns@cern.ch) (H. Hillemanns).

<https://doi.org/10.1016/j.nima.2026.171297>

Received 14 October 2025; Received in revised form 19 December 2025; Accepted 15 January 2026

Available online 23 January 2026

0168-9002/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

- <sup>2</sup> European Organisation for Nuclear Research (CERN), Meyrin, Switzerland  
<sup>3</sup> Università degli studi di Torino, Torino, Italy  
<sup>4</sup> Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Torino, Torino, Italy  
<sup>5</sup> Università degli studi di Trieste, Trieste, Italy  
<sup>6</sup> Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Trieste, Trieste, Italy  
<sup>7</sup> Dipartimento Interateneo di Fisica 'M. Merlin', Bari, Italy  
<sup>8</sup> Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Bari, Bari, Italy  
<sup>9</sup> Ruprecht Karls Universität Heidelberg, Heidelberg, Germany  
<sup>10</sup> STFC Daresbury Laboratory, Daresbury, United Kingdom  
<sup>11</sup> University of Oxford, Oxford, United Kingdom  
<sup>12</sup> Sejong University, Seoul, South Korea  
<sup>13</sup> University of Liverpool, Liverpool, United Kingdom  
<sup>14</sup> Central China Normal University (CCNU), Wuhan, China  
<sup>15</sup> Rheinische Friedrich-Wilhelms-Universität Bonn, Bonn, Germany  
<sup>16</sup> Centre National de la Recherche Scientifique, Strasbourg, France  
<sup>17</sup> Università degli studi di Padova, Padova, Italy  
<sup>18</sup> Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Padova, Padova, Italy  
<sup>19</sup> Norwegian University of Science and Technology (NTNU), Trondheim, Norway  
<sup>20</sup> Nikhef National institute for subatomic physics, Amsterdam, Netherlands  
<sup>21</sup> Yonsei University, Seoul, South Korea  
<sup>22</sup> Lawrence Berkeley National Laboratory, Berkeley, CA, United States  
<sup>23</sup> University of California, Berkeley, CA, United States  
<sup>24</sup> Nuclear Physics Institute of the Czech Academy of Sciences, Husinec-Řež, Czechia  
<sup>25</sup> Suranaree University of Technology, Nakhon Ratchasima, Thailand  
<sup>26</sup> Lund University, Lund, Sweden  
<sup>27</sup> Inha University, Incheon, South Korea  
<sup>28</sup> Synchrotron Light Research Institute, Nakhon Ratchasima, Thailand  
<sup>29</sup> Aarhus University, Aarhus, Denmark  
<sup>30</sup> Università degli studi di Catania, Catania, Italy  
<sup>31</sup> Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Catania, Catania, Italy  
<sup>32</sup> University of Birmingham, Birmingham, United Kingdom  
<sup>33</sup> Pusan National University, Pusan, South Korea  
<sup>34</sup> University of Tokyo, Tokyo, Japan  
<sup>35</sup> Università degli studi di Cagliari, Cagliari, Italy  
<sup>36</sup> Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Cagliari, Cagliari, Italy  
<sup>37</sup> Università degli Studi di Foggia, Foggia, Italy  
<sup>38</sup> University of Kansas, Lawrence, KS, United States  
<sup>39</sup> EPFL, Lausanne, Switzerland  
<sup>40</sup> University of Illinois Chicago, Chicago, IL, United States  
<sup>41</sup> Czech Technical University, Prague, Czechia  
<sup>42</sup> University of Oslo, Oslo, Norway  
<sup>43</sup> Universität Münster, Münster, Germany  
<sup>44</sup> Fudan University, Shanghai, China

## ARTICLE INFO

### Keywords:

Monolithic active pixel sensors  
 Solid state detectors  
 Silicon sensors  
 CMOS stitching

## ABSTRACT

This paper presents the characterisation and testing of the first wafer-scale monolithic stitched sensor (MOSS) prototype developed for the ALICE ITS3 upgrade that is to be installed during the LHC Long Shutdown 3 (2026–2030). The MOSS chip design is driven by the truly cylindrical detector geometry that imposes that each layer is built out of two wafer-sized, bent silicon chips. The stitching technique is employed to fabricate sensors with dimensions of 1.4 cm × 25.9 cm, thinned to 50 μm. The chip architecture, the in-pixel front-end, the laboratory and in-beam characterisation, the susceptibility to single-event effects, and the series testing are discussed. The testing campaign validates the design of a wafer-scale stitched sensor and the performance of the pixel matrix to be within the ITS3 requirements. The MOSS chip demonstrates the feasibility of the ITS3 detector concept and provides insights for further optimisation and development.

## 1. Introduction

Monolithic Active Pixel Sensors (MAPS) are used in high-energy physics experiments as they enable the construction of ultra-thin, large-scale detectors. They were first used in a collider environment in the STAR pixel detector [1]. The ALPIDE sensor implemented in a 180 nm CMOS imaging process [2–4] was used on a much larger scale in the Inner Tracking System 2 (ITS2) of the ALICE experiment at the CERN Large Hadron Collider (LHC) [5]. Since 2020, the ALICE Collaboration [6], in synergy with CERN EP R&D [7], has been carrying out extensive R&D for its ITS3 upgrade [8], which will replace the three inner layers of the ITS2 and which is scheduled for installation during the LHC Long Shutdown 3 (LS3, 2026–2030). The main objective of the ITS3 detector is to reduce the material budget from the current 0.36%  $X_0$  per layer down to about 0.09%  $X_0$  per layer, coming mostly

from the sensors themselves. Furthermore, a new beam pipe with an inner radius of 16 mm and wall-thickness of 500 μm is foreseen, which allows the innermost layer to be installed as close as possible to the beam axis. The required<sup>1</sup> radiation tolerance of the detector is 4 kGy of Total Ionising Dose (TID) and  $4 \times 10^{12}$  1 MeV  $n_{eq}$  cm<sup>-2</sup> of Non-Ionising Energy Loss (NIEL).

To meet these constraints, the ALICE Collaboration has adopted the 65 nm CMOS imaging technology developed by Tower Partners Semiconductor Co. [9]. This technology was validated for ALICE applications and beyond with a set of test structures [10–13]. The 300 nm wafer size and the stitching technique [14,15] allow the fabrication of sensors substantially larger than the typical reticle size (about 3 cm ×

<sup>1</sup> Radiation tolerance requirement has been updated w.r.t. Ref. [8], while the corresponding measurements presented in this paper have been carried out at 2.5 times higher doses compared to those quoted in Ref. [8].

2 cm), with sensor dimensions reaching 10 cm × 27 cm. Thinning to 50 μm thickness allows these wafer-scale sensors to be bent and form a self-supporting, truly cylindrical structure. Dedicated studies have demonstrated that MAPS maintain their performance after bending [16, 17]. Due to a relatively low power consumption of about 40 mW/cm<sup>2</sup>, air cooling becomes feasible, thereby minimising the need for complex cooling and mechanical support structures.

The ALICE ITS3 detector will consist of two truly cylindrical half-barrels. Each half-barrel consists of three sensor layers with a length of 26.6 cm at radii of 19.0 mm, 25.2 mm, and 31.5 mm [8]. Each layer of a half-barrel is formed from a single silicon sensor, built as an array of repeated smaller layout components described in the next chapter. The sensor is physically bent into a half-cylindrical shape and supported by ultra-light carbon foam structures. Electrical interconnection is provided solely through wire-bonding at the ends of the half-cylinder. The production of such large-area pixel sensors is a novel development in the field of high-energy physics experiments, and the prototype sensor discussed in this article aims to assess the feasibility and performance characteristics of this technology and concept.

The Monolithic Stitched Sensor (MOSS) was fabricated in 2023 as part of the Engineering Run 1 (ER1). MOSS measures 1.4 cm × 25.9 cm and served as a demonstrator for the stitching process and the performance of the pixel matrix under ITS3 operating conditions. The development goals of the MOSS design included: (a) gaining experience with the stitching technique to design large sensors that meet the integration requirements of ITS3; (b) studying topologies for distributing power and signals using metal interconnects that span wafer-scale distances; (c) investigating yield and constraints related to Design for Manufacturability (DfM) rules; (d) evaluating the performance of large-area pixel arrays; and (e) analysing noise, power consumption, leakage, and variability in electrical characteristics across very large sensors. A comprehensive characterisation campaign was carried out on the MOSS sensor to gain knowledge and to assess its compliance with ITS3 requirements [8]. The MOSS sensor and the findings from this campaign are presented and discussed in this article.

## 2. MOSS sensor

The MOSS prototype sensor, exploring the application of stitching for the fabrication of wafer-scale MAPS, is schematically shown in Fig. 1. The design is made of three layout components in the design reticle: the Left End-Cap (LEC), the Right End-Cap (REC), and the Repeated Sensor Unit (RSU). The full sensor consists of a linear array of ten abutting RSUs, completed by one LEC and one REC at the respective ends, resulting in a one-dimensional stitched assembly. Each RSU is subdivided into two symmetrical sections referred to as *top and bottom half-units*. Every half-unit contains four *regions*, each containing one pixel matrix and the related biasing, control and readout. The matrices of top regions have 256 × 256 square pixels with a pitch of 22.5 μm. The ones of bottom regions have 320 × 320 square pixels with a pitch of 18.0 μm. The area occupied by pixels within each RSU is about 265 mm<sup>2</sup> out of the total RSU area of 357 mm<sup>2</sup>. In total, the MOSS sensor comprises approximately 6.72 million pixels. The large and small pixel pitches in the top and bottom halves of the sensor allow testing of low and high integration densities, respectively. This feature was intended to evaluate how integration density may affect the functional yield of a large-area stitched sensor.

One 300-mm-diameter wafer comprises six MOSS sensors as shown in Fig. 2. The physical area covered by six MOSS sensors is similar to the area of the outermost ITS3 layers, while five and four adjacent MOSS sensors cover surfaces that are similar to the ones of the middle and innermost layers, respectively. Additionally, each wafer contains 23 *babyMOSS* structures, comprising one RSU, LEC, and REC each. Although smaller than the MOSS, these are effectively fully functional sensors with only one RSU.

The functional block diagram of the MOSS sensor is provided in more detail in Fig. 3, illustrating a bottom half-unit of a RSU, a LEC, and a REC. Each of the 20 half-units (10 top, 10 bottom) can also function autonomously via separate wire bond pads situated along the sensor's long edge, enabling individual operation and characterisation. This modular design of the first stitched prototype enables the distinction between potential faults at the individual half-unit level and those originating from the stitching process. Individual half-units can be independently powered, isolated, and tested even if others malfunction. Signal and power routing between RSUs and the end-caps is achieved by metal-wiring crossing the RSU boundaries, which is produced by stitching. The communication lines, referred to as the stitched communication backbone, are routed over the periphery outside the regions and close to the wire bond pads, as illustrated in Fig. 3. The powering lines are routed as a grid across the entire sensor area and above the pixel arrays.

Two modes of operation are supported: one provides independent control and readout for each half-unit via the bond pads along the sensor's long edge, while the other enables control and readout through the I/Os in the LEC. Power is supplied via the top and bottom bond pads for all RSUs. Because of the fine subdivision of supply nets, the limitations of the MOSS-sensor's metal stack in terms of conductivity, and to prevent excessive voltage drops along the sensor length, power can be supplied via the LEC and REC only to the leftmost (RSU 1) and rightmost (RSU 10), respectively. Each half-unit comprises individual analogue (AVDD, AVSS) and digital (IOVDD, DVDD, DVSS) power nets.<sup>2</sup> The IOVDD net supplies the level-shifting circuitry that translates on-sensor 1.2 V to off-sensor 1.8 V signal levels. Separate global power nets (BBVDD, BBVSS) are available for powering the stitched backbone circuitry, with one net for the top half and a separate one for the bottom half of the sensor. On the LEC, the control and readout I/Os of the top and bottom backbones have dedicated supply nets (BBIOVDD). The sensor substrate biasing net (PSUB) is global for the entire sensor. It is used to reverse bias the charge-collection diodes. Each half-unit has an additional, multiplexed analogue I/O pad on the long edge, used for monitoring and characterising the on-sensor DACs of the four regions within the half-unit.

### 2.1. In-pixel front-end

The in-pixel front end is designed to bias the collection electrode, it amplifies the charge signal and applies a threshold to it to determine whether the pixel was hit or not [18]. The front-end is not a charge amplifier in the classical sense: it profits from the low capacitance at the input node, and hence the relatively large voltage excursion caused by the collected signal charge. The amplification is carried out both by M1 and M2, who each contribute to the signal of the amplifier output. Its operating point is defined by four currents  $I_{\text{bias}}$ ,  $I_{\text{biasn}}$ ,  $I_{\text{reset}}$ , and  $I_{\text{db}}$ , and four voltages  $V_{\text{casb}}$ ,  $V_{\text{casn}}$ ,  $V_{\text{shift}}$ , and  $V_{\text{rcas}}$ , indicated in the simplified schematic in Fig. 4. To meet the stringent power consumption constraints, most transistors operate in weak inversion and with very low biasing currents.<sup>3</sup> A brief overview of the operating principle of the in-pixel front-end is outlined in the following.

