

System simulations for the ALICE ITS detector upgrade

Simon Voigt Nesbo^{1,2,*}, *Johan Alme*², *Matthias Bonora*⁵, *Piero Giubilato*^{3,5}, *Håvard Helstrup*¹, *Matteo Lupi*⁵, *Gianluca Aglieri Rinella*⁵, *Dieter Röhrich*², *Joachim Schambach*⁴, *Ruben Shahoyan*⁵, and *Arild Velure*⁵

¹Western Norway University Of Applied Sciences, Bergen, Norway

²University of Bergen, Bergen, Norway

³Universita e INFN, Padova, Italy

⁴The University of Texas at Austin, Physics Department, Austin, Texas, United States

⁵European Organization for Nuclear Research (CERN), Geneva, Switzerland

Abstract. The ALICE experiment at the CERN LHC will feature several upgrades for Run 3, one of which is a new Inner Tracking System (ITS). The ITS upgrade is currently under development and commissioning, and will be installed during the ongoing long shutdown 2.

A number of factors will have an impact on the performance and readout efficiency of the ITS in run 3, and to that end, a simulation model of the readout logic in the ALPIDE pixel sensor chips for the ITS was developed, using the SystemC library for system level modeling in C++. This simulation model is three orders of magnitude faster than a normal HDL simulation of the chip and facilitates simulations of an increased number of events for a large portion of the detector.

In this paper, we present simulation results, where we have been able to quantify detector performance under different running conditions. The results are used for system configuration as well as for the ongoing development of the readout electronics.

1 Introduction

The ALICE experiment will be operating at increased luminosities in the third run of the LHC. In preparation for this, upgrades to the experiment are installed during the Long Shutdown 2 (LS2) of the LHC, which is currently ongoing. The Inner Tracking System (ITS) is among the systems which are being upgraded [1]. The specification for the ITS upgrade calls for event rates of up to 100 kHz for Pb–Pb and 400 kHz pp, which is two orders of magnitude higher than the existing system. A new Monolithic Active Pixel Sensor (MAPS), the ALPIDE [2], was developed specifically for this purpose. The ITS detector consists of 24 120 ALPIDE chips arranged in seven cylindrical layers, which are grouped in an Inner Barrel (IB) comprising the three innermost layers, and an Outer Barrel (OB) for the remaining four layers. The layers are constructed from long and narrow *staves* of ALPIDE chips that run along the length of the barrels. The data from the sensor chips are read out by 192 Readout Unit (RU) boards, one RU per *stave*. The RU is also responsible for configuration and control of the chips, as well as trigger distribution. Several factors have an impact on the

*e-mail: svn@hvl.no

performance and readout efficiency of the system, in addition to the different configurations of RUs and staves. These factors are not just limited to operating conditions such as run type and event rates, but also include a number of sensor configuration parameters.

1.1 ALPIDE Monolithic Active Pixel Sensor

The readout logic in the ALPIDE chip is illustrated in Fig. 1 and has been explained in detail in Ref. [3]. The chip has a pixel matrix of 1024 x 512 pixels, which is divided into 32 regions. Pixel hit discrimination is performed in each pixel. There are three event buffers, which are in-pixel digital buffers that can store a full event. The event buffers are collectively referred to as the Multi Event Buffer (MEB). When a trigger is received, the chip initiates a strobe window. The duration of the strobe window is configurable. Discriminated pixel-hits asserted during the strobe are latched into an event buffer for that strobe, as shown in Fig. 2. The system will either use minimum-bias triggers with a short strobe (Fig. 2a), or periodic triggers with a long strobe (Fig. 2b). The chip has two operating modes to go with this; triggered and continuous mode. In principle, the two modes can be used with any type of trigger; they only differ in situations of high occupancy where the three event buffers are not sufficient to capture all events. In triggered mode, new events are lost when all event buffers are already full. In continuous mode, there is a flushing mechanism that flushes the oldest event buffer to make room for new events.

