Published by IOP Publishing for Sissa Medialab

Received: *November 3, 2023*
Accepted: *March 8, 2024*
Published: *April 23, 2024*

Topical Workshop on Electronics for Particle Physics
Geremeas, Sardinia, Italy
1–6 October 2023

Data transmission architecture of the ALICE ITS3 stitched sensor prototype MOSAIX

P. Dorosz on behalf of the MOSIAX design team and the ALICE collaboration

CERN,
Esplanade des Particules, 1217 Meyrin, Switzerland

E-mail: piotr.dorosz@cern.ch

Abstract: The ALICE Collaboration will replace the three innermost layers of the Inner Tracker System (ITS) at the LHC with an innovative vertex detector. A single-die stitched monolithic pixel detector segment of 1.85 cm × 26.6 cm designed in 65 nm CMOS imaging technology will be used as a building block for these layers. The pixel detector segment consists of 144 data transmitters that are evenly distributed over the full area. The data communication is done via the 1.85 cm short edge of the detector. This contribution will focus on the architecture, challenges, and techniques used to aggregate up to 30.72 Gb/s of data flux arriving at the short edge of the chip and to send it off-chip.

Keywords: Data acquisition circuits; VLSI circuits; Electronic detector readout concepts (solid-state)

https://doi.org/10.1088/1748-0221/19/04/C04050

2024 JINST 19 C04050

Contents

1 **Introduction** 1

2 **On-chip data transmission** 2

3 **Off-chip data transmission** 4

4 **Performance of the phase aligner** 4

5 **Power consumption of the left endcap** 5

6 **Summary** 6

2024 JINST 19 C04050

1 Introduction

The three innermost layers of the ALICE Inner Tracker System (ITS) will be replaced by the ITS3 [1]. Each layer is divided into identical half-layers (see figure 1). The large area thin silicon sensor is a 12″ wafer-size sensor that is bent to a half-cylinder shape. The bent sensor is 26.6 cm long and 50 µm thick. It is electrically connected by means of wire bonding to two flexible printed circuits (FPC) from A-side and C-side. Both FPCs provide power to the sensor but only the C-side FPC is used for control signals and data transmission. The half-layer of the ITS3 is divided into 26.6 cm × 1.85 cm segments. The three different ITS3 half-layers are constructed by 3, 4 or 5 segments for layers 0, 1, and 2, respectively.

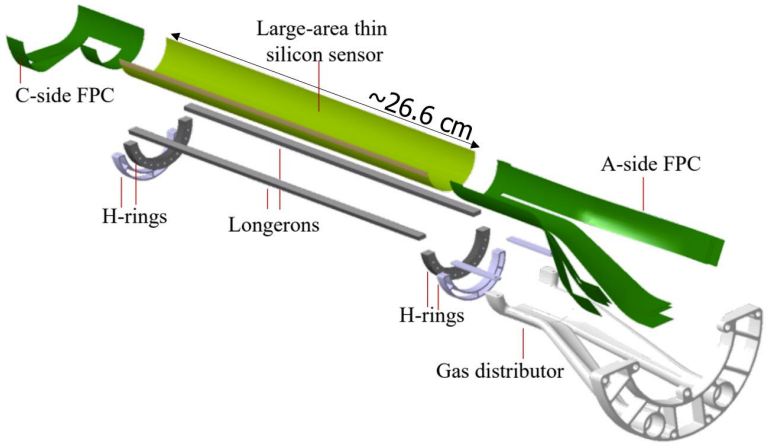

Figure 1. ALICE ITS3 half-layer view. Figure courtesy ITS3 project Work Package 5.

The segment is an independent monolithic stitched sensor (MOSAIX) and consists of three building blocks (figure 2): a repeated sensor unit (RSU) containing the pixel array, a left endcap (LEC), used for powering and data transmission, and a right endcap (REC), used for powering only. These building blocks form a wafer-scale design, where the RSU, the middle part, is repeatedly placed 12 times next to each other and connected via wafer stitching to the LEC for data transmission and power delivery.