Charge collected by the pixel diode causes a voltage drop at the input node of the front-end (gate of M1). The input source-follower transistor M1 provides a high input impedance. The source of M1 reproduces the input voltage and is connected to the gate of the main amplifying transistor M2. The M2 amplifier, M4 cascode, and M9 current sink (together with its related cascode M8) form a folded-cascode amplification stage. M8 stabilises the  $I_{\text{biasn}}$  current-sink branch,

<sup>2</sup> The convention of \*VDD and \*VSS suffixes was chosen, representing the positive supply and corresponding ground, respectively.

<sup>3</sup> The precise values of biasing currents and voltages are discussed in Section 4.3.




Fig. 1. MOSS layout. Left End-Cap (LEC), Repeated Sensor Unit (RSU), Right End-Cap (REC), and Stitching Boundaries (SB) are indicated.



Fig. 2. Processed wafer with 6 numbered MOSS sensors in the centre, and one of the overall 23 babyMOSS sensors labelled near the top of the wafer.

and its gate voltage  $V_{casn}$  is set based on the operating conditions of M7 and M9. M0 provides  $I_{bias}$ , the main current of the front-end. Increasing  $I_{bias}$  increases the gain and decreases equivalent input noise and response time, at the cost of increased power consumption.  $I_{biasn}$  is set significantly lower than  $I_{bias}$ , nominally 1/10 of  $I_{bias}$ , to boost the output impedance of the amplifier and therefore its gain.

The level-shifting transistor M3, steered by  $V_{shift}$ , increases the bias on the collection diode thus reducing its capacitance, up to the limit where M0 is pushed out of saturation, i.e. the main front-end biasing current is reduced. Transistors M5, M6, and M7 constitute the feedback to the input node that sets the input voltage. M5 and M6 together form a cascoded current source. This topology provides a better control of the small  $I_{reset}$  current over large matrices. The M6 gate voltage  $V_{rcas}$  is derived from  $I_{reset}$  in the biasing unit (see Section 2.2), and its value is typically close to the rail voltage (AVDD). The  $I_{reset}$  current must be set larger than the sensor leakage current, but it must be limited to avoid additional shot noise and degradation of the amplifier gain.  $I_{reset}$  is therefore in the pico-ampere range. The  $V_{casb}$  voltage on the gate of M7 and  $I_{reset}$  establish the baseline voltage at the output of the amplifier. Upon charge collection, the rising voltage at the amplifier output reduces the gate-source voltage of M7, turning it off and redirecting the  $I_{reset}$  current to reset the collection diode. A reverse bias of  $-1.2$  V is applied to the sensing diodes via the substrate biasing net PSUB. This voltage is unavoidably applied also to the bulk of the NMOS transistors in the pixels. This reverse substrate/bulk bias for the NMOS transistors in the pixel matrix results in a transistor threshold or VT shift, reducing the current drive capability of the NMOS transistors. This is not a show-stopper, but the effect has to be taken into account in the design of the in-pixel circuitry.

The discrimination of the amplifier output signal is implemented with transistors M10 and M11. In static operation, the  $I_{db}$  current is larger than the standby current in M11, and the discriminator output is thus kept close to the supply voltage. When the amplifier output voltage raises following a hit, the current in M11 increases, and if it exceeds  $I_{db}$ , the drain voltage of M10 drops to almost ground, indicating a hit. The charge threshold, which is set using one value for all pixels of the entire half unit of a RSU, therefore depends on the amplifier gain (influenced

by  $I_{bias}$  and  $I_{reset}$ ), its output baseline defined by  $V_{casb}$  and  $I_{reset}$ , and the discriminator current  $I_{db}$ .

For testing and calibration, charge can be injected at the circuit input by applying a voltage step to the injection capacitance  $C_{inj}$ . This step is produced by switching the capacitor node between  $V_{pulseh}$  and  $V_{pulsel}$ , where  $V_{pulsel}$  corresponds to the potential of the AVSS ground net of the DAC that generates  $V_{pulseh}$  at the periphery of the matrix (see Section 2.2).

The MOSS sensor integrates four variants of the in-pixel front-end in addition to the baseline design, to investigate potential optimisation. One variant uses a larger input transistor M1 to reduce random telegraph and  $1/f$  noise, another features an enlarged discriminating transistor M11 to reduce threshold dispersion, and one includes a larger amplifying transistor M2 to increase gain, albeit with the trade-off of additional input capacitance. The fourth variant employs a modified layout to study the influence of inter-device parasitic capacitances.

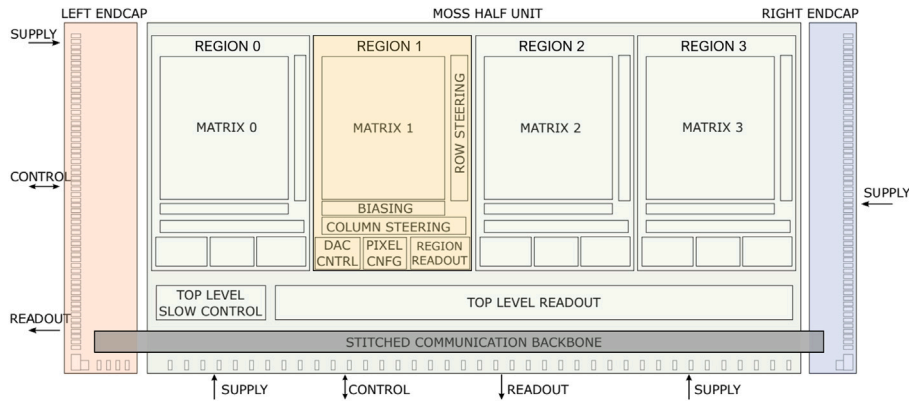
## 2.2. Pixel matrix and biasing

Each pixel in the matrix integrates the previously described analogue section together with a dedicated digital section. The pixels are controlled and read out via a network of digital lines organised into orthogonal buses running along each column and row. These signals are routed and buffered through the ROW STEERING and COLUMN STEERING blocks located at the array periphery, as illustrated in Fig. 3. Configuration functions such as masking, pulsing, and resetting of individual pixels are managed by the PIXEL CNFG block in the peripheral region.

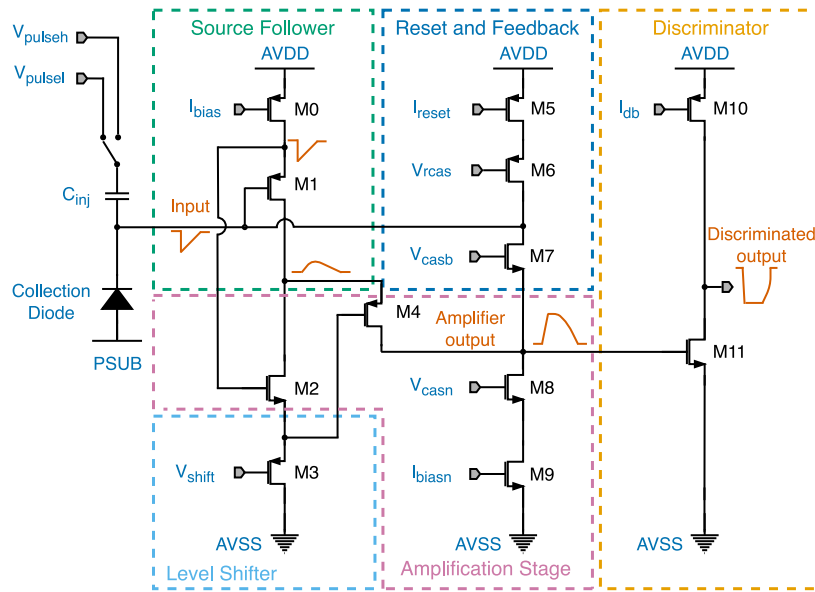
Each pixel includes a readout latch that stores the detection of a hit. A global strobe signal, distributed across all pixels and regions within a half-unit, governs the sampling of the pixel discriminator output into the readout latch. A hit is registered when the strobe signal and the discriminated output of the pixel front-end (see Section 2.1) are simultaneously asserted. The strobe signal is initiated by a user command through the control interface, with its duration and an additional internal delay being configurable. The hit information remains latched in the pixels until a readout command is issued via the control interface, triggering the readout sequence.

Pixel hit readout is managed by the TOP LEVEL READOUT and REGION READOUT blocks. The positions of hit pixels within the array are sequentially encoded through a two-step process: first, scanning rows that contain at least one hit, and then scanning the hit pixels within each selected row. Row and column positions are determined by priority encoders located at the array periphery. For each hit pixel, its row and column addresses are written to a memory buffer in the region readout periphery, after which the corresponding pixel latch is cleared, allowing the encoders to advance to the next hit. This procedure repeats until all pixel latches are cleared. While the readout of the array progresses, the collected hit addresses are assembled into a data frame by the top-level readout unit and transmitted to the data backbone.

On the analogue side, all pixels within a region are connected to ten nets that control the biasing of the front-end and the test charge



**Fig. 3.** MOSS sensor block diagram with one bottom half-unit of a RSU, a LEC, and a REC. Supply, control, and readout lines are schematically indicated. Functional blocks within the half-unit are labelled. The stitched communication backbone spans the full length of the sensor, crossing the stitching boundaries. Region 1 is highlighted, illustrating the contained blocks.



**Fig. 4.** Simplified front-end schematic. Control voltages are applied to the gates of the corresponding transistors, while bias currents are provided via current mirrors (indicated as transistors with a current supplied to the gate, e.g. M0). The orange traces illustrate the characteristic voltage signals at key nodes within the front-end circuit.

injection circuitry. Four currents,  $I_{bias}$ ,  $I_{biasn}$ ,  $I_{reset}$ , and  $I_{db}$ , together with four voltages,  $V_{casb}$ ,  $V_{casn}$ ,  $V_{shift}$ , and  $V_{pulseh}$ , are generated by a set of 8-bit DACs in a biasing unit located in the region-specific periphery (see Fig. 3). The voltage  $V_{rcas}$  is derived from the value of  $I_{reset}$ , while  $V_{pulsel}$  is connected to the ground of the  $V_{pulseh}$  DAC. Each biasing unit includes two internal bandgap circuits that generate reference voltages and currents for the DACs, which can be fine-tuned using dedicated 4-bit DACs.

A scheme of the routing of the biasing and strobe wires is shown in Fig. 5. The digital strobe signal is routed vertically on the side of the pixel matrix and then distributed horizontally across every two pixel rows. Analogue wires distribute the biasing nets horizontally at the bottom of the matrix and then vertically along each column to the pixels. The DACs driving each biasing net are distributed horizontally and spaced by  $450\mu\text{m}$  approximately. Some implications of this layout on the pixel performance are discussed in Section 4.2.

### 2.3. Integration of the pixel sensor within the pixel

The design of the pixel sensor shown as a schematic cross-section in Fig. 6 is similar to that of the earlier prototype, the DPTS sensor [11].

It features a low-doped, deep n-type implant in the pixel area which helps to deplete the pixel over its full area. The optimisation of the process to implement this pixel and the low-doped, deep n-type implant has been described in [19]. The gap between the implants of adjacent pixels enhances the lateral field accelerating the charge collection for charges generated near the pixel edges. To investigate the effect of gap size on charge sharing, a layout variation was implemented. In a subset of wafers, and only for pixels with a  $22.5\mu\text{m}$  pitch, the pixel-to-pixel gap width was increased from the baseline  $2.5\mu\text{m}$  to  $5.0\mu\text{m}$ .

### 3. Test setup

A dedicated test system, shown in Fig. 7, has been developed to functionally characterise the MOSS sensor. The sensor is wire bonded onto a passive printed circuit *carrier board*. Custom-designed *Proximity boards* connect to the carrier board, supplying power, and enabling current and voltage monitoring. Commercial *FPGA boards* are used to operate the Proximity boards and the MOSS sensor. The sensor slow-control and data-readout lines pass through the Proximity boards and connect directly to the FPGAs. Communication between the FPGA boards and a PC is established via a USB3 interface. Custom FPGA




Fig. 5. A schematic representation (not to scale) of one region in the bottom half-unit, illustrating the distribution of the tunable analogue bias and the strobe signal to the pixels.

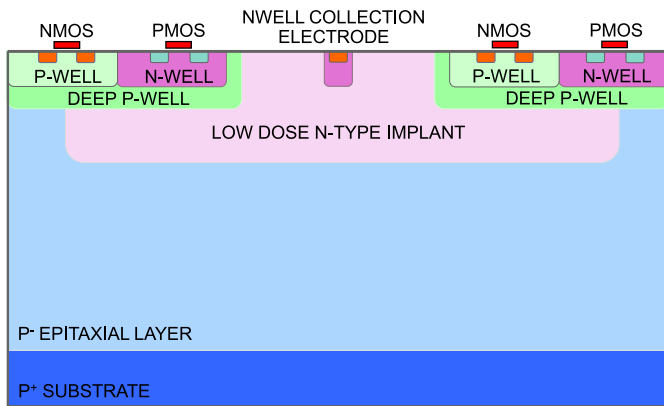


Fig. 6. Schematic cross-section of the pixel sensor.

firmware and Python-based software were developed to steer the setup and the sensor.

The MOSS carrier board has five connectors for FPGA–Proximity board pairs: one is dedicated to the sensor interconnects on the LEC, and four are for the I/Os along the top and bottom long edges of the sensor. One single FPGA–Proximity board pair is sufficient to test a set of five half-units connected to one of the long-edge connectors. The global PSUB voltage is supplied by an external power supply. A full configuration with five pairs of FPGA–Proximity boards connected to all the carrier connectors is used for series testing, and allows for independent half-unit characterisation, simultaneous operation of all the half-units, and testing via the LEC (see Section 5).