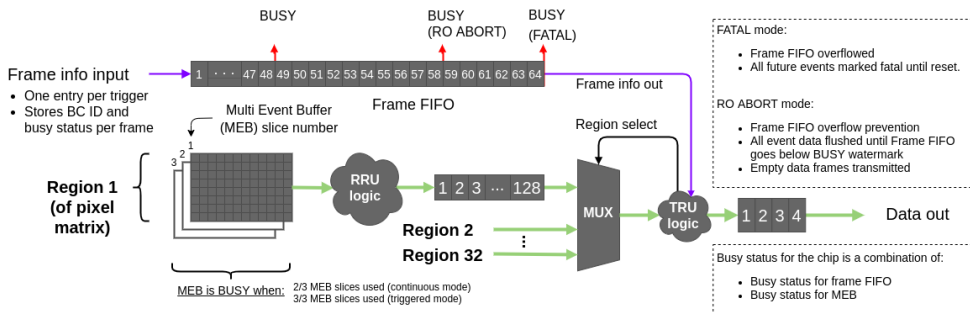


Figure 1: Readout logic in the ALPIDE chip

A Region Readout Unit (RRU) reads data from one region of the pixel matrix into a region FIFO. Informations for each trigger, such as Bunch Crossing (BC) ID and busy status, are stored in the frame FIFO, and data from different events are delimited in the region FIFOs. The last step of the readout before encoding and transmission is performed by the Top Readout Unit (TRU), which multiplexes between the RRU's in a round-robin fashion and performs framing of event data with the informations from the frame FIFO.

1.2 ITS Readout Electronics

Each RU is responsible for the readout of one staff. The staves come in three different configurations, and the RU was designed for use with any of the three. There are 9 ALPIDE chips in the IB staves, and each chip has a dedicated 960 Mbps¹ data link. In the OB staves, which come in two configurations, one 320 Mbps¹ data link is shared by seven chips. In the first two layers of the OB, there are 4 x 28 chips and 16 data links per staff. And in the final

¹The 8b10 encoded bit rate is 1200 Mbps for the IB, and 400 Mbps for the OB. The corresponding unencoded bit rate is 960 Mbps and 320 Mbps, for the IB and OB, respectively.

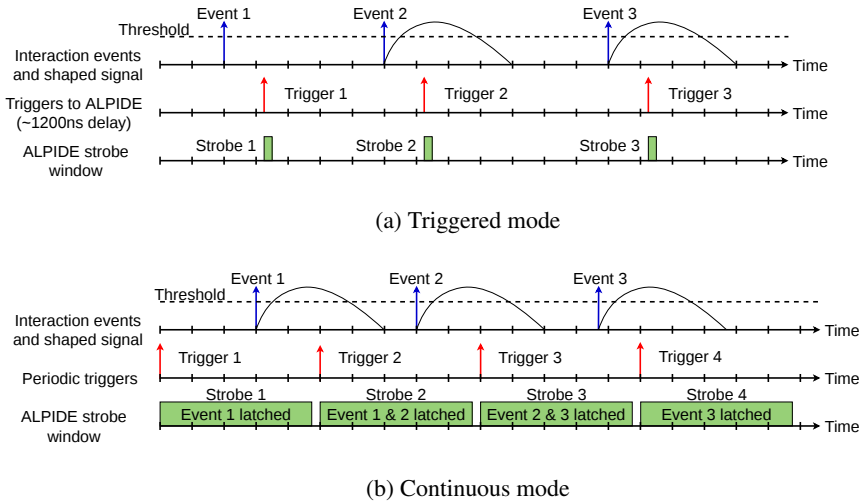


Figure 2: Trigger modes for the ALPIDE

two layers, there are 4×49 chips and 28 data links. Consequently, the RUs are used in very different configurations; in the IB, the link count is low, but data rates are high due to the proximity to the interaction point; in the OB, the link count is high, but the data rate per link is rather low.

Data from the sensor chips are aggregated in the RU, which has three 3.2 Gbps optical GBT links dedicated to data transfer. The capacity of these links must not be exceeded.