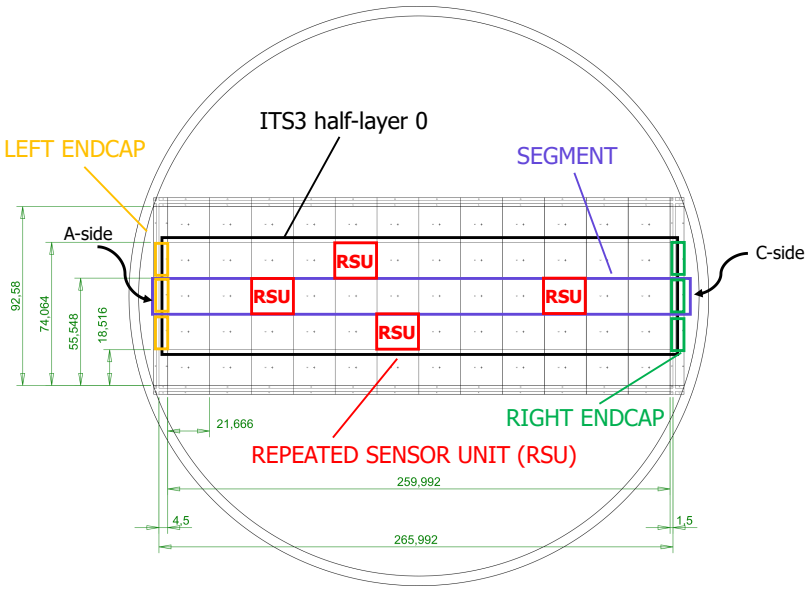

Figure 2. Wafer stitching plan with dimensions of the ITS3 half-layer in mm.

MOSAIX has 26.6 cm of length and 1.85 cm of width with 10.7 million pixels in total. An external asynchronous strobe signal of programmable duration determines an integration window for the pixel signals to be stored, processed, and sent to the left endcap. The LEC is the only interface between MOSAIX and the off-chip domain. The LEC has 144 on-chip inputs with transfer rate 160 Mb/s coming from the pixel matrix and 6 outputs with transfer rate 5.12 Gb/s to the off-chip domain. This gives a total data rate of 23.04 Gb/s of the input and 30.72 Gb/s of the output transmission. The extra bandwidth of the output transmission is used for link redundancy. This paper will present a solution for two main design challenges:

- How to transmit data at 160 Mb/s and a clock at 160 MHz across a 26.6 cm long chip?
- How to transfer 30.72 Gb/s off-chip?

2 On-chip data transmission

The RSU of MOSAIX is divided into 12 tiles (figure 3a). Each tile can be considered an independent pixel sensor with its own pixel array, biasing, and data transmitter. The tiles are distributed along the 26.6 cm length of the MOSAIX, hence the distances between tiles and the LEC vary. Each tile stores data from the pixel array in an internal memory, then serializes it at 160 Mb/s rate in formatted packets, and transmits it over a dedicated on-chip link to the LEC. There are one-to-one direct connections between each of the 144 tiles in MOSAIX and the LEC. The lengths of the on-chip links correspond to the distances between the tiles and the LEC and that is why they vary between 1 mm and 26 cm. The bit stream on each on-chip link arrives to the LEC with a different unknown phase due to various propagation delays and temperature and process variations. For these reasons the LEC needs to perform independent phase alignment of the incoming data for each of the 144 links. The data in the LEC has to be sampled in the middle of the eye opening to prevent errors occurring due to the jitter that is present in the incoming data.

