# A First-level Event Selector for the CBM experiment

*J. de Cuveland and V. Lindenstruth*
FIAS Frankfurt Institute for Advanced Studies, Goethe-Universität, Frankfurt am Main, Germany

The First-level Event Selector (FLES) is the central physics selection system in CBM. It receives all hits and performs online event selection on the 1 TByte/s input data stream. The event selection process requires high-throughput event building and full event reconstruction using fast, vectorized track reconstruction algorithms. The current FLES architecture foresees a scalable supercomputer. To achieve the high throughput and computation efficiency, all available computing devices will have to be used, in particular FPGAs at the first stage of the system and heterogeneous many-core architectures such as GPUs for efficient track reconstruction. A high-throughput network infrastructure and flow control in the system are other key aspects.

The FLES system will contain approximately 1000 input nodes with fast PCIe interfaces receiving data from the detector at a total rate of 1 TByte/s. Online event analysis will be performed by processing nodes estimated to require in total approximately 60000 cores. A high-throughput network infrastructure will connect all nodes and enables interval building at the full input data rate. The FLES architecture is summarized in Fig. 1.

The interface between the custom data acquisition network and the FLES PC farm is implemented as an add-on card to the host PC (see Fig. 2). This add-on card has to feature high-bandwidth links both from the readout electronics to the host PC and a large and high-bandwidth buffer memory. The current baseline assumption is to use a custom FPGA-based add-on card with a PCIe interface at the maximum achievable PCIe data rate. The FPGA design implements the data acquisition protocol for receiving data, merges input links, performs memory management and controls the PCIe transfer. To allow for efficient time-indexed data access, it additionally analyzes the incoming data's time stamp information and builds corresponding look-up tables.

While there are CBM running scenarios that require full event building prior to selection, other selection scenarios are feasible with a reduced load on the event building network. This is especially important during the early building phases of the experiment. Therefore, the FLES implements

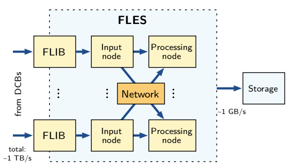


Figure 1: Architecture of the FLES

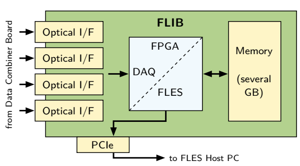


Figure 2: The FLES interface board (FLIB)

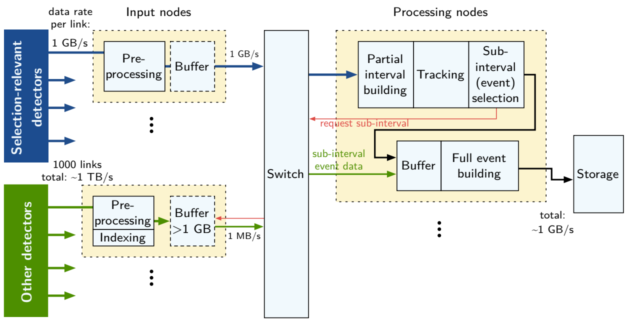


Figure 3: Data flow in the First-level Event Selector

a two-staged event building and event selection scheme. The main data flow in the FLES is illustrated in Fig. 3. Data from all detectors that contribute to the initial stage of event selection are forwarded to the interval-building network between input nodes and processing nodes. The goal here is to concentrate data that have a similar time stamp from all detector links in a single processing node. For that reason, the input nodes package all input data corresponding to fixed time intervals (containing, e. g., 1000 events). Each processing node receives data for a different time interval. This interval building corresponds to the event building process commonly used in other experiments. To optimize for maximum throughput, several routing concepts are currently under investigation. Using a network technology such as InfiniBand QDR, FLES implements high-throughput event building at the full data rate of 1 TByte/s.

For each time interval, the associated processing node performs the required feature extraction and online analysis. At an average event rate of 10 MHz, discarding all events that cannot be cleanly separated in time would result in a considerably reduced efficiency. To avoid this, the track reconstruction algorithm applied by the processing node directly matches tracks not only in the three spacial dimensions, but also in time as an additional dimension. After interesting events have been identified, the two-staged event building process allows the processing node to request a corresponding sub-interval from additional detectors not participating in the initial interval building. The challenge in online event reconstruction is the high event rate of 10 MHz in combination with non-trivial trigger criteria requiring partial or full event reconstruction. To fully exploit the hardware potential of heterogeneous many-core systems, reconstruction codes have to be rewritten with respect to efficient parallel processing. To efficiently utilize resources, the FLES building network is also used as an interface to the mass storage system.