2 Simulation Model of ALPIDE and ITS

The simulation model for the system was written in C++ using the SystemC framework. The model of the ALPIDE chip in the simulation was designed to be simple enough that it would provide substantially faster simulation than the full HDL implementation, while accurately modeling the behavior of the readout logic in the chip. To that end, the model has an accurate implementation of the Finite State Machines (FSM) and parts of the internal readout logic shown in Fig. 1, specifically the Region Readout Unit (RRU) and Top Readout Unit (TRU). Other aspects of the readout logic were simplified to mimic the behavior of the ALPIDE chip. Instead of analog pulse shaping, the pixel hits are modeled as a rectangular pulse, with a fixed dead time and active time². Features that were not relevant for the simulation, such as the 8b10 encoding, were omitted entirely.




Figure 3: Overview of SystemC simulation model for ITS

²Time over threshold

Figure 3 shows an overview of the simulation model. A pool of Monte Carlo (MC) events is used as input data to the model. Quantities such as the number of triggers accepted and rejected by the chips, busy events, and data word counts are continuously monitored and stored to file at the end of the simulation. These data are later analyzed to calculate the readout efficiency, data rates, and other performance figures for a given simulation setup.

2.1 Monte Carlo Input Data

The MC events used by the SystemC simulation were generated using an MC simulation test bench for the ITS upgrade, which is part of the AliRoot framework [4]. The test bench simulates the physics for pp and Pb–Pb events³ in ITS, and the resulting hit data for each event is stored to file and later used as input to the SystemC simulation.

The MC event pool generated for the SystemC simulations consists of 10 000 Pb–Pb events, along with data for the Quantum Electrodynamics (QED) background⁴. The QED background is inputted to the simulations continuously, irrespective of the Pb–Pb interactions. The pp event pool consists of 100 000 events.

Pixel noise is not included in the MC data and is not part of the simulation model. Recent measurements have revealed that the fake-hit rate of the ALPIDE is on the order of 1.0×10^{-9} pixels/frame [5], depending on threshold settings. Since the ALPIDE has around a half million pixels, this amounts to one additional pixel firing every 2000 frames or so, for each chip. This is practically negligible in the context of our simulations and was consequently omitted in the simulations.

2.2 Simulation Setup

The simulation setup is summarized in table 1. The simulations were run at the event rates listed in the table, for Pb–Pb and pp interactions. Note that the ALICE Run-3 requirement for pp is 200 kHz, but this represents such a light load on the detector that it was omitted from the simulations.

Table 1: Simulation setup for Pb–Pb and pp simulations

Parameter	Pb–Pb	pp
Event rates [kHz]	50 ^a , 100 ⁱ , 150 ^b , 200 ^b	400 ⁱ , 1000 ^b , 2000 ^b , 5000 ^b
Number of events simulated	100 000	1 000 000
Strobe length (min-bias trigger) [ns]	100	100
Strobe length (periodic trigger) [ns]	4 900, 9 900, 19 900	4 900, 9 900, 19 900
Period (periodic trigger) [ns]	5 000, 10 000, 20 000	5 000, 10 000, 20 000

^a ALICE Run-3 requirement ⁱ ITS specification ^b Beyond specification

For each of the event rates that are listed, one simulation was run using minimum-bias triggers with the specified strobe length. Simulations were also run with periodic triggers with the corresponding periods and strobe lengths. The simulations with periodic triggers were run with the chips in both triggered and continuous mode, to evaluate the performance of the flushing mechanism in the continuous mode. The minimum-bias simulations were run

³The AliRoot simulation used Pythia for pp event generation and Hijing for Pb–Pb.

⁴The QED background is due to electromagnetic interactions and pair-production in ultra-peripheral Pb–Pb collisions. These interactions occur on every bunch crossing regardless of whether there are hadronic collisions or not.

with the chips in triggered mode only. A total of 56 simulations were run to cover the different combinations of parameters listed in the table. Because of the large amount of simulations, the number of simulated collisions was limited to 100 000 for Pb–Pb, and 1 000 000 for pp, and only one stave was simulated per layer (a total of 643 ALPIDE chips per simulation).