2024 JINST 19 C04050

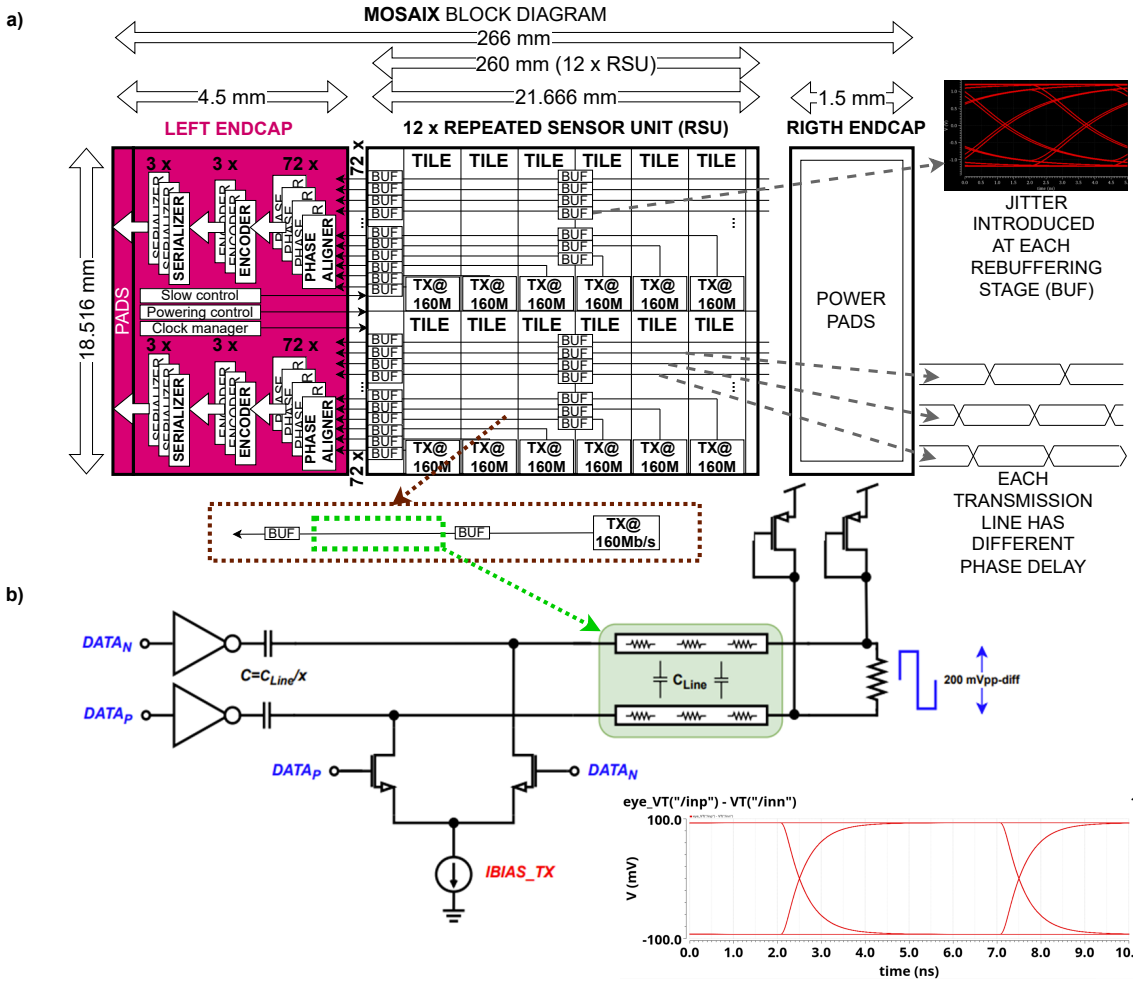

Figure 3. Block diagram of MOSAIX: a) data transmission path from tiles (TX @ 160 Mb/s) to the left endcap b) simplified schematic of one section of an on-chip transmission link with the simulated eye-diagram of the analog model of the section.

The on-chip transmission link has a full-custom differential design (figure 3b). The voltage amplitude over the line is limited to 200 mV to reduce the power consumption. The line capacitors (C) are placed after the inverters to limit the voltage swing over the line. The diode-connected PMOS' transistors provide the common mode voltage level and the NMOS transistors are used in a differential pair. The IBIAS_TX current is divided through the two branches of the differential pair depending on $DATA_N$ and $DATA_P$. Due to the signal attenuation the link needs signal re-buffering by a receiver-transmitter module (BUF) that is placed every 3 tiles (11 mm). Based on the extracted line parameters of the link, in order to achieve the required 200 MHz of bandwidth, the section cannot be longer than 11 mm. The extracted simulation of the inter-symbol-interference jitter with a pseudo-random sequence is of 12.5 ps peak-to-peak per section. In the long links that consist of multiple sections the jitter accumulates at each section and reduces the eye opening of 160 Mb/s data bits. The phase aligners in LEC are designed to cope with an eye opening that is limited to at least half the period of the data width.

