FPGA-based real-time data processing for accelerating reconstruction at LHCb

Friday 24 September 2021 14:00 (16 minutes)

In Run-3 beginning in 2022, the LHCb software trigger will start reconstructing events at the LHC average crossing rate of 30 MHz. Within the upgraded DAQ system, LHCb established a testbed for new heterogeneous computing solutions for real-time event reconstruction, in view of future runs at even higher luminosities.

One such solution is a highly-parallelized custom tracking processor ("Artificial Retina"), implemented in state of the art FPGA devices connected by fast serial links.

We describe the status of the development of a life-size demonstrator system for the reconstruction of pixel tracking detectors, that will run on real data during Run-3.

Summary (500 words)

With the slowdown of Moore's law, HEP experiments are looking at heterogeneous computing solutions as a way to face ever-increasing data flows and complexity. LHCb is on the frontier of these developments due to its specific physics needs, calling for the full software reconstruction of events in real-time at the LHC average rate of 30 MHz (40 Tb/s), already in the next physics run starting in 2022. LHCb has already adopted a GPU-based solution for HLT1 for the next run, and is further researching solution for its future Upgrade-II, with a significant increase of luminosity by a factor 5÷10. To this purpose, a coprocessor testbed has been established, to allow parasitical testing of new processing solutions in realistic DAQ conditions during the 2022 run.

One such solution under development is a highly-parallelized custom tracking processor ("Artificial Retina"). The "Artificial Retina"architecture takes advantage of FPGAs parallel computational capabilities, by distributing the processing of each event over an array of FPGA cards, interconnected by a high-bandwidth (~15 Tb/s) optical network. This is expected to allow operation in real-time at the full LHC collision rate, with no need for time-multiplexing or extra buffering thanks to its brief latency (<1 μs).

This level of performance has never been attained before in a complex track-reconstruction task, and achieving it opens the door to early reconstruction of track primitives transparently during detector readout. These data can be used as seeds by the High Level Trigger (HLT1/HLT2) to find tracks and perform trigger decisions with much lower computational effort than possible by starting from the raw detector data. This can free an important fraction of computing power of the conventional event-processing farm, allowing more powerful and faster reconstruction at higher luminosities than otherwise possible. Implementation of this technology could enhance the physics potential of LHCb already from the following physics run (Run-4), by expanding its trigger capability with the inclusion of long-lived tracks in the HLT1 decision.

In this talk we describe the status of commissioning, and the performance of the first realistic prototype of this system. The prototype is configured to process data from the VELO pixel detector, that is compact enough to be processed by a small system (~40 boards), and yet accounts for a significant fraction (~1/2) of the HLT Level 1 (HLT1) computational load. The 2-dimensional pixel geometry of the VELO requires a demanding cluster-finding task to be performed at the 30 MHz event rate as a preliminary step, before performing track reconstruction proper. This has been addressed by a derivation of the same "Artificial Retina"approach, implemented in firmware within the already-existing readout boards, and will be a default part of the standard LHCb reconstruction already in Run-3. The track pattern-recognition stage is instead implemented in independent commercial boards, equipped with Stratix-10 FPGAs (2.8 MLE) and 16 * 28 Gbps links each, connected in a network topology of 4 cyclically-connected full-mesh nets, each carrying a total ~4 Tb/s flow of point-to-point connections.

Authors: LAZZARI, Federico (Universita di Siena & INFN Pisa (IT)); CONTU, Andrea (INFN); BASSI, Giovanni (SNS & INFN Pisa (IT)); PUNZI, Giovanni (Universita & INFN Pisa (IT)); TUCI, Giulia (Universita & INFN Pisa (IT)); GIAMBASTIANI, Luca (Universita e INFN, Padova (IT)); STICCHI, Marco (Universita & INFN Pisa

(IT)); MORELLO, Michael J. (SNS and INFN-Pisa (IT)); DORIGO, Mirco (INFN Trieste); FANTECHI, Riccardo (Universita & INFN Pisa (IT)); BALDINI, Wander (Universita e INFN, Ferrara (IT))

Presenter: LAZZARI, Federico (Universita di Siena & INFN Pisa (IT))

Session Classification: Trigger

Track Classification: Trigger