3 Simulation Results

The ITS specifications for the interaction rate are 100 kHz in Pb–Pb and 400 kHz in pp, and Figs. 4a and 4b show the average data rate per ALPIDE data link under these conditions.

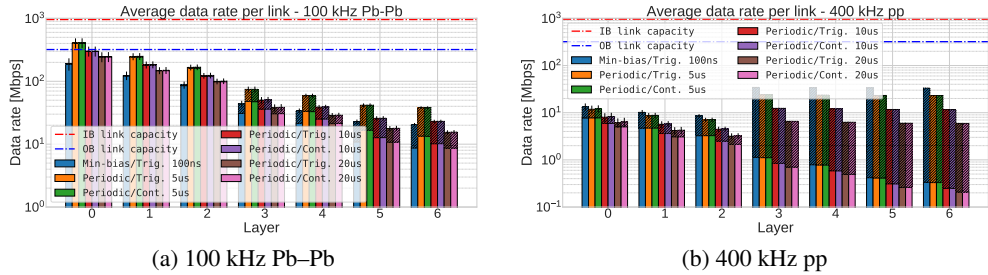


Figure 4: Average data rate per link. Note: y-axis range differs between the plots.

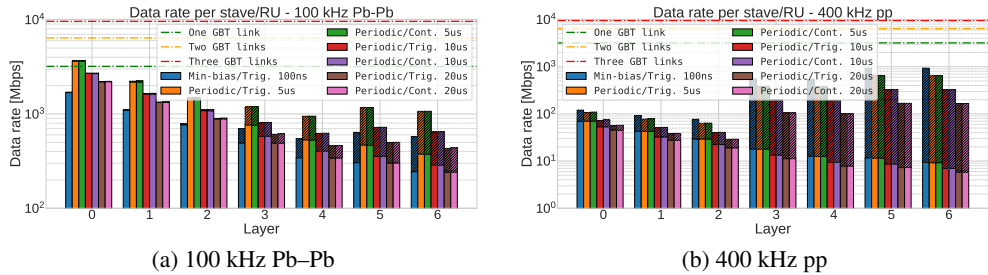


Figure 5: Total data rate per stave/RU. Note: y-axis range differs between the plots.

The data rate is shown versus layer for the combinations of trigger and strobe parameters that were simulated. The legend of the figures indicate trigger type first (minimum-bias or periodic), chip mode (triggered or continuous), and finally strobe length and trigger period in the case of periodic triggers. This applies to all of the figures in this paper.

A stave covers a large range of pseudorapidity, and occupancy is not constant over the whole range. Since the data rate is averaged over all links in the stave, the difference in occupancy leads to the relatively large error bars shown for the inner layers. The solid area of the bars corresponds to data that is associated with actual pixel hits. The shaded area represents other types of data words and consists predominantly of the protocol overhead associated with each readout frame. This overhead is proportional to the trigger rate and the number of chips per link, and is consequently more significant for the links in the outer barrel. In fact, for 400 kHz pp there is more protocol overhead than actual data for the outermost layers. The capacity of the ALPIDE data links is 960 Mbps for the first three layers (IB), and 320 Mbps for the last four layers (OB), and this is also indicated in the figures. The average data rate is more than an order of magnitude higher for Pb–Pb than it is for pp, but it is still well within the capacity of the links for all the simulations in Fig. 4. However, a larger margin reduces the chance of FIFOs and event buffers filling up, which should make

event data loss less likely. The worst case is in the innermost layer where the occupancy is the highest, around 40% of the link capacity is utilized here for 100 kHz Pb–Pb.

The total data rate for the one stave (or RU) that was simulated is shown in Fig. 5. As mentioned before, each RU has up to three 3.2 Gbps GBT links, and the link capacity of one, two, and all three GBT links is indicated in the figures. With the exception of 100 kHz Pb–Pb with 5 μ s strobe, the data rate is within the capacity of one GBT link in these simulations. Because of the large number of links in the outer barrel, the relative difference in data rate between the layers is not as profound in Fig. 5 as it was for individual links in Fig. 4.

