HK 54: Computing II

Time: Wednesday 17:30–19:00 Location: HBR 19: C 103

HK 54.1 Wed 17:30 HBR 19: C 103

Real-time calibrations for future detectors at FAIR — •Valentin Kladov[1,2], Johan Messchendorp[2], and James Ritman[1,2,3] — [1]Ruhr-Universität Bochum — [2]GSI Helmholtzzentrum fur Schwerionenforschung GmbH — [3]Forschungszentrum Jülich

The online data processing of the next generation of experiments, such as those conducted at FAIR, requires a reliable reconstruction of event topologies and, therefore, will depend heavily on in-situ calibration procedures. In this study we present a neural network-based tool designed to provide real-time predictions of calibration constants, which rely on continuously available environmental data. To enhance regularization, we incorporate information about previous environmental states into the Long Short-Term Memory (LSTM) architecture. LSTM is combined with Graph Convolutions to facilitate predictions across multiple channels simultaneously and to account for correlations between the channels. A proof-of-principle of this approach has been demonstrated using data from the Drift Chambers of the HADES detector obtained during the February 2022 experiment. Our method demonstrated the ability to provide fast and stable calibration predictions with a precision comparable to that obtained using traditional offline, time-consuming approaches. We plan to apply the proposed methodology in a real-time experimental setting during the next HADES beam time scheduled for Feb-Mar 2024.

HK 54.2 Wed 17:45 HBR 19: C 103

ASAPO: A high-speed streaming framework to support an automated data-processing pipeline. — •Mikhail Karnevskiy — DESY, Hamburg, Germany

Modern high-speed and high-resolution detectors produce a significant data rate, thereby increasing the need for online data processing and data reduction. A well-developed common data processing framework, provided as a service, minimizes the cost of implementing different use cases, enables more efficient scientific work due to a higher degree of automation, and enhances the reliability of the overall system.

ASAPO is a high-performance streaming framework actively developed at DESY to enable online and offline data processing using TCP/IP and RDMA over Ethernet and Infiniband. It efficiently facilitates high-bandwidth communication between detectors, the storage system, and independent analysis processes. User-friendly interfaces are available for C/C++, Python on all major platforms. A high-level Python library reduces boilerplate code when writing independent analysis workers, which can be combined into complex pipelines. ASAPO supports users with automatic retransfer, trivial parallelization on a per-image basis, multi-module detectors, and web-based monitoring.

Several experimental facilities at Petra III already use ASAPO for various data-processing pipelines. Involved algorithms include azimuthal integration of X-ray scattering data, peak finding, and indexing of diffraction patterns.

HK 54.3 Wed 18:00 HBR 19: C 103

The Event Processing Nodes: technical operation and performance of the ALICE GPU-based processing farm and computing model for synchronous and asynchronous data reconstruction — •Federico Ronchetti for the ALICE Germany-Collaboration — CERN, Esplanade des Particules 1, 1211 Geneva, CH

The Large Hadron Collider (LHC) returned to operation on July 5^{th}, 2022. During LHC Long Shutdown 2 (2019-2021), the ALICE detector underwent a major upgrade that increased the sustainable hadronic rate from 1 to 50 kHz for Pb-Pb collisions in continuous readout mode.

The improved detector performance and the change of the data taking paradigm required the operation of a completely new computing model which merges online (synchronous) and offline (asynchronous) data processing into a single software framework.

After a short introduction of the ALICE upgrade an overview of the current ALICE computing model will be given together with a technical description with the supporting hardware facility which makes extensive use of GPU computing since due to the increased data volumes, storing all the produced raw data is infeasible. The ALICE computing model has the unique feature to be able to exploit GPU processing also for offline data reconstruction hence energy efficiency considerations on the wide-spread use of GPUs will be formulated.

HK 54.4 Wed 18:15 HBR 19: C 103

Modern C++ with SYCL as Multi Paradigm Programming Language for FPGA-Based Detector Readout — •Thomas Janson and Udo Kebschull for the ALICE Germany-Collaboration — IRI, Goethe-Universität Frankfurt am Main, Max-von-Laue-Straße 12, 60438 Frankfurt am Main, Germany

Recent developments in high-level synthesis for FPGA targets enable new methods of implementing detector readout in high-energy physics. In this talk, different methods are shown to develop complex algorithms using high-level synthesis and the Intel oneAPI framework based on SYCL2020, which can be used to develop, test, and implement complex algorithms. SYCL is a programming model using C++ for heterogeneous hardware like GPUs, CPUs, and FPGAs. SYCL inherits many of the Modern C++ features, like generic programming with templates, lambda expressions for functional programming, and many more. We evaluate the usability of Modern C++ features with SYCL and the Intel oneAPI FPGA IP Authoring Flow for FPGAs. First experiences and results are shown and discussed.

HK 54.5 Wed 18:30 HBR 19: C 103

An online GPU hit finder for the STS detector in CBM — •Felix Weiglhofer for the CBM-Collaboration — Frankfurt Institute for Advanced Studies, Goethe University, Frankfurt, Germany

The CBM experiment is expected to run with a data rate exceeding 500 GB/s after averaging. At this rate storing raw detector data is not feasible and an efficient online reconstruction is required instead. GPUs have become essential for HPC workloads. Higher memory bandwidth and parallelism can provide significant speedups over traditional CPU applications. These properties also make them a promising target for the planned online processing in CBM.

We present an online hit finder for the STS detector capable of running on GPUs. The hit finder consists of four steps using STS Digis (timestamped detector messages) as input. Digis are sorted by sensor, within each sensor, they are sorted by channel and timestamp. Neighboring Digis are combined into clusters. Finally, after time sorting clusters on each sensor are combined into hits.

Each of those steps is trivially parallel across STS sensors or even sensor sides. To fully utilize GPU hardware, we modify the algorithms to be parallel on Digi or cluster level. This includes a custom implementation of parallel merge sort allowing full parallelism within GPU blocks.

Our implementation achieves speedup of 24 on mCBM data compared to the same code on a single CPU core. The exact achieved throughput will be shown during the presentation.

This work is supported by BMBF (05P21RFFC1).

HK 54.6 Wed 18:45 HBR 19: C 103

Data challenges at CBM - towards scalable workflows — •Andreas Redelbach for the CBM-Collaboration — Frankfurt Institute for Advanced Studies, Goethe University, Frankfurt, Germany

Operating the CBM experiment at interaction rates up to 10 MHz requires data reduction in real-time. This necessitates highly efficient online processing of measurements and the underlying algorithms. More specifically, since the free-streaming readout data are processed in software, the performance of the reconstruction algorithms is a critical issue. A promising option for acceleration is based on an efficient parallelization of data processing developed for both CPU and GPU architectures. A number of measures have been taken to minimize runtimes of reconstruction algorithms and to optimize the scaling of some time-critical workflows.

Using the mCBM full-system test setup at SIS18 allows testing of all relevant components connected to study all processing steps. It is interesting to note that progress has been achieved in particular using test data from mCBM beamtimes. In this contribution, some of the concepts and recent progress towards high throughput processing in the CBM reconstruction chain are summarized.

This work is supported by BMBF (05P21RFFC1).