### 3.1. Testbeam setup

A testbeam telescope was set up to investigate the MOSS in-beam performance in different operating conditions. The system used a single FPGA–Proximity board pair connected to one of the connectors 1–4 in Fig. 7. Multiple test campaigns took place at the CERN PS and SPS testbeam facilities from July 2023 to April 2025. The results, based on data taken at the CERN PS with a set of representative sensors and a beam of 7 GeV/c negative hadrons, are presented in Sections 4.7 and 4.8.

Fig. 8 shows a schematic diagram of the beam telescope. Six reference planes equipped with ALPIDE sensors [2–4] were used to reconstruct particle tracks, with the MOSS sensor placed as a Device Under Test (DUT) in between the reference arms. An aluminium cooling jig placed on the back of the MOSS carrier was used to keep the sensor at

a constant temperature of 27 °C, corresponding to the operating conditions envisaged for the ITS3 [8]. The carrier board and the cooling jig featured a cutout corresponding to the location of the pixel matrices in order to limit multiple scattering. The coincidence of the amplified and discriminated signals of two scintillators triggered the data acquisition. To reduce the selected events to those containing a single particle track, only trigger signals spaced more than 50  $\mu$ s, with an additional dead time of 100  $\mu$ s after sending the trigger, were accepted.


The EUDAQ [20] and Corryvreckan [21] frameworks were used for the data acquisition and analysis, respectively. Particle tracks were reconstructed by fitting the positions of the clusters found in the reference planes with General Broken Lines [22]. A clean data sample was selected by requiring exactly one reconstructed track per event, a single hit on each reference plane, and a good track fit ( $\chi^2 < 3$ ). Clusters recorded by the DUT were associated with tracks if they fell within a circular acceptance window of 30  $\mu$ m radius, centred on the interpolated intersection of the track with the DUT.

## 4. Pixel matrix characterisation


The pixel matrix was characterised both in the laboratory and with charged-particle beams, with the aim of measuring detection efficiency, fake-hit rate, and spatial resolution before and after irradiation. BabyMOSS sensors were used interchangeably with MOSS sensors, particularly in measurements where size was a limiting factor, such as irradiation campaigns. Under identical conditions, the performance of the two sensor sizes was indistinguishable within the sensor-to-sensor variations, as expected given that they differ only in the number of Repeated Sensor Units (see Section 2). The performance of all front-end variants was studied in detail. However, since the differences are minimal and consistent with expectations (see Section 2.1), only the standard (baseline design) front-end variant will be discussed in the following.

### 4.1. Threshold, noise, and fake-hit rate

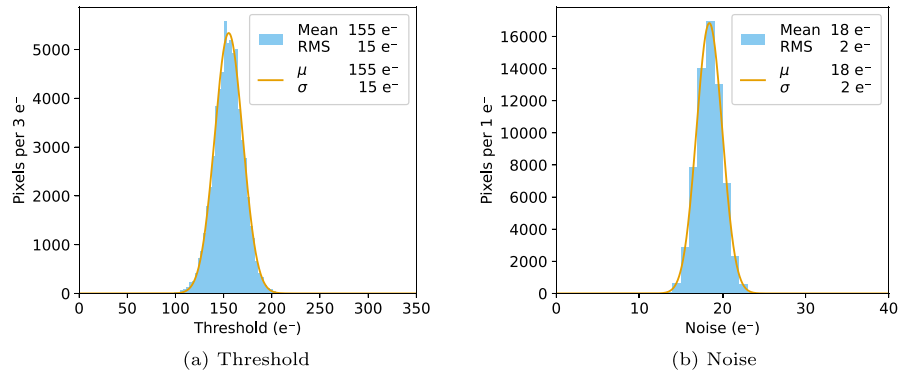
The effective discrimination threshold applied to the charge-collection diode signal is determined by the in-pixel front-end operating point, primarily by the amplifier output baseline and the discriminator current, adjusted via  $V_{casb}$  and  $I_{db}$ , respectively (see Section 2.1). To measure the charge threshold, injection capacitance  $C_{inj}$  (see Fig. 4) is used. The injected charge is incrementally increased by adjusting  $V_{pulseh}$ , and the pixel output state is recorded for each charge level over multiple repeated injections. The hit probability as a function of the injected charge follows the characteristic S-curve response which can be described by a Gaussian error function [11]. The two parameters of



**Fig. 7.** Functional test system. One pair of FPGA board and Proximity board is connected to the carrier connector for the MOSS interfaces on the LEC (5). The additional four connectors 1–4 can connect to FPGA and Proximity boards for powering and characterisation through the long-edge interfaces of the sensor.



**Fig. 8.** Schematic representation of the setup used for in-beam measurement with ionising particles (not to scale). Six ALPIDE sensors are used as reference planes. A Device Under Test (DUT) is placed between two reference arms. The coincidence of the two scintillator signals is used for triggering.



**Fig. 9.** Threshold and noise distributions of a reference region ( $22.5\mu\text{m}$  pixel pitch,  $2.5\mu\text{m}$  gap) under typical operating conditions. Both distributions are well described by a Gaussian fit. The threshold dispersion is comparable to the average pixel noise.

this function, the mean and standard deviation, correspond to the pixel threshold and noise.

Fig. 9 shows the distributions of threshold and noise values measured for all pixels in a region under typical operating conditions (see Section 4.3). The threshold dispersion, originating from fabrication-induced variations in the in-pixel circuitry, is significant, with an RMS of about 10% of the mean threshold value, yet it remains consistent with the earlier prototypes [11]. The measured noise includes contributions from the in-pixel front-end and the sensing node, and its average is comparable to the threshold dispersion.

The threshold dispersion is mostly uncorrelated with pixel position, as seen in Fig. 10(a) where the threshold data is mapped to pixel positions in the matrix. Eleven distinct vertical lines are observed in Figs. 10(a) and 10(b), and correspond to columns in which the front-end input is coupled to the digital signals steering the charge-injection circuit [23]. Therefore, this is an artefact of the threshold measurement, and is not affecting the sensing performance of the pixels.

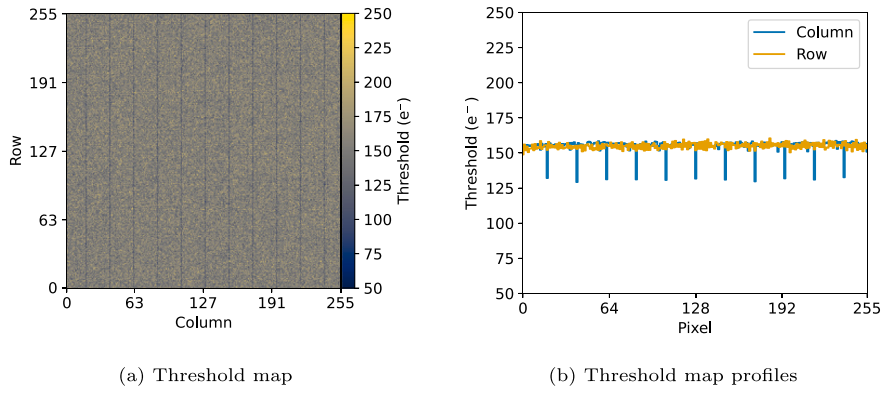
The threshold exhibits a temperature dependence. In measurements with a non-irradiated sensor, a linear decrease of approximately  $2\text{e}^- \text{C}^{-1}$  was observed over the range 20–35 °C, consistent with simulations of the in-pixel front-end.

As a result of in-pixel discrimination, noise ultimately manifests as fake hits. The fake-hit rate quantifies how often the sensor registers hits in the absence of particles or charge injection. It accounts for all

previously discussed noise sources, including random telegraph noise, as well as systematic effects such as coupling to signal distribution nets (see Section 4.2). The measurement of the fake-hit rate is performed by issuing several tens of thousands of triggers and recording the number of hits per pixel per unit time. Pixels registering a fake hit in more than 1% of issued triggers are classified as noisy, excluded from the fake-hit rate calculation, and masked from further analysis. In all but a small fraction of cases (see Section 5.5), there are fewer than three such pixels per region, comprising 65 536 and 102 400 pixels in top and bottom half-units, respectively. The numbers of pixels that require masking increases after exposing sensors to ionising or non-ionising radiation (see Section 4.7).

#### 4.2. Strobe effect on threshold and fake-hit rate

A parasitic capacitive coupling was identified at the crossings of the analogue-bias and strobe-distribution lines (see Fig. 5). Each crossing introduces a capacitance that was found to cumulatively have a measurable impact on the analogue front-end performance. Consequently, the falling edge of the strobe signal (see Section 2.2, active when low) introduces perturbation to all analogue biases. At first, the distributed biasing voltage drops rapidly due to coupling to the strobe signal. The biasing unit then tries to compensate this drop, i.e. to restore the voltage to its original level. This dynamic response depends on spatial



**Fig. 10.** Threshold map of a reference region ( $22.5\ \mu\text{m}$  pixel pitch,  $2.5\ \mu\text{m}$  gap) and its profiles along column and row direction. The threshold dispersion is uniform across the matrix, except of the eleven columns where front-end input is coupled to digital signals steering the charge-injection circuit [23].


factors, specifically the column number, given by the location of the biasing units (see Fig. 5). The driving strength varies with the distance from the bias origin due to resistance and parasitic capacitance along the path, leading to differences in the amplitude and duration of the resulting perturbations. A perturbation of a bias connected to the front-end input node, i.e.  $V_{\text{shift}}$  and  $I_{\text{bias}}$  biases (see Section 2.1), produces a signature on the input line similar to that of an injected charge. As a result, the amplifier output resembles that of a small injected charge, with a typical front-end peak time  $\mathcal{O}(1\ \mu\text{s})$  and a recovery time of approximately  $10\ \mu\text{s}$ , driven by the recovery time of the biases.

The effect of the perturbation on the threshold is visible in Fig. 11, showing the average threshold variation as a function of column, at different delays between the strobe signal and the charge injection. To isolate the effect of the perturbation, the threshold measured at a delay of  $8.8\ \mu\text{s}$  when the system is assumed to have recovered from the perturbation, is subtracted from the measurements at shorter delays.

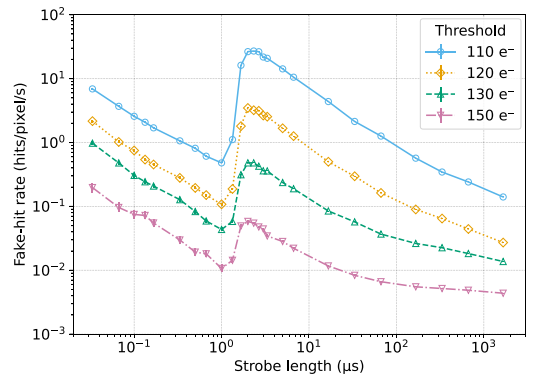
The profiles, especially looking at short delays (e.g.  $1.3\ \mu\text{s}$ ), reveal characteristic symmetries at the columns associated with the biasing unit connection to the matrix of  $V_{\text{shift}}$  (column 150) and  $I_{\text{bias}}$  (column 55). Furthermore, the threshold shows different patterns at different delays, reflecting the temporal evolution of the perturbation. The difference to the reference is largest at short delays (close to the start of the strobe) and smallest at large delays indicating the recovery from the perturbation. The reference threshold at  $8.8\ \mu\text{s}$  is chosen because no remaining column dependence or time dependence is observed beyond this point, which also agrees with the expected bias recovery time from simulation.

Fig. 12 shows the dependence of the fake-hit rate on the strobe signal, illustrating the cumulative effect of perturbations over time on the probability of observing a fake hit. For random noise, the number of fake hits per unit time is expected to be independent of strobe length. In practice, however, two systematic effects modify the measured fake-hit rate. At very short strobe lengths, the measurement is dominated by the few noisiest, continuously active pixels. At very long strobe lengths, the fake-hit rate is underestimated because the in-pixel latch can be asserted only once, limiting the contribution of the noisiest pixels to a single fake hit. Consequently, the dependence of the fake-hit rate on strobe duration is expected to decrease monotonically. Instead, a sudden increase is observed approximately  $2\ \mu\text{s}$  after the strobe is asserted. This behaviour is consistent with the timescale for amplifying a small charge (noise injection) at the input when the strobe begins.

This finding leads to the conclusion to operate the sensor during the pixel matrix characterisation with a short strobe length of  $0.6\ \mu\text{s}$ . It mitigates the impact of the perturbation by closing the strobe window before the perturbation has a significant effect on the fake-hit rate and ensures that the fake-hit rate is not underestimated, as it would be the case for a longer strobe length ( $>100\ \mu\text{s}$ ). The final ITS3 sensor will use a layout without strobe-bias crossings and edge-based latching of a hit, ensuring a fake-hit rate independent of acquisition time.