A full comparison of data rates per stave for all the simulations is shown in Figs. 6a and 6b for Pb–Pb and pp, respectively. These figures have been limited to layer 0 since that is the most critical layer in all of the simulations.

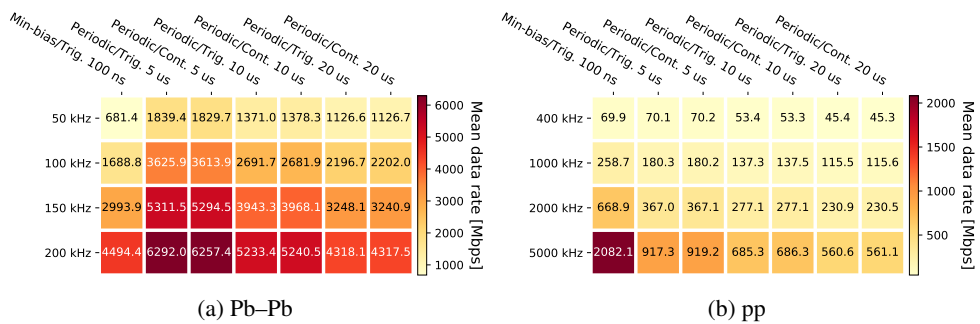


Figure 6: Data rates per stave/RU in layer 0 for all Pb–Pb and pp simulations.

The readout efficiency in terms of readout frames, at 100 kHz and 200 kHz Pb–Pb, is shown in Figs. 7a and 7b. At 100 kHz Pb–Pb, the ITS specification, the efficiency is near 100% for all combinations of trigger modes and other parameters. However, the efficiency is a bit worse for certain parameter sets, but only by a small fraction of a percent. For pp the efficiency is approximately 100% regardless of the chosen parameters, even at 5 MHz event rate. Efficiency figures for pp are hence not included. Looking at the efficiency for 200 kHz Pb–Pb in Fig. 7b, it is clear that using periodic triggers with a short period of 5 μ s is the least efficient. Since the pulse shaping time is also around 5 μ s, most events coincide with two strobe windows, which means that events are essentially “double-sampled” with a 5 μ s strobe length. This can also be observed in Figs. 4a and 5a; the data rate for simulations with a 5 μ s strobe is around twice as large as for simulations with minimum-bias trigger and short strobe (100 ns). This double-sampling effect is also present for periodic triggers with longer strobes, but it becomes less likely as the strobe length increases.

