Status of CBM First-level Event Selector Prototype Developments*

J. de Cuveland[1], *D. Hutter*[1], *and V. Lindenstruth*[1]
[1]FIAS Frankfurt Institute for Advanced Studies, Goethe-Universität Frankfurt am Main, Germany

The First-level Event Selector (FLES) is the central event selection system in the CBM experiment. Its task is to select data for storage based on online analyses including a complete event reconstruction. To do so, the FLES timeslice building has to combine data from all input links to time intervals and distribute them via a high-performance network to the compute nodes. Data rates at this point are expected to exceeds 1 TByte/s. It has proven practical to use an InfiniBand network for data transfer between the FLES nodes. Even so, further investigations are needed to design the final system and develop the required software algorithms.

Micro-FLES Cluster

To provide a small scale, highly customizable platform for these studies, the *Micro-FLES* cluster was installed at GSI (see Fig. 1). Eight identical compute nodes provide a total of 192 logic cores and 512 GB memory. While consuming only 1 U of installation space, the nodes still provide PCIe 3.0 expansion slots for up to 2 FLES Interface Boards (FLIB) and 3 GPU cards. This allows additional tests of the full data transport chain and reconstruction algorithms in the future. In addition to the compute nodes, one head node provides infrastructure services such as storage and boot images. Although there is local storage in the compute nodes for test purposes, it is currently not used for operating the cluster.

For timeslice building the cluster is equipped with a state-of-the-art InfiniBand FDR network. Each node offers two 56 Gbit/s 4x FDR ports currently connected to one core switch. The availability of two ports per node easily allows the study of more complex network topologies.

Complementarily, the development of a timeslice-building test software based on InfiniBand Verbs has started. Especially questions regarding data structures, buffer management, and network scheduling are addressed. Preliminary tests using the same data structures as intended for the final setup have shown promising results.

FLES Interface Board (FLIB)

In the process of upgrading the CBMnet-based prototype read-out chain, a replacement for the currently used PC interface board (namely the ABB) is needed. For this purpose the FLIB prototype board derived from the ALICE C-RORC [1] has been chosen. Based on a Xilinx Virtex-6 FPGA, it features an eight-lane PCIe Gen 2 interface, up to 12 optical links and two DDR3 memory sockets.

* Work supported by BMBF (05P12RFFCP) and HIC for FAIR

Figure 1: The *Micro-FLES* cluster installed at GSI

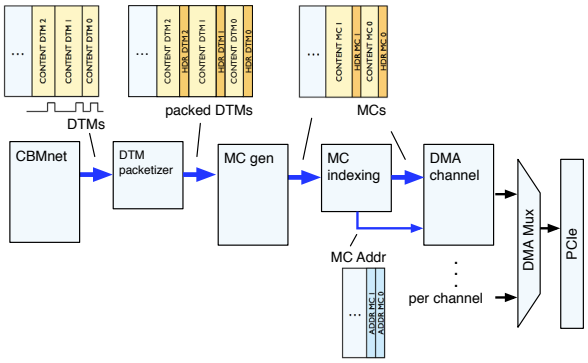

Figure 2: Schematic view of the FLIB read-out data path

In contrast to the final system, early prototype setups will not support the creation of microslice containers (MC) (as described in [2]). Hence, a firmware was developed that is capable of reading out raw CBMnet messages as delivered from the CBM front-end electronics.

The corresponding data path is shown in Figure 2. In the first step, incoming CBMnet messages are preprocessed in hardware to facilitate subsequent software processing. After preprocessing, a separate module packages the incoming data into simplified MCs, whose reduced data content requirements (in comparison to ordinary MCs) support not yet fully synchronized detector electronics. Finally, an address index table is created and MCs are written to the PC host memory via the PCIe interface. The chosen partitioning of the design and the creation of simplified MCs enables reuse of developed hardware and software components when migrating to a fully microslice-based read-out chain.

References

[1] H. Engel *et al*, "ALICE C-RORC as CBM FLES Interface Board Prototype", this report

[2] J. de Cuveland *et al*, "CBM First-level Event Selector Dataflow Architecture and Microslice Concept", GSI Sci. Report 2011, PHN-NQM-CBM-41, p 63.