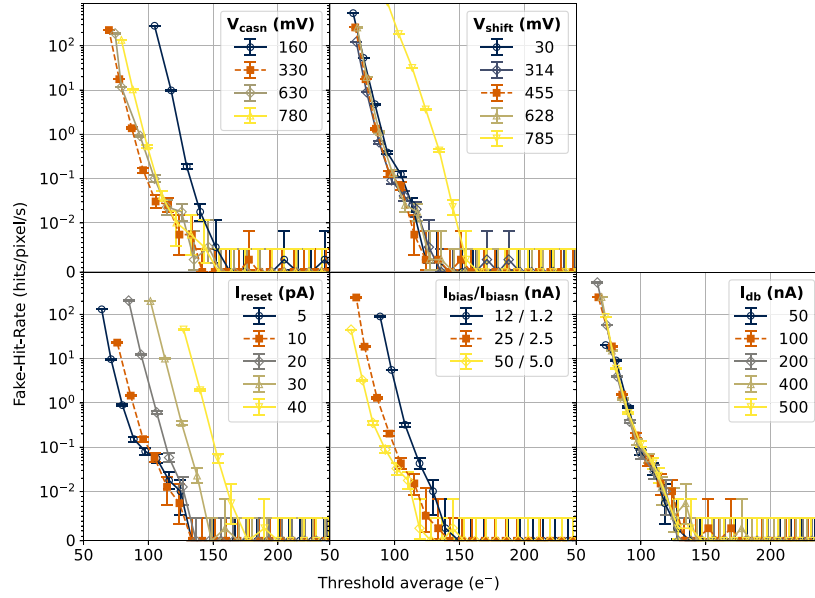
**Fig. 11.** Column profile of threshold variation measured at different time delays after the strobe is asserted. To isolate the effect of the perturbation, the threshold measured at a delay of  $8.8\ \mu\text{s}$  when the system is assumed to have recovered from the perturbation, is subtracted from the measurements at shorter delays. The small peak substructure arises from the different coupling with the steering signals discussed in Section 4.1.



**Fig. 12.** Fake-hit rate as a function of strobe length at different thresholds. The strobe perturbation causes a sharp increase in the fake-hit rate at a strobe length of about  $2\ \mu\text{s}$ . The apparent rise at very short strobe lengths and the apparent decrease at very long strobe lengths are misleading due to the construction of the fake-hit rate observable.

#### 4.3. Working point validation of the in-pixel front-end

The front-end operating point, as a function of bias parameters, was extensively studied to validate the design simulations and optimise the signal-to-noise ratio. Fig. 13 illustrates the relation of the fake-hit rate and the threshold as individual front-end parameters are varied. The parameter  $V_{\text{casb}}$  serves as a linear handle to compensate for threshold



**Fig. 13.** Fake-hit rate as a function of the threshold average per region (set via  $V_{casb}$ ) for different front-end parameter variations. The plots, arranged from top left to bottom right, show the effect of varying  $V_{casb}$ ,  $V_{shift}$ ,  $I_{reset}$ ,  $I_{bias}$  (together with  $I_{biasn}$  at a ratio 10:1), and  $I_{db}$ . The error bars indicate the statistical uncertainty of the fake-hit rate measurement. The nominal operating point extracted from simulations (red dashed lines) was confirmed to yield optimal fake-hit rate performance for a given threshold.

shifts induced by other front-end settings, allowing consistent threshold tuning without affecting overall front-end operation (see Section 2.1).

Starting with the parameter  $I_{reset}$ , a reduction in the fake-hit rate at a given threshold is observed as  $I_{reset}$  decreases. This behaviour is consistent with an increased front-end gain, attributed to a higher resistance in the reset and feedback branch, and with a lower shot noise from a reduced current injected into the collection electrode. is further supported by the expectation that a lower  $I_{reset}$  reduces shot noise at the collection electrode. Since  $I_{reset}$  also influences the diode biasing, it can be adjusted to compensate for the increased leakage current in sensors exposed to non-ionising radiation. To maintain uniformity in the biasing conditions across different irradiation levels, without compromising the noise performance, a nominal setting of 10 pA is adopted.

An increase in the front-end bias current,  $I_{bias}$ , reduces the fake-hit rate at a given threshold. This is consistent with an increase in the overall front-end voltage gain, and a reduction in the thermal noise of the input transistor M1. Since  $I_{bias}$  is the primary contributor to the power consumption of the in-pixel front-end, a value of 25 nA is chosen as a trade-off between the ITS3 power budget constraints [8] and fake-hit rate performance.

Variations of discriminator current  $I_{db}$ , cascode voltage  $V_{casb}$ , and level-shifting voltage  $V_{shift}$  around their nominal values (100 nA, 330 mV, and 460 mV, respectively) do not lead to significant changes in the observed performance. However, at elevated values of  $V_{shift}$ , the M0 transistor in the front-end circuit (see Fig. 4) enters the ohmic region, reducing the front-end gain. Consequently, the fake-hit rate increases for a given threshold. A similar effect is observed for low values of  $V_{casb}$ , where the M9 transistor also transitions into the ohmic region, again leading to degraded performance. While choosing a lower  $I_{db}$  current can reduce dynamic power consumption, it requires a corresponding decrease in  $V_{casb}$  to compensate for the associated threshold reduction (see Section 2.1). However, this limits the available adjustment range of  $V_{casb}$ , particularly for sensors affected by threshold shifts due to ionising radiation, hence, the nominal value was kept at 100 nA.

Overall, the nominal operating point extracted from simulations was confirmed to yield optimal fake-hit rate performance for a given threshold and is adopted for all subsequent measurements and analyses presented in this work.

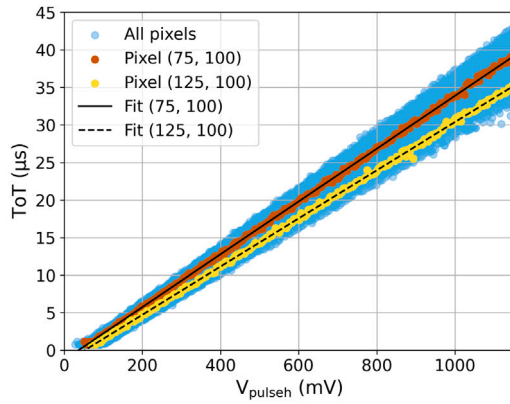
#### 4.4. Time-over-threshold measurement

The Time-over-Threshold (ToT) quantifies the duration for which the signal at the discriminator input exceeds the threshold, thereby encoding the charge information in time. The ToT of a MOSS front-end signal can be assessed in a specific readout configuration. Here, the strobe signal is set to last much longer than a typical front-end pulse. When charge is collected, the discriminator output activates, leading the hit to be stored in the corresponding in-pixel latch (see Section 2.2). This triggers a global signal indicating data presence in the pixel matrix. As soon as this signal is propagated to the test system, a readout command is sent. This initiates the readout of the pixel address(es) by the region periphery and resets the in-pixel latch. As long as the strobe and discriminator output remain active, the pixel-hit latch will immediately reassert and the region readout will read out the same pixel address repeatedly. The number of generated pixel addresses is then proportional to the duration of the front-end signal, thereby providing ToT measurement.

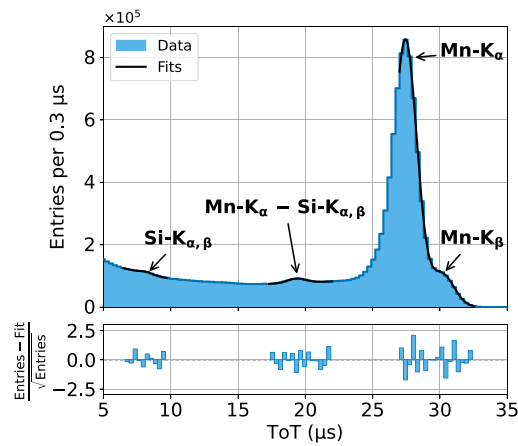
The ToT response, similar to the threshold (see Section 4.1), exhibits pixel-to-pixel variations due to fabrication-related non-uniformities. Fig. 14 presents the ToT measured across a number of pixels within a selected region as a function of the pulsing voltage  $V_{pulseh}$ , modulating the injected charge. For each pixel, the ToT response increases linearly with  $V_{pulseh}$ , but the slope and intercept of this relationship vary across pixels. To ensure a uniform matrix response, each pixel's ToT response is individually calibrated by fitting a straight line. This calibration is applied to all the following measurements.

#### 4.5. Measurements with X-rays

Fig. 15 shows the sensor response to X-ray emissions from a  $^{55}\text{Fe}$ , plotted as a ToT spectrum of events where the charge is detected by a single pixel. Besides the primary 5.9 keV Mn- $K_{\alpha}$  and 6.5 keV Mn- $K_{\beta}$  emissions from the  $^{55}\text{Fe}$  decay, spectral features arising from secondary interactions within the silicon sensor are also resolved: namely, the silicon fluorescence line Si- $K_{\alpha}$  and the escape peak Mn- $K_{\alpha,\beta}$  - Si- $K_{\alpha,\beta}$ . The Mn- $K_{\alpha}$  and Mn- $K_{\beta}$  peaks are fitted with a sum of two Gaussians. The Si- $K_{\alpha}$  and Mn- $K_{\alpha,\beta}$  - Si- $K_{\alpha,\beta}$  peaks are fitted with a Gaussian added to a linear background. Based on the fit to the dominant Mn- $K_{\alpha}$  emission



**Fig. 14.** ToT as a function of the pulsing voltage  $V_{\text{pulseh}}$ . Two arbitrary pixels with distinct slopes are highlighted to better illustrate pixel-to-pixel variations. Pixels are fitted with a linear function, which is used to calibrate the pixel ToT response.



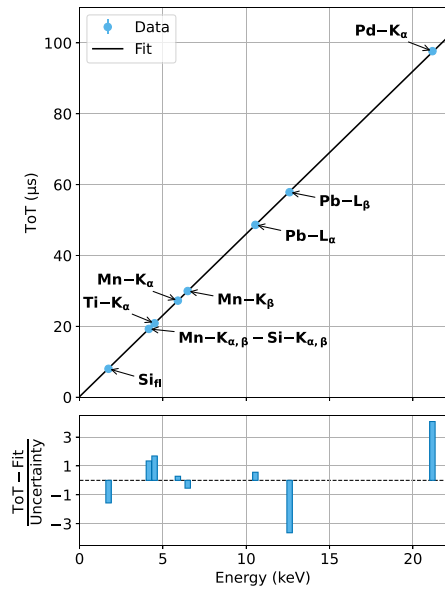
**Fig. 15.** Single-pixel cluster ToT spectrum measured with an  $^{55}\text{Fe}$  source. The primary  $\text{Mn-K}_{\alpha}$  and  $\text{Mn-K}_{\beta}$  X-ray emissions are resolved, as well as the secondary  $\text{Si-K}_{\alpha}$  and  $\text{Mn-K}_{\alpha,\beta} - \text{Si-K}_{\alpha,\beta}$  peaks. Fit residuals shown in the bottom plot are within 2.5 standard deviations.

line, the energy resolution is determined to be  $\text{FWHM}/\text{Mean} = (7.3 \pm 0.2)\%$ , consistent with values reported for the previous prototype [11].

X-ray fluorescence emissions from titanium, lead, and palladium were also measured and fitted. The measurement took place at the OptimaTo laboratory [24], located at Elettra Sincrotrone Trieste [25]. The experimental conditions were equivalent to those described in Ref. [12]. By comparing the fitted peak positions with literature energy values, the ToT response was confirmed to be linear across the energy range from 1.7 keV to 21.2 keV, as shown in Fig. 16.

#### 4.6. Injection capacitance calibration

Due to fabrication process variations, the injection capacitance (denoted as  $C_{\text{inj}}$  in Fig. 4) must be calibrated to enable accurate comparison of threshold values across different pixels and regions. Given the confirmed linearity of the ToT response with both injected charge and deposited energy (see Figs. 14 and 16, respectively), the injection capacitance can be determined by measuring the position of the  $\text{Mn-K}_{\alpha}$  peak (see Fig. 15) and calculating the corresponding pulsing voltage,  $V_{\text{pulseh}}$  (see Fig. 14). The injection capacitance is then obtained as  $C_{\text{inj}} = Q_{\text{Mn-K}_{\alpha}}/V_{\text{pulseh}}$ , where  $Q_{\text{Mn-K}_{\alpha}}$  is the charge deposited by photoelectric absorption of a  $\text{Mn-K}_{\alpha}$  emission.



**Fig. 16.** Linearity of the in-pixel front-end response to collected charge. The observed emission line positions are plotted against corresponding literature values, with a linear fit to the data shown in black.

Figs. 17 and 18 show the distributions of measured injection capacitance values for a subset of pixels<sup>4</sup> in a non-irradiated region and the average values across regions subjected to different irradiation levels, respectively. With an RMS of approximately 1% of the mean value, the pixel-to-pixel capacitance spread can be assumed to contribute only marginally to the measured threshold spread (see Section 4.1). However, the average capacitance dispersion across different regions cannot be attributed to statistical pixel-to-pixel variation alone and instead indicates systematic region-to-region differences.

To mitigate this spread, the injection capacitance values for all regions presented in this work are calibrated using the described procedure. The average capacitance of  $(272 \pm 2)$  aF, measured across different regions and irradiation levels, is reasonably close to the design value<sup>5</sup> of  $(258 \pm 22)$  aF. The quoted uncertainty corresponds to the statistical uncertainty on the mean across regions, while the design uncertainty reflects variations from parasitic extraction corners. As expected, no impact of irradiation is observed.