2024 JINST 19 C04050

3 Off-chip data transmission

The LEC is divided into a top and a bottom parts where each part is handling data from 72 tiles (figure 4). Data arriving from the tiles is first phase aligned in a data aggregator block. Each link has a dedicated phase aligner. The phase alignment is performed with the use of a delay line that generates a set of phases from the incoming data bit. The logic locks to the phase that is between the edges of the incoming data bit (the middle of an eye opening). Afterwards, the data from 24 links is deserialized and grouped into a bus of 24×4 bits which goes directly to a data encoder. The A-side FPC is equipped with VTRx+ opto-transmitters [2] and the communication between MOSAIX and off-chip domain will be performed using lpGBT protocol as a transport layer [3]. This means that the lpGBT encoder RTL is instantiated in the LEC and from the perspective of the off-chip domain, it is as if MOSAIX has 6 on-chip lpGBTs. The built-in parity bits with a Forward Error Correction (FEC) technique in the lpGBT protocol reduce the transmission errors that appear due to noise, inter symbol interference, or single-event upsets. Scrambling is used to keep the high density of transitions in the serial bit stream ensuring that the link is DC balanced [3]. Interleaving is used to increase the error correction capability [3]. The encoder also has a pseudo-random bit generator for testing the link. The lpGBT-link encoder logic and phase-alignment logic (ePort-Rx) are based on the lpGBT design. The last stage of data handling is to further serialize the data and send it off-chip at 5.12 Gb/s data rate. For this purpose a custom Giga Wire Transfer-Power Supply Immune (GWT-PSI) serializer is used [4]. GWT-PSI has a dual-port FIFO for clock domain crossing. Each GWT-PSI is a self-contained macrocell that has an internal phase-lock loop and a delay-locked loop to generate all clock frequencies required by the serialization process. There are 6 serializers in total (in the top and bottom part of the LEC) to enable transmission from all 144 tiles (1 serializer per 24 tiles). The total off-chip data rate of the MOSAIX is 30.72 Gb/s.

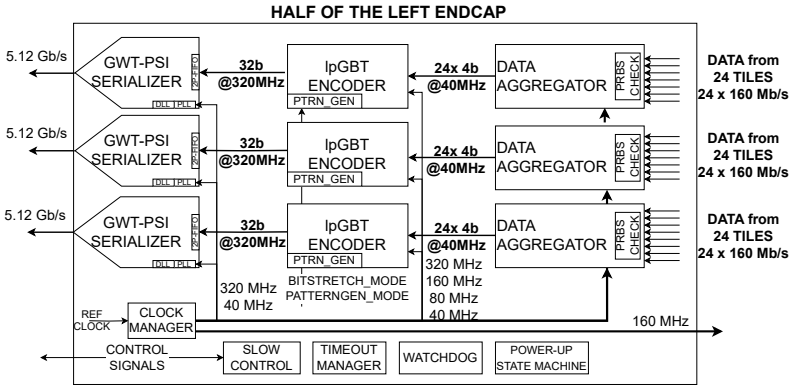

Figure 4. Block diagram of a half of the left endcap.

4 Performance of the phase aligner

The phase aligner is based on a delay line, which generates a set of phases for each incoming data bit (figure 5a). The logic extracts an edge vector, which stores the information on the location of bit edges: "1" in the vector (figure 5b). These vectors are added to form a histogram with the highest

2024 JINST 19 C04050

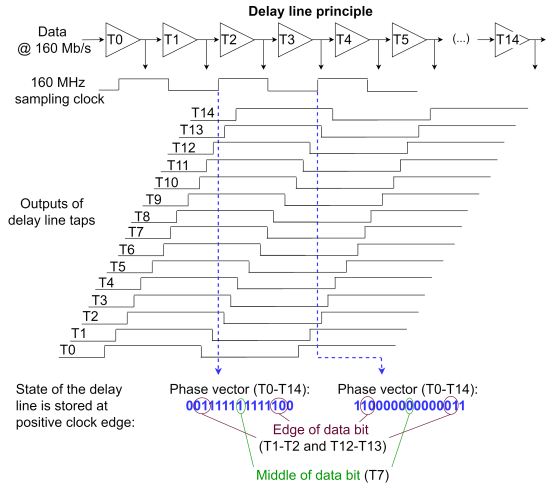

(a) Extraction of the phase vector from a delay line.