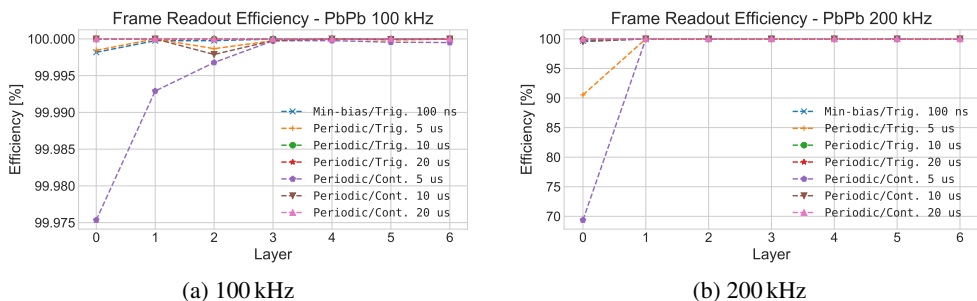


Figure 7: Frame readout efficiency in Pb–Pb simulations

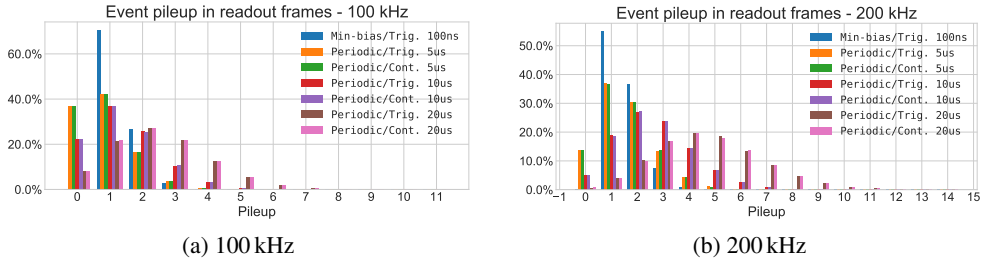


Figure 8: Pileup of events in readout frames in Pb–Pb simulations

Pileup is also shown for 100 kHz and 200 kHz Pb–Pb in Figs. 8a and 8b. A pileup of zero on the x-axis indicates empty frames⁵; a pileup of one indicates exactly one interaction event in a readout frame; pileup larger than one indicates more than one event per readout frame, i.e. there is a pileup of events in that frame. With minimum-bias triggers, there are no empty frames⁶ since only interaction events are triggered on, and minimum-bias triggers have smaller pileup values⁷. Pileup values larger than three are rarely observed for minimum-bias triggers or periodic triggers with 5 μ s period. But with longer strobe lengths, it becomes more prevalent.

4 Conclusions

Several combinations of running modes and parameters were simulated, which enabled estimations of readout efficiencies, data rates, and expected pileup per frame. Interaction rates of 50 kHz for Pb–Pb and 200 kHz pp are planned for ALICE in Run-3, using periodic triggers, and possibly with higher rates for pp. As a baseline, the strobe lengths will be inverse of the interaction rate, i.e., 20 μ s for Pb–Pb and 5 μ s for pp, but this will be optimized when the detector is commissioned with beams. According to our results, the upgraded ITS detector is capable of operating with very high efficiency in these configurations, even with shorter strobe lengths in Pb–Pb and at much higher interaction rates in pp. It should also be possible to operate at rates beyond the specifications for Pb–Pb runs, but in this case, it may be necessary to use a longer strobe to achieve high efficiency. However, a longer strobe comes at the expense of a higher pileup of events in readout frames, which makes reconstruction of events more challenging. Compared to simulations performed at an earlier point in the development of the ITS upgrade [6], the results presented here are a bit more optimistic. This is primarily because those simulations included pixel noise, which was omitted in our simulations. The fake-hit rate was expected to be 3-4 order of magnitude higher when the earlier simulations were performed, and made a significant impact on those results.

As a final remark, the simulation model has also been used to simulate data rates and readout efficiency for the proton-CT project at the University of Bergen (UiB) [7] and the planned Forward Calorimeter (FoCal) in ALICE.

⁵Empty in this context means that no interaction events were associated with the frame.

⁶The readout frames may still be empty for individual chips in terms of pixel hits, but the strobe for the frame was still associated with an actual interaction event.

⁷The figures may falsely give the impression that a pileup of two is higher with minimum-bias triggers, but this is only because zero pileup is included in the figures.

References

- [1] The ALICE Collaboration, *Journal of Physics G: Nuclear and Particle Physics* **41** (2014)
- [2] The ALICE Collaboration, *Nucl. Instrum. Methods Phys. Res., A* **824**, 434 (2016)
- [3] S.V. Nesbo, J. Alme, M. Bonora, P. Giubilato, H. Helstrup, S. Hristozkov, G. Aglieri Rinella, D. Röhrich, J. Schambach, R. Shahoyan et al. (ALICE ITS), *PoS TWEPP-17*, 147. 5 p (2017)
- [4] The ALICE Collaboration, *AliRoot github repository - itsuTestBench* (2019), [Online; accessed October 28, 2019], <https://github.com/alisw/AliRoot/tree/master/ITSMFT/ITS/itsuTestBench>
- [5] S. Kushpil, F. Krizek, A. Isakov, *IEEE Transactions on Nuclear Science* **66**, 2319 (2019)
- [6] A. Szczepankiewicz, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **824**, 465 (2016)
- [7] O.S. Groettvik, J. Alme, R. Barthel, T. Bodova, V. Borshchov, A. van den Brink, V. Eikeland, A. Herland, N. Van Der Kolk, S. Voigt Nesbø et al., *PoS TWEPP2019*, 090 (2020)