#### 4.7. Detection efficiency and fake-hit rate

In-beam measurements were conducted to evaluate detection efficiency at different thresholds, using the setup described in Section 3.1. The uncertainty of the threshold is calculated by accounting for the statistical uncertainty of the threshold measurement in  $V_{\text{pulseh}}$  DACs (see Section 4.1) and the statistical uncertainty from the conversion into electrons based on the calibration of the pulsing capacitance (see Section 4.4). The uncertainty of the detection efficiency and fake-hit rate is evaluated using a Clopper-Pearson interval with a 66.3% confidence level.

The target performance for ITS3 sensors is a detection efficiency higher than 99%, with a fake-hit rate lower than 0.1 hits/pixel/s [8]. This performance must be maintained after the radiation doses expected during ITS3 operation, which, including a safety factor, are<sup>1</sup>

<sup>4</sup> To avoid artefacts from non-idealities in the injection circuit [23], only the central quarter of the region is shown, excluding also affected columns.

<sup>5</sup> The design value of the injection capacitance and the corresponding uncertainty are determined from the typical and extreme parasitic extraction corners.




Fig. 17. Distribution of injection capacitances measured for a subset of pixels in a non-irradiated region. The orange line represents a Gaussian fit. The measured spread is approximately 1% of the mean value.

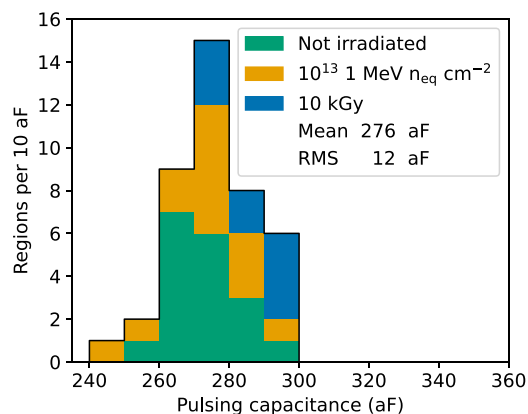


Fig. 18. Distribution of the average injection capacitance values for non-irradiated, 10 kGy, and  $10^{13}$   $1 \text{ MeV } n_{\text{eq}} \text{ cm}^{-2}$  irradiated regions (stacked). The spread of the distributions reflects region-to-region variations, while no impact of the irradiation is visible.

of 4 kGy Total Ionising Dose (TID) and  $4 \times 10^{12}$   $1 \text{ MeV } n_{\text{eq}} \text{ cm}^{-2}$  Non-Ionising Energy Loss (NIEL).

Fig. 19 compares the detection efficiency and the fake-hit rate as a function of the average sensor threshold for two pixel pitches, 22.5  $\mu\text{m}$  and 18.0  $\mu\text{m}$ , and shows the effect of increasing the size of the deep implant gap (see Section 2.2) from 2.5  $\mu\text{m}$  to 5.0  $\mu\text{m}$  for the 22.5  $\mu\text{m}$  pitch. A larger pixel pitch is observed to increase detection efficiency. This is consistent with the fact that larger pixels have proportionally less border area than smaller ones; hits in the border region are statistically more likely to fall below the threshold due to charge sharing between multiple pixels and energy straggling. This effect was previously observed in small-scale prototypes [10]. The pixel matrix with 22.5  $\mu\text{m}$  pixel pitch shows a slightly lower efficiency for a larger gap, in agreement with the expectation that when the gap size is increased, the electric field at the border is decreased, increasing the charge sharing in the corresponding region [26].

The fake-hit rate is higher for the larger pixel pitch. The shot noise from leakage current, which increases with pixel volume, is insufficient to account for this.<sup>6</sup> This is consistent with the expectation that leakage current is not the dominant contributor in non-irradiated samples. In the smaller-pitch matrix, the higher pixel count increases the capacitive

load on the biasing circuit and the number of bias-strobe crossings. This modulates the strobe-induced perturbation (see Section 4.2), which can explain the different fake-hit rates. Overall, a larger pixel pitch achieves the target performance over a wider threshold range.

Figs. 20(a) and 20(b) show detection efficiency and fake-hit rate for 22.5  $\mu\text{m}$  and 18.0  $\mu\text{m}$  pitches with 2.5  $\mu\text{m}$  gap, as a function of the threshold, for different irradiation levels: non-irradiated, an ionising radiation (TID) dose<sup>7</sup> of 10 kGy, and a non-ionising radiation (NIEL) fluence<sup>8</sup> of  $10^{13}$   $1 \text{ MeV } n_{\text{eq}} \text{ cm}^{-2}$ . The use of Xray photons and neutrons for TID and NIEL irradiation, respectively, instead of charged hadrons was motivated in order to study the impact of both types of irradiation separately. After irradiation, the operational margin, defined as the threshold range where ITS3 requirements are met, is reduced as expected. Whereas, for example, for the non-irradiated pixel matrix with pitch 22.5  $\mu\text{m}$  shown in Fig. 19, the operational margin is between  $110 e^-$  and  $200 e^-$  for a total range of about  $90 e^-$ , it reduces to less than  $50 e^-$  at the radiation levels used in this study (see Fig. 20(a)). For the pixel matrix with a 18.0  $\mu\text{m}$  pitch, the observed operational margin of less than  $20 e^-$  (see Fig. 20(b)) can practically be considered as disappeared.

The sample irradiated with non-ionising radiation shows an increased fake-hit rate, expected due to higher sensor leakage from irradiation damage and the associated rise in shot noise. Furthermore, the fake-hit rate plateau at higher thresholds is consistent with the residual radioactivity of the sensors after irradiation. The TID-irradiated sample also shows a higher fake-hit rate, as expected from ionising radiation impacting transistor performance and thus the front-end signal-to-noise ratio. Both pixel pitches exhibit comparable trends across different irradiation levels.

#### 4.8. Spatial resolution and average cluster size

A particle hit on the sensor may cause one or more pixels to register a signal above threshold. Adjacent pixels registering a hit form a cluster, whose size corresponds to the number of pixels above threshold. To calculate the spatial resolution, the RMS of the residual (the distances between the track intercept on the device under test and the cluster centre of mass<sup>9</sup>) is computed in both column and row directions. The spatial resolution for the two directions is then obtained by quadratically subtracting the estimated telescope tracking resolution at the DUT position, which is about 2  $\mu\text{m}$ . The spatial resolution referred to in the remainder of this paper is the average of the resolutions along the column and row directions. The uncertainty of the spatial resolution is derived from the statistical uncertainty of the RMS of the residual. The error on the cluster size is a statistical error on the mean.

Fig. 21 compares spatial resolution and average cluster size for 22.5  $\mu\text{m}$  and 18.0  $\mu\text{m}$  pixel pitches and shows the effect of increasing the gap size for the 22.5  $\mu\text{m}$  pitch. The ITS3 spatial-resolution target value of 5  $\mu\text{m}$  [8] is shown as a dashed line. The dotted lines represent the “hit/no-hit resolution”, expected if the deposited charge were collected by a single pixel. The measured spatial resolution is consistently better than this limit, showing a decreasing trend towards lower thresholds, where larger cluster sizes are observed. Hits near pixel borders, where

<sup>7</sup> The TID irradiation was done using the CERN Xray machine. The sensors were tested within a week after receiving the ionising radiation dose and again after six months of annealing at room temperature. Comparable results were obtained. The results after the annealing are shown, as they were obtained under the same conditions as the other results in this work.

<sup>8</sup> The NIEL irradiation was done using neutrons from the JSI TRIGA Mark II reactor in Ljubljana. The sensor irradiated with a NIEL fluence was kept at  $-20^\circ\text{C}$  between the irradiation and the testbeam in order to avoid annealing, as the process of procuring new samples required a significant time and labour investment.

<sup>9</sup> With equal weights on each pixel since no analogue information is available.

<sup>6</sup> Estimated using the thermal model of the fake-hit rate described in Ref. [12].



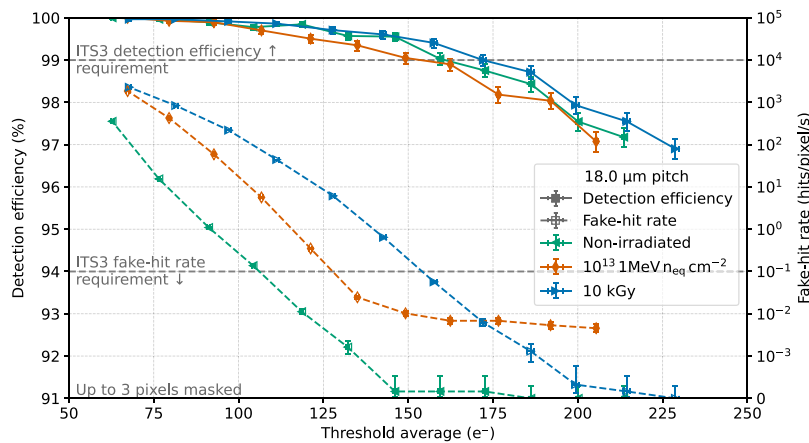


Fig. 19. Detection efficiency (filled symbols, solid lines) and fake-hit rate (open symbols, dashed lines) as a function of the average threshold, comparing a non-irradiated 22.5 μm pitch, 18.0 μm pitch, and 22.5 μm pitch with an increased gap size from the 2.5 μm to 5.0 μm.



(a) Pixel matrix with pitch 22.5 μm irradiated to different levels.



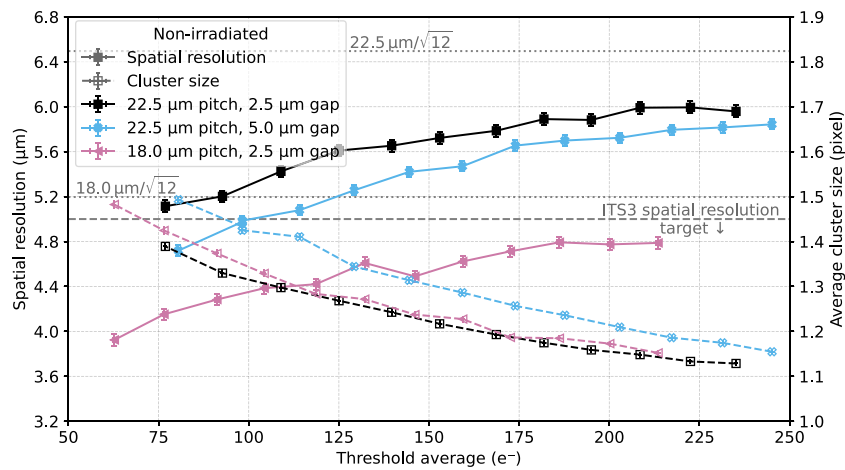
(b) Pixel matrix with pitch 18.0 μm irradiated to different levels.

Fig. 20. Detection efficiency (filled symbols, solid lines) and fake-hit rate (open symbols, dashed lines) as a function of average threshold for different irradiation levels: non-irradiated, an ionising radiation dose of 10 kGy and a non-ionising radiation fluence of  $10^{13}$  1 MeV  $n_{eq}$   $cm^{-2}$ .

charge is shared between multiple pixels, often fall below threshold and are not registered, reducing the average cluster size at higher thresholds. Across all thresholds, the 18.0 μm pitch provides better spatial resolution due to the finer size. At a threshold of 160  $e^-$  (the lowest threshold within the operational range after irradiation, see Fig. 20), the 18.0 μm pitch achieves a spatial resolution of about 4.5 μm,

satisfying the ITS3 requirements. In contrast, the 22.5 μm pitch reaches only about 5.7 μm.

Comparing the pixels with a 2.5 μm and a 5 μm gap sizes shows that a larger gap increases the average cluster size, especially at low thresholds. At higher thresholds, the cluster size tends to converge to that of the 2.5 μm gap, as the threshold becomes too high to detect



**Fig. 21.** Spatial resolution (filled symbols, solid lines) and average cluster size (open symbols, dashed lines) as a function of the average threshold, comparing a non-irradiated 22.5  $\mu\text{m}$  pitch, 18.0  $\mu\text{m}$  pitch and 22.5  $\mu\text{m}$  pitch with an increased gap size from the 2.5  $\mu\text{m}$  to 5  $\mu\text{m}$ .

shared signals and border effects contribute less to the average cluster size. The increased cluster size leads to an improvement of spatial resolution by 0.1–0.4  $\mu\text{m}$  at thresholds below 160  $e^-$ , resulting in about 5.4  $\mu\text{m}$  at a threshold of 160  $e^-$ . While a 5  $\mu\text{m}$  gap shows potential to enhance spatial resolution, the 22.5  $\mu\text{m}$  pixel pitch remains above the ITS3 target value.

Figs. 22(a) and 22(b) show the spatial resolution and average cluster size for 22.5  $\mu\text{m}$  and 18.0  $\mu\text{m}$  pitches with a 2.5  $\mu\text{m}$  gap, as functions of the average threshold for different irradiation levels: non-irradiated, an ionising radiation dose<sup>7</sup> of 10 kGy, and a non-ionising radiation fluence<sup>8</sup> of  $10^{13}$  1 MeV  $n_{\text{eq}} \text{cm}^{-2}$ . After non-ionising irradiation, the average cluster size decreases at low thresholds, causing a corresponding degradation in spatial resolution of up to 0.3  $\mu\text{m}$ . At higher thresholds, no significant difference is observed between non-irradiated and NIEL-irradiated sensors, as signals from charge sharing at pixel borders already fall below threshold even without irradiation. Since ionising radiation affects mostly the in-pixel front-end, shifting the threshold and noise without impacting charge collection, cluster size or spatial resolution shows no significant effect caused by the delivered ionising dose.