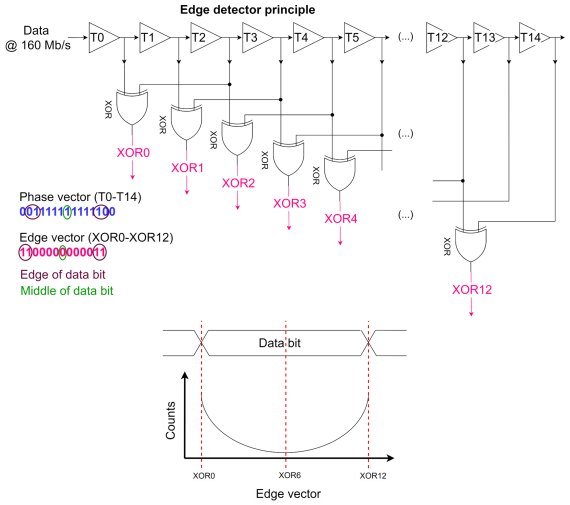

(b) Extraction of the edge vector from a delay line and XOR gates. The edge vector gets registered and detected edges are counted to form a histogram.

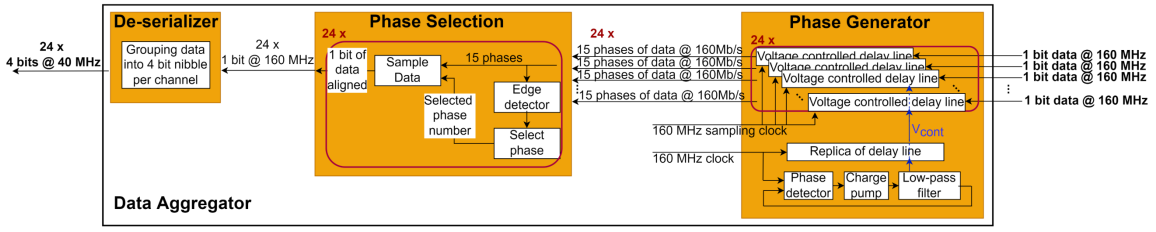

(c) Block diagram of a data aggregator.

Figure 5. Phase alignment principle: (a) A delay line extracts 15 phases from each incoming data bit; (b) An edge vector is extracted from the delay line to locate the edges and the middle of data bit; (c) Edge and phase vectors are propagated to a phase selection state machine and only one phase is chosen for data sampling.

values around bit edges (T1-T2, T12-T13 or XOR0, XOR12) and the lowest values around the middle of a bit (T7 or XOR6). The phase selection logic chooses only one phase for data sampling (figure 5c). This phase corresponds to the middle position of the data bit.

The functional verification of the data aggregator RTL has confirmed that if the jitter of incoming data does not exceeded 50% of the data bit width (3.125 ns) it is still possible to identify the correct phase of incoming data.

5 Power consumption of the left endcap

The power budget of the LEC is limited to 400 mW. The power verification of the implemented data transmission components exhibits the following power consumption: 180 mW from 6 × GWT-PSI; 26.4 mW from 6 × lpGBT encoder; 90 mW from 6 × Data aggregator. This gives a total of 300 mW. The remaining 100 mW is reserved for control logic of the LEC: clock manager, slow control, and power-up state machines.

2024 JINST 19 C04050

6 Summary

This paper presents the design of the stitched monolithic silicon sensor MOSAIX from the perspective of on-chip and off-chip data transmission. The monolithic stitched particle sensor MOSAIX has 144 data transmitters distributed along 26 cm sending data to the left endcap of the chip. The data arriving at the left endcap has different distances to travel depending on the physical location of the source, making the phase of the data unknown. The left endcap performs phase alignment, deserialisation, grouping, and encoding of data before serializing it again at 5.12 Gb/s and sending it off-chip. The framework chosen for the data transfer from the LEC to an off-chip domain is the lpGBT protocol. The total data rate of MOSAIX is 30.72 Gb/s. Power consumption of the left endcap is within the planned 400 mW budget.

References

[1] L. Musa, *Letter of Intent for an ALICE ITS Upgrade in LS3*, CERN-LHCC-2019-018, CERN, Geneva (2019) [DOI:10.17181/CERN-LHCC-2019-018].

[2] J. Troska et al., *The VTRx+, an optical link module for data transmission at HL-LHC*, *PoS* **TWEPP-17** (2017) 048.

[3] lpGBT Design Team, *lpGBT documentation*, (2022), https://cds.cern.ch/record/2809058/.

[4] V. Gromov et al., *Prototype of a 10.24Gbps Data Serializer and Wireline Transmitter for the readout of the ALICE ITS3 detector*, in the proceedings of the *Topical Workshop on Electronics for Particle Physics*, Geremeas, Sardinia, Italy, 1–6 October 2023.

2024 JINST 19 C04050