## 5. Stitched design validation

To validate the stitched sensor design and evaluate production yield, a test campaign was carried out on 82 non-irradiated MOSS sensors from 14 wafers. This section summarises the results in terms of the yield of different sensor components and the overall yield of wafer-scale sensors. The tests were performed by powering, controlling, and reading out the MOSS sensor via the long edge, one half-unit at a time (see Section 2). Half-units that passed these tests were subsequently retested via the LEC (short edge) to verify communication across the stitching boundary. The test procedure begins by powering the half-units, followed by verifying the functionality of digital and analogue periphery, and concludes with characterising the pixel matrix. Testing was conducted in a laboratory with centrally controlled air temperature. The sensor temperature was continuously monitored but not actively controlled. The 4  $^{\circ}\text{C}$  maximum temperature difference between sensors recorded across tests is not considered to have a significant impact on the interpretation of the following results (see Section 4.1).

### 5.1. Powering

During the powering test, each net of a half-unit was activated sequentially, followed by the application of a clock signal, a reset procedure, and the configuration of the nominal operating point. The spatial distribution of successfully powered half-units is visualised as

a wafer map in Fig. 23. A radial gradient is observed in the powering yield, with reduced functionality concentrated near the wafer centre. The incidence of these faults varies significantly across wafers, as will be shown in Section 5.7.

These patterns suggest underlying manufacturing issues. The MOSS design introduced a novel metal-stack configuration, implemented for the first time during the ER1 production run. This configuration was specifically customised by the foundry in a collaborative effort to meet the requirements of the ITS3 sensor development project. A detailed investigation, reported in Ref. [27], correlated the failures with features of the new metal stack and facilitated the implementation of corrective measures by the foundry. The findings also provided valuable insights for future design iterations aimed at mitigating similar risks.


The current on each power net is measured after all powerable half-units and the LEC are switched on and configured. The distributions of current values are shown in Fig. 24. These data indicate the typical operating conditions of the sensor and provide a baseline for understanding sensor-to-sensor performance differences (see Section 5.5). Power nets AVDD, DVDD, and IOVDD are measured per half-unit. Power nets BBVDD and BBIOVDD are measured for the top and bottom half of each sensor. The PSUB current is global and measured per MOSS sensor.

The larger current spread in the analogue net (AVDD) is attributed to differing consumption between the top and bottom half-units (with lower and higher pixel density, respectively), and to varying region front-end implementations at default biasing conditions. Both digital level-shifting nets, IOVDD and BBIOVDD, show the expected low current consumption when no signals are being transferred. All cases, both within the plotted ranges and in the overflow bin, remain fully operational with the increased current attributable to the aforementioned metal-stack issue. Similarly, cases on the opposite side of the AVDD and DVDD spectra, with notably lower current, occur when current from one net is sunk by the other. The large spread in substrate current (PSUB) is also attributed to the same issue, and, additionally, to protection diode structures between power nets.

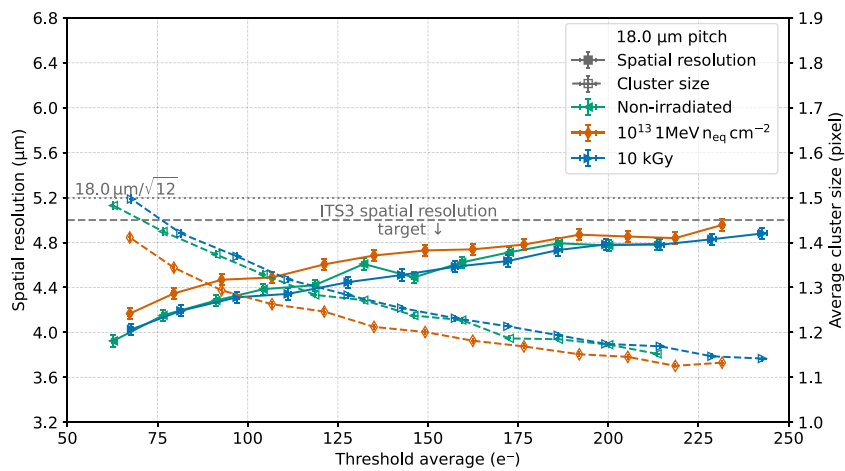
### 5.2. Digital periphery tests

Successfully powered half-units are tested for proper functionality in the digital periphery, including slow control communication and register access. The test involves writing and reading back different patterns to all registers.<sup>10</sup> In total, 0.1% of regions exhibited one or more

<sup>10</sup> Each half-unit contains 402 registers. Registers that could place the sensor in an unstable condition, for example by setting the pixel front-end to extremely high power consumption, were only read out and not written to.



(a) 22.5 μm pixel pitch.



(b) 18.0 μm pixel pitch.

Fig. 22. Spatial resolution (filled symbols, solid lines) and average cluster size (open symbols, dashed lines) as a function of the average threshold for different irradiation levels: non-irradiated, an ionising radiation dose of 10 kGy and a non-ionising radiation fluence of  $10^{13}$  1 MeV  $n_{eq}$   $cm^{-2}$ .

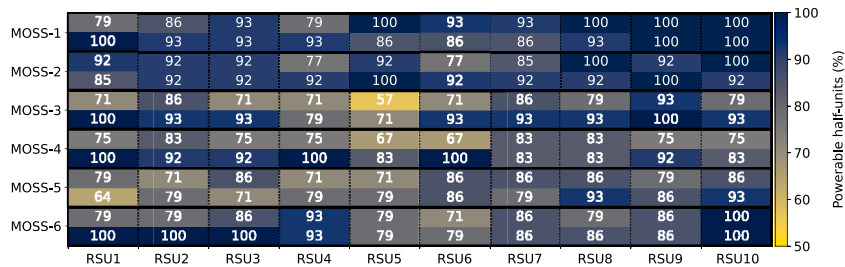



Fig. 23. Fraction of half-units that can be powered, as a function of the location on the six MOSS sensors on each wafer. A radial gradient is observed, with reduced yield concentrated near the wafer centre. A detailed investigation correlated the failures with features of the new metal stack [27].

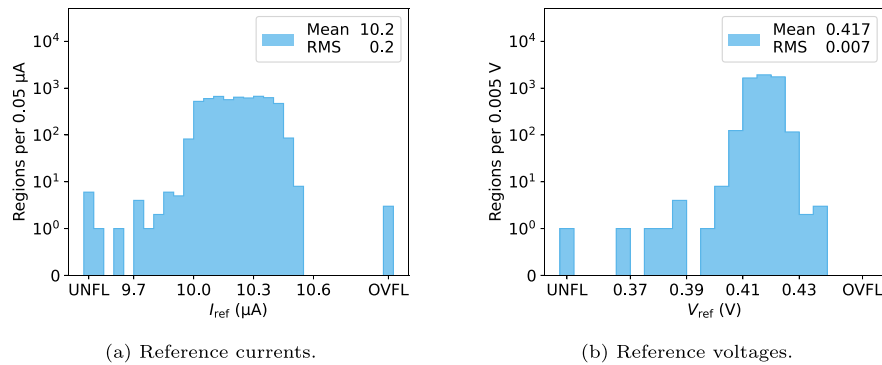
register readback errors. Since these errors affected only a single out of four regions in a half-unit, they are attributed to defects within the sensor periphery rather than slow-control communication. Additionally, two shift-registers used for masking pixels were tested, with 0.1% of regions showing failures. These failures are not considered critical, since the structure was implemented solely to simplify the prototype design and will not be included in the final ITS3-sensor design. The radiation sensitivity of these registers is discussed in Section 6.

### 5.3. Analogue biasing

Testing of the analogue biasing block begins with tuning the bandgap reference voltages (see Section 2.2), followed by measuring the DAC reference voltages and currents. Each DAC (8 per region, 640 per MOSS sensor) is then varied over its full range, and its output is measured via external ADCs connected to dedicated bonding pads on



**Fig. 24.** Current distribution for all power nets across all successfully powered sensors. AVDD, DVDD, and IOVDD nets are measured per half-unit (1353 entries). BBVDD and BBIOVDD nets are measured for the top and bottom half of each sensor (145 entries). The PSUB current is measured per MOSS sensor (76 entries). Currents above overflow ('OVFL') are not shown for better visualisation of the individual distributions.



**Fig. 25.** Distributions of reference currents and voltages for all regions, with values outside the range placed in underflow and overflow bins. The average values are close to the target values of  $10.2 \mu\text{A}$  and  $0.4 \text{ V}$  for  $I_{\text{ref}}$  and  $V_{\text{ref}}$ , respectively.

the long edge of the sensor. The purpose of this test is to verify that the in-pixel front-end can be biased within the designed range.

The distributions of reference currents and voltages for all regions are shown in Fig. 25. The average values are close to the target values of  $10.2 \mu\text{A}$  and  $0.4 \text{ V}$  for  $I_{\text{ref}}$  and  $V_{\text{ref}}$ , respectively. The spread, depending also on the limited precision of the bandgap tuning, is small enough (FWHM less than 5%) to provide reliable biasing references across regions and sensors. The few outliers are not expected to affect sensor operation, as their impact can be compensated by adjusting the corresponding DAC settings.

An analysis was performed for each DAC to assess linearity and operational range. As an example, Fig. 26 shows the distribution of integral non-linearity for all measured current and voltage DACs. In general, linearity stays within 3 DAC counts for the 8-bit DACs (see Section 2). The same result was observed for the differential non-linearity.

DACs with integral non-linearity above 3 counts were considered non-compliant.  $I_{\text{reset}}$  and  $V_{\text{casb}}$  DACs show larger non-linearity than the others. For  $I_{\text{reset}}$ , this is due to the difficulty of supplying and measuring very small currents of the order of a few picoamperes. For  $V_{\text{casb}}$ , changing its value in a critical range can strongly affect the operating point, altering the power consumption of the pixel matrix and causing supply voltage drops that affect the voltage measurement itself.

A region passes the analogue biasing test if all its DACs perform within predefined ranges, ensuring that reliable biasing conditions of the pixel matrix can be established within the operational range (see Section 4.3). In total, 0.65% of regions did not meet this criterion.

#### 5.4. Pixel matrix readout

Pixel matrix readout testing verifies the correct propagation of hit pixel addresses from the pixel matrix to the digital circuitry in the

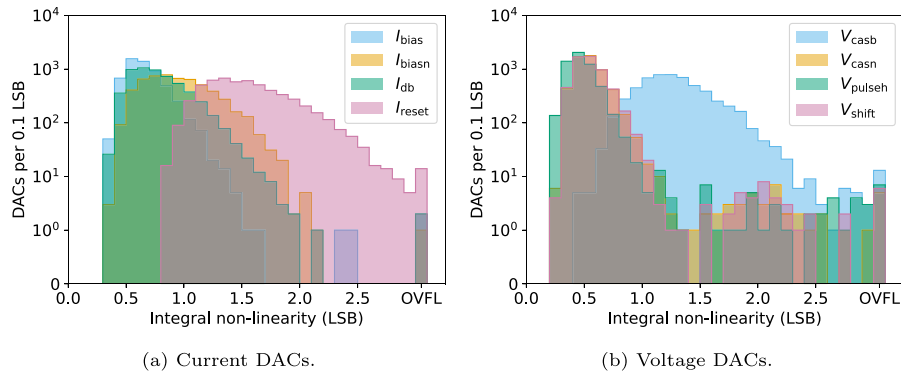


Fig. 26. Distributions of DAC integral non-linearity for all tested regions. DACs with integral non-linearity exceeding 3 DAC counts are considered non-compliant.

sensor periphery and the successful transmission of data packets to the acquisition system. First, using dedicated testing circuitry, in-pixel latches are asserted (see Section 2.2), and pixel addresses are read out row by row. This stage identifies half-units with faulty data-readout interfaces (0.1% of tested half-units) as well as various pixel matrix issues, including faulty pixels, rows, and columns. The faulty columns and rows arise from the simplistic matrix-steering and readout architecture adopted in the MOSS sensor, given that its design did not target to achieve readout performance or resilience to faults. Specifically, a pixel whose latch cannot be de-asserted blocks the readout at its address, and if such a pixel also cannot be masked, it forces masking of an entire column, row, or of a full region. Overall, 449 out of 5544 regions (8.1%) had to be fully masked and excluded from readout. Since these failures were anticipated at the design stage, and the final ITS3 sensor will employ a different readout architecture, they are not considered critical.

The operation of the pixels is tested in two phases. First, the in-pixel hit latches are set via digital configuration, overriding the analogue front-end and validating the functionality of the digital pixel section and of the matrix readout. Then, testing with the injection of a test charge at the input of the front-end is executed (see Section 2.1). Pixels are first classified into four categories with the direct digital test: dead pixels (never fire or always fire, requiring masking), noisy pixels (fire when not expected), inefficient pixels (do not fire reliably when expected), and good pixels. Good pixels are further tested with the charge injection to verify the functionality of the analogue in-pixel front-end. Based on the response to this test, these pixels are reclassified accordingly. As illustrated in Fig. 27 for the dead pixel category, faulty pixels are rare, with only a small number of regions containing them. Regions with more than 1% faulty pixels are classified as failing this test, which corresponds to 0.6% of all tested regions.

### 5.5. Pixel matrix performance

Threshold and noise are measured per pixel, while the fake-hit rate is measured per region (see Section 4.1). Pixel matrix performance is then evaluated in terms of threshold and noise uniformity, and fake-hit rate after masking the noisiest pixels. Regions with more than 1% faulty pixels, either excessively noisy or with undetermined thresholds, are classified as having unsatisfactory performance, accounting for 0.33% of the tested regions.

Fig. 28 shows the distributions of region-average threshold, threshold RMS, and average noise for all regions with 22.5  $\mu\text{m}$  pixel pitch and standard front-end implementation using nominal biasing settings (see Section 4.3). Values are reported in  $V_{\text{pulseh}}$  DAC counts, as calibrating the injected charge at this scale (about 5000 regions in total, 642 shown here) was not feasible within the available time. The observed distribution spread of about 10–15% reflects variations in pulsing capacitance (see Section 4.6), differences in analogue biasing (see Section 5.3), and variation in reverse bias voltage (discussed below).

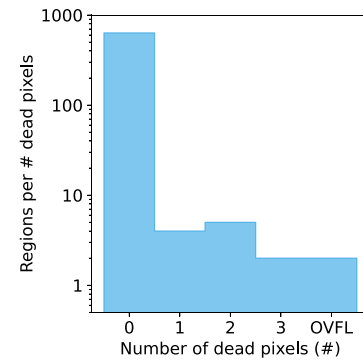



Fig. 27. Distribution of dead pixels per region selecting only the top region 0. The other regions exhibit nearly identical behaviour. In most regions, there are no dead pixels.

Although the threshold can be set independently for each region, reducing the spread in Fig. 28 to a few percent, studying its variation across regions before any adjustment provides insight into potential design, manufacturing, or operational systematic effects. For example, the required operational margin for biasing parameters can be determined, and observed spatial patterns can be correlated with manufacturing steps or test setup effects. Fig. 29 shows a wafer-shaped map of average deviations from the mean region threshold across all wafers. The maximum average deviation per region is about  $25 e^-$ , or roughly 13%. Two patterns are apparent, a horizontal gradient, and a lower average threshold in all regions corresponding to the position of MOSS number 5 on the wafer (see Fig. 2). The latter is linked to three specific sensors (out of fourteen) in the MOSS-5 position. Their lower average threshold is explained by the high substrate bias currents of 26 mA, 44 mA, and 53 mA (see the PSUB distribution tail in Fig. 24). Higher currents reduce the substrate bias seen by the front-end, lowering the threshold. The origin of the horizontal gradient remains unclear, though effects related to the test setup have been excluded.


Fig. 30 shows the fake-hit rate measurements for all regions with 22.5  $\mu\text{m}$  pixel pitch and standard front-end implementation (the same as in Fig. 28). As discussed in Section 4.1, noisy pixels are masked and excluded from the fake-hit rate calculation. In most cases, no noisy pixels are identified, and if present, masking a few is usually sufficient to significantly reduce the fake-hit rate. The spread in the fake-hit rate distribution corresponds to the spread of the threshold distribution (see Fig. 28).

### 5.6. Left End-Cap

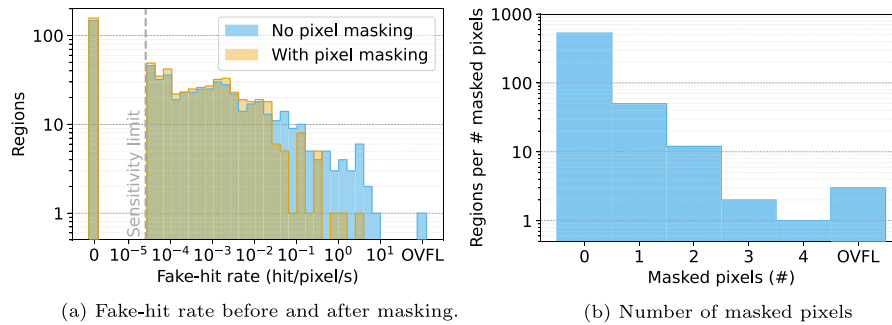
One of the development goals for the MOSS sensor was to prototype and validate the on-sensor data transfer with circuits extending over



**Fig. 28.** Distribution of the threshold average, threshold RMS, and noise average for all the tested regions with  $22.5\ \mu\text{m}$  pixel pitch and standard front-end implementation.  $1\ V_{\text{pulseh}}\ \text{DAC}$  corresponds to about  $8\ e^-$ .



**Fig. 29.** Average deviation from the mean region threshold as a function of wafer position. Two patterns are apparent, a horizontal gradient, and a lower average threshold in all regions corresponding to the MOSS-5 position.  $1\ V_{\text{pulseh}}\ \text{DAC}$  corresponds to about  $8\ e^-$ .



**Fig. 30.** Fake-hit rate performance of regions with  $22.5\ \mu\text{m}$  pixel pitch and standard front-end variant. Masked pixels are those that occur in more than 1% of events (see Section 4.1). The sensitivity limit is determined by the number of events, i.e. the duration of the measurement.

the full sensor length and crossing multiple stitching boundaries (see Section 2) between the RSU and reaching the LEC. Since future ITS3 sensors are expected to rely exclusively on readout via one of the extremities, it is important to validate the response of MOSS half-units when read out through the stitched communication backbone and the data interfaces on the LEC. To do so, several tests were repeated with half-units steered and read out via the LEC instead of their individual interfaces on the long edges. These tests inspect digital periphery functionalities, pixel matrix readout, and pixel matrix performance. Comparing the results between the two operating modes makes it possible to determine whether any issues arise from (1) the communication between the half-units and the LEC over the stitched communication backbone or (2) from the LEC itself. Out of 78 sensors tested via the LEC, one issue for each type was identified, resulting in 21 additional regions (about 0.7%) not passing this check. In pixel matrix performance tests, threshold and fake-hit rate measurements were found to be consistent across both readout schemes.

### 5.7. Half-unit and region yield

**Fig. 31** summarises the yield losses, normalised per region, across the different testing phases described in the previous sections. The largest yield loss occurs during the powering test due to the faults in

the power network, an issue which is not expected to affect forthcoming sensor designs (see Section 5.1). As discussed in Section 5.4, yield loss in the matrix readout can result from issues in the simplistic readout architecture used in this prototype. Since these issues were anticipated in the sensor design, they can be excluded from the yield loss estimate. All other tests indicate a robust design, with yield losses remaining below one percent at each stage. Overall, about 76% of regions pass the full test sequence considering all issues, and 85% when excluding the failures related to the readout-architecture limitations. Excluding also the powering issues, the region yield observed in functional tests is above 98%.

**Fig. 32** shows the functional yield of half-units, including readout architecture issues, as a function of their wafer position. No specific patterns are apparent. Significantly, no top-bottom asymmetry within MOSS sensors is observed, indicating that the higher matrix-layout density of the bottom half-units due to the smaller pixel pitch did not cause any statistically significant yield loss.

The half-unit yield values per wafer are shown in **Fig. 33**, for each of the 14 wafers systematically tested. The region yields are reported with four categories: failures of powering, functional but not meeting performance criteria, failures attributable to the sensor readout architecture, and fully functional. Powering yield loss is more common in odd-numbered wafers than in even-numbered wafers. An analysis of processing control and monitoring data confirmed the existence of small

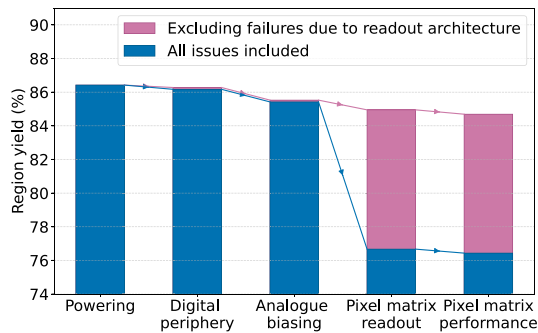


Fig. 31. Yield losses, normalised per region, across the different sensor block verifications.

differences between the odd-numbered and even-numbered wafers for some of the characteristics related to the metal interconnects. This correlates with the observed yield fluctuations and is attributed to the two wafer subsets being processed under slightly different conditions.

### 5.8. Discussion on the yield of ITS3 layers

The series-testing data can be used to estimate the counts of ITS3 half-layers that could be produced in the hypothetical scenario of building them with MOSS sensors and with the yield observed within the ER1 engineering batch. Each half-layer can be produced from a single wafer by dicing a set of adjacent MOSS sensors that meet the acceptance criteria. The ITS3 half-layers for Layer 0, 1, and 2 would need respectively four, five or six adjacent functional MOSS sensors on one wafer. Each ER1 wafer contains six adjacent sensors (see Section 2), allowing multiple possible mappings of half-layers to wafers. The ITS3 requirements tolerate loosing up to 2% of the sensitive area [8] due to local defects, which for MOSS corresponds to 2% region failures. The following estimates of acceptable half-layer counts are based on 12 fully tested<sup>11</sup> MOSS wafers, a half of a typical production lot, limiting the total number of possible combinations for each half-layer. Priority was given to achieving equal numbers of different half-layers rather than favouring a particular layer size. Finally, it should be emphasised that these estimates are based on an exploratory prototype sensor, which still lacks features needed for ITS3, such as high-speed serializers and the ability to power the entire sensor exclusively from the end-caps.

Fig. 34 shows the number of ITS3 half-layers that can be successfully assembled from a set of twelve fully analysed wafers, under different acceptance assumptions. Two scenarios are presented, following the reasoning in the previous section: (a) excluding only region failures attributed to the readout architecture, and (b) excluding both readout architecture and powering-related failures. The latter condition is less restrictive for sensor quality as it assumes that manufacturing imperfections associated with the metal stack can be fully mitigated in the future productions, which is yet to be demonstrated. Three cases are considered, (i) region failures amount to less than 2%, (ii) less than 3%, and (iii) less than 5%. If 2% of failures are tolerated than the scenario which excludes both readout architecture and powering-related failures would provide enough sensors to assemble the full ITS3 barrel, given that two half-layers per layer size are needed and three for each one are found. In the first scenario, which excludes only region failures attributed to the readout architecture, one would have to tolerate 5% of region failures to be able to obtain two half-layers of each type from the given twelve wafers. However, increasing the initial number of wafers would likely allow assembly of a full ITS3 barrel even under the most stringent conditions of less than 2% region failures.

<sup>11</sup> The remaining two wafers contained one mechanically broken sensor each, and thus were excluded from this exercise.

## 6. Single-event effect measurements

The collisions in the LHC create a radiation environment with a flux of particles that can generate stochastic Single-Event Effects (SEE), and in particular Single-Event Upset (SEU) and Single-Event Latch-up (SEL) in CMOS circuits. The future ITS3 sensor shall exhibit sufficient robustness against SEU to achieve reliable operation under these operating conditions. It shall also have low sensitivity to SEL to prevent potentially destructive radiation-induced over-currents. Dedicated SEE tests were carried out on the MOSS sensor, which does not implement TMR redundancy techniques, to quantitatively evaluate the impact of Single Event Effects and to assess the potential need for mitigation measures in the final ITS3 sensor design.

Sensitivity to SEU was measured by programming specific data patterns into sensor registers and recording bit flips as a function of irradiation type, flux (typically hadrons), and exposure time. SEU characterisation tests were performed at the NPI cyclotron in Řež, using 30 MeV proton beams [28]. Several MOSS registers implemented with flip-flops, performing various functions and located in different regions of the sensor, were simultaneously irradiated. All monitored registers exhibited SEU cross sections between  $10^{-14}$  cm<sup>2</sup> bit<sup>-1</sup> and  $4 \times 10^{-14}$  cm<sup>2</sup> bit<sup>-1</sup>, consistent with previous results from dedicated SEU test sensors fabricated in the same CMOS technology.

While SEUs are generated directly by the products of the primary collisions, SEL events are caused indirectly by such particles through the production of nuclear-recoil fragments inside the silicon sensor. Such fragments can have a Linear Energy Transfer (LET) sufficiently high to trigger a SEL. On the other side the LET of these fragments has been shown to not exceed 15 MeV cm<sup>2</sup> mg<sup>-1</sup> under the operational conditions of the LHC [29]. The sensitivity to SEL is assessed by detecting and measuring the frequency of latch-up events as a function of a flux of particles with a well defined LET. A latch-up event is characterised by a persistent increase of one of the sensor supply currents due to the creation of a parasitic, self-sustaining thyristor structure in the sensor. Resetting such structures requires power-cycling either the entire sensor or the smallest sub-sensor power domain where SEL has occurred. The sensitivity to SEL is highly dependent on specificities of the detailed layout of a given circuit and has therefore be tested on the MOSS sensor to identify locations eventually requiring modifications of the circuit layout in the final sensor.

SEL tests were conducted at the Heavy Ion Facility of UCL (Louvain-la-Neuve, BE) and at the BASE facility of LBNL (Berkeley, US). Ion beams with LET between 3.3–62.5 MeV cm<sup>2</sup> mg<sup>-1</sup> and fluxes between 500–15 000 ions/cm<sup>2</sup>/s were used, with about 10% homogeneity over an irradiation area with a diameter of 2.5 cm. To avoid beam degradation, the setup was installed in a vacuum vessel, and the sensor was connected to a cooling circuit for temperature stabilisation. During irradiation, the supply currents and the correct functioning of the sensor were continuously monitored. Collimators and movable stages were also used to expose selectively small areas of the sensors to better identify locations with high SEL sensitivity. The threshold for the detection of over-currents was set to 50 mA to enable clear identification of typical latch-up current pulses.

Fig. 35 presents the measured SEL cross-section for a babyMOSS device irradiated without any collimators as a function of the LET of the incident ion species. The data exhibit the characteristic steep rise in cross section at low LET values, followed by a more gradual increase for LET values above approximately 20–30 MeV cm<sup>2</sup> mg<sup>-1</sup>. Despite the relatively high detection threshold of 50 mA, SEL events were observed at LET values below 15 MeV cm<sup>2</sup> mg<sup>-1</sup>. Targeted irradiations using various collimators enabled the localisation of SEL-sensitive regions to the periphery of the sensor. Notably, no SEL events were detected when irradiating the pixel array. An insufficient density of well contacts was identified in certain peripheral components, prompting corrections for subsequent iterations of the sensor design. SEL-sensitivity testing will be repeated on future designs to assess the likelihood of latch-up under the expected operational conditions of ITS3.

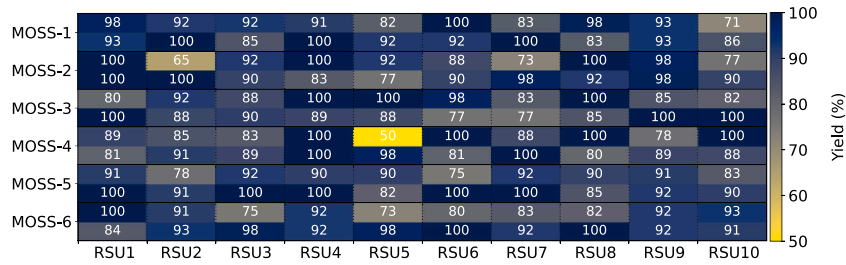


Fig. 32. Half-unit yield as a function of wafer position. No specific patterns are apparent.




Fig. 33. Yield overview per wafer. For each wafer, the fraction of regions in different categories is shown as a stacked representation. Powering yield loss is more common in odd-numbered wafers than in even-numbered wafers.




Fig. 34. The number of ITS3 half-layers (HL) that can be successfully assembled with MOSS sensors from a set of 12 wafers<sup>11</sup> under different assumptions on the allowed region failure rate and the types of excluded failures. Both scenarios are expected to be compatible with the final production.

7. Conclusions

This paper presents an overview of the design, as well as the results of a comprehensive testing and characterisation campaign, of the MOSS sensor, a novel wafer-scale prototype CMOS MAPS sensor, assembled from ten identical sensors using sophisticated stitching techniques. Its design and testing represented key milestones in the development of stitched MAPS and the prototyping of future ALICE ITS3 sensors, while also generating valuable know-how that will be reflected in the subsequent sensor design iteration. The sensor is fully functional, and demonstrates the viability of implementing the stitching technique as a means to manufacture sensing devices of different size with one modular design, in line with the integration requirements of the ITS3 half-layers. Complex power distribution with multiple power




Fig. 35. SEL cross section per total area irradiated babyMOSS as a function of LET measured at the Heavy Ion Facility of UCL (Louvain-la-Neuve, BE) and at the BASE facility of LBNL (Berkeley, US). The dashed line indicates the maximum LET that fragments generated in silicon can reach under LHC operational conditions.

nets, metal interconnects across stitching boundaries, and transmission of signals on-sensor over about 25 cm have been validated.

A total of 82 full MOSS sensors were tested systematically in the lab, each one containing 10 Repeated Sensor Units, amounting to 6560 powered and controlled pixel matrices, and over half a billion individually characterised pixels. The testing of the ER1 lot identified an excess rate of faults in the sensor power grid. A detailed failure analysis was conducted and enabled the identification of manufacturing imperfections associated with the metal stack, newly introduced by the foundry to meet the specific requirements of the project. Feedback was shared with the foundry, enabling the identification of the root cause and the implementation of corrective measures that are expected to drastically reduce the likelihood of these defects in future productions with the same metal stack [27].

Despite these faults, the qualification demonstrated correct operation and acceptable performance of the individual building blocks and a resilient design, with a comprehensive functional yield exceeding 98%. Extrapolations based on the systematic functional qualification of the MOSS sensors show a yield sufficient to construct ITS3 half-layers from single wafers. The results of the systematic laboratory analysis did not find significant variations of characteristics inside a wafer and from wafer-to-wafer, once the faults related to the metal-stack issue were excluded. Significantly, no change of the functional failure rates was observed in relation to the different layout densities prototyped in the two halves of the MOSS sensor, providing experimental evidence of the possibility to relax some of the conservative margins adopted in the layout of the pixel array with larger pitch.

A subset of sensors underwent extensive pixel matrix characterisation in the laboratory and with ionising-particle beams. The required performance, a detection efficiency above 99% and a fake-hit rate below 0.1 hits/pixel/s was achieved. Spatial resolution was found to be

5–5.5  $\mu\text{m}$  for 22.5- $\mu\text{m}$ -pitch pixels and 4–4.5  $\mu\text{m}$  for 18- $\mu\text{m}$ -pitch pixels, indicating that an intermediate pixel pitch will be sufficient to meet the ITS3 target of 5  $\mu\text{m}$ . No significant differences of key performance figures were observed across the pixel variants prototyped in MOSS. Cross-coupling effects between digital signals and sensitive analogue nodes were identified and thoroughly studied, providing useful knowledge for mitigation in future design iterations. Measurements of SEL cross sections with heavy-ion beams revealed a sensitivity of MOSS to single-event latch-ups for LET below 15 MeV cm<sup>2</sup> mg<sup>-1</sup>, localised in specific digital peripheral blocks and attributed to sub-optimal contacting of wells.

In conclusion, the MOSS sensor demonstrates the viability of stitching to implement large area MAPS and exhibits promising performance and yield for the ALICE ITS3 upgrade. Remaining challenges for the next sensor include demonstrating high-speed data transmission, both on- and off-sensor, and managing the voltage drops resulting from power being supplied solely at the short sensor edges — features that are absent in the MOSS design. The comprehensive characterisation of MOSS provides valuable insights into the sensor's behaviour under different operating conditions and irradiation levels, paving the way for the design of the final ITS3 sensor and for further developments and applications.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgements

Filip Křížek acknowledges the support by the Ministry of Education, Youth and Sports of the Czech Republic project LM2023040.

Zijun Zhao acknowledges the support by the National Key Research and Development Program of China (2022YFA1602103).

Tommaso Fagotto and Giovanni Vecil acknowledge the support by the Italian Ministry of Foreign Affairs and International Cooperation (project PGR12370).

This research has received funding support from NSRF via the Program Management Unit for Human Resources & Institutional Development, Research and Innovation [B46G680100], and NSTDA under the contract number JRA-CO-2563-14066-TH.

The authors thank the Optimato team at the Elettra Sincrotrone for providing their facility and expert support in conducting the X-ray fluorescence measurements.

### Data availability

Data will be made available on request.

### References

- [1] G. Contin, et al., The STAR MAPS-based PiXeL detector, NIM-A 907 (2018) 60–80, <http://dx.doi.org/10.1016/j.nima.2018.03.003>.
- [2] M. Šuljić, ALPIDE: The monolithic active pixel sensor for the ALICE ITS upgrade, J. Inst. 11 (11) (2016) C11025, <http://dx.doi.org/10.1088/1748-0221/11/11/c11025>.
- [3] G. Aglieri Rinella, The ALPIDE pixel sensor chip for the upgrade of the ALICE inner tracking system, NIM-A 845 (2017) 583–587, <http://dx.doi.org/10.1016/j.nima.2016.05.016>.
- [4] M. Mager, ALPIDE, the monolithic active pixel sensor for the ALICE ITS upgrade, NIM-A 824 (2016) 434–438, <http://dx.doi.org/10.1016/j.nima.2015.09.057>.
- [5] ALICE Collaboration, ALICE upgrades during the LHC long shutdown 2, J. Inst. 19 (05) (2024) P05062, <http://dx.doi.org/10.1088/1748-0221/19/05/p05062>.
- [6] ALICE Collaboration, URL <https://alice-collaboration.web.cern.ch>. (Accessed 22 September 2025).
- [7] CERN EP R&D, 2025, URL <https://ep-dep.web.cern.ch/rd-experimental-technologies>. (Accessed 22 September 2025).
- [8] ALICE Collaboration, Technical Design Report for the ALICE Inner Tracking System 3 - ITS3; A Bent Wafer-Scale Monolithic Pixel Detector, Tech. rep., CERN, Geneva, 2024, URL <https://cds.cern.ch/record/2890181>.
- [9] Tower Partners Semiconductor Co., URL <http://www.towersemi.com/>. (Accessed 22 September 2025).
- [10] G. Aglieri Rinella, et al., Characterization of analogue monolithic active pixel sensor test structures implemented in a 65 nm CMOS imaging process, NIM-A 1069 (2024) 169896, <http://dx.doi.org/10.1016/j.nima.2024.169896>.
- [11] G. Aglieri Rinella, et al., Digital pixel test structures implemented in a 65 nm CMOS process, NIM-A 1056 (2023) 168589, <http://dx.doi.org/10.1016/j.nima.2023.168589>.
- [12] G. Aglieri Rinella, et al., Further characterisation of digital pixel test structures implemented in a 65 nm CMOS process, 2025, [arXiv:2505.05867](https://arxiv.org/abs/2505.05867).
- [13] G. Aglieri Rinella, et al., Time performance of analog pixel test structures with in-chip operational amplifier implemented in 65 nm CMOS imaging process, NIM-A 1070 (2025) 170034, <http://dx.doi.org/10.1016/j.nima.2024.170034>.
- [14] Tower semiconductor ltd, stitching design rules for forming interconnect layers, us patent 6225013b1, 2001, 2025, URL [http://refhub.elsevier.com/S0168-9002\(23\)00008-6/sb11](http://refhub.elsevier.com/S0168-9002(23)00008-6/sb11). (Accessed 15 July 2025).
- [15] P. Vicente Leitao, et al., Development of a stitched monolithic pixel sensor prototype (MOSS chip) towards the ITS3 upgrade of the ALICE inner tracking system, J. Inst. 18 (01) (2023) C01044, <http://dx.doi.org/10.1088/1748-0221/18/01/C01044>.
- [16] ALICE ITS project, First demonstration of in-beam performance of bent monolithic active pixel sensors, NIM-A 1028 (2022) 166280, <http://dx.doi.org/10.1016/j.nima.2021.166280>.
- [17] A. Andronic, et al., Detection efficiency and spatial resolution of monolithic active pixel sensors bent to different radii, 2025, [arXiv:2502.04941](https://arxiv.org/abs/2502.04941).
- [18] F. Piro, et al., A compact front-end circuit for a monolithic sensor in a 65-nm CMOS imaging technology, IEEE Trans. Nucl. Sci. 70 (9) (2023) 2191–2200, <http://dx.doi.org/10.1109/tns.2023.3299333>.
- [19] W. Snoeys, et al., Optimization of a 65 nm CMOS imaging process for monolithic CMOS sensors for high energy physics, PoS Pixel2022 (2023) 083, <http://dx.doi.org/10.22323/1.420.0083>.
- [20] P. Ahlburg, et al., EUDAQ—a data acquisition software framework for common beam telescopes, J. Inst. 15 (01) (2020) P01038, <http://dx.doi.org/10.1088/1748-0221/15/01/p01038>.
- [21] D. Dannheim, et al., Corryvreckan: A modular 4D track reconstruction and analysis software for test beam data, J. Inst. 16 (03) (2021) P03008, <http://dx.doi.org/10.1088/1748-0221/16/03/P03008>.
- [22] V. Blobel, A new fast track-fit algorithm based on broken lines, NIM-A 566 (1) (2006) 14–17, <http://dx.doi.org/10.1016/j.nima.2006.05.156>.
- [23] S. Emiliani, et al., Investigation of non-idealities of pulsing circuitry in the ALICE ITS3 MOSS monolithic sensor, J. Inst. 19 (12) (2024) C12021, <http://dx.doi.org/10.1088/1748-0221/19/12/C12021>.
- [24] V.D. Trapani, et al., Speckle-based imaging (SBI) applications with spectral photon counting detectors at the newly established OPTIMATO (OPTimal IMAGING and tomography) laboratory, J. Inst. 19 (01) (2024) C01018, <http://dx.doi.org/10.1088/1748-0221/19/01/C01018>.
- [25] Elettra Sincrotrone Trieste, URL <https://www.elettra.eu/it/index.html>. (Accessed 22 September 2025).
- [26] H. Wennlöf, et al., Simulating monolithic active pixel sensors: A technology-independent approach using generic doping profiles, NIM-A 1073 (2025) 170227, <http://dx.doi.org/10.1016/j.nima.2025.170227>.
- [27] G. Aglieri Rinella, et al., Fault detection and failure analysis of CMOS metal stacks, IEEE Trans. Nucl. Sci. (2025) In submission to.
- [28] F. Křížek, et al., Irradiation setup at the U-120M cyclotron facility, NIM-A 894 (2018) 87–95, <http://dx.doi.org/10.1016/j.nima.2018.03.066>.
- [29] R. García Alía, et al., Ultraenergetic heavy-ion beams in the CERN accelerator complex for radiation effects testing, IEEE Trans. Nucl. Sci. 66 (1) (2019) 458–465, <http://dx.doi.org/10.1109/TNS.2018.2883501>.