Nuclear Instruments and Methods in Physics Research A (xxxx) xxxx–xxxx

Contents lists available at ScienceDirect

Nuclear Instruments and Methods in Physics Research A

journal homepage: www.elsevier.com/locate/nima

Real-time track-less Cherenkov ring fitting trigger system based on Graphics Processing Units

R. Ammendola[e], A. Biagioni[d], S. Chiozzi[a], P. Cretaro[d], A. Cotta Ramusino[a], S. Di Lorenzo[c,h], R. Fantechi[c], M. Fiorini[a,g,*], O. Frezza[d], A. Gianoli[a], G. Lamanna[c], F. Lo Cicero[d], A. Lonardo[d], M. Martinelli[d], I. Neri[a,g], P.S. Paolucci[d], E. Pastorelli[d], R. Piandani[c], M. Piccini[b], L. Pontisso[c], D. Rossetti[f], F. Simula[d], M. Sozzi[c,h], P. Vicini[d]

[a] *INFN Sezione di Ferrara, Ferrara, Italy*
[b] *INFN Sezione di Perugia, Perugia, Italy*
[c] *INFN Sezione di Pisa, Pisa, Italy*
[d] *INFN Sezione di Roma, Roma, Italy*
[e] *INFN Sezione di Roma Tor Vergata, Roma, Italy*
[f] *NVIDIA Corp., Santa Clara, CA, USA*
[g] *Università degli Studi di Ferrara, Ferrara, Italy*
[h] *Università degli Studi di Pisa, Pisa, Italy*

ARTICLE INFO

Keywords:
Ring Imaging Cherenkov detectors
Trigger system
Graphics Processing Units
Real-time selection

ABSTRACT

The parallel computing power of commercial Graphics Processing Units (GPUs) is exploited to perform real-time ring fitting at the lowest trigger level using information coming from the Ring Imaging Cherenkov (RICH) detector of the NA62 experiment at CERN. To this purpose, direct GPU communication with a custom FPGA-based board has been used to reduce the data transmission latency. The GPU-based trigger system is currently integrated in the experimental setup of the RICH detector of the NA62 experiment, in order to reconstruct ring-shaped hit patterns. The ring-fitting algorithm running on GPU is fed with raw RICH data only, with no information coming from other detectors, and is able to provide more complex trigger primitives with respect to the simple photodetector hit multiplicity, resulting in a higher selection efficiency. The performance of the system for multi-ring Cherenkov online reconstruction obtained during the NA62 physics run is presented.

1. Introduction

Recent years witnessed the development of a new trend in information technology, i.e. General Purpose computing on Graphics Processing Units (GPGPU). GPUs are based on a massively parallel architecture in which thousands of cores process large blocks of data in parallel, which makes them more efficient than general-purpose CPUs. Despite being originally designed for 3D computer graphics, modern commercial GPUs are often exploited to efficiently perform scientific computations in several fields, e.g. high-energy physics (HEP) or medical imaging.

2. GPUs for HEP trigger systems

In HEP experiments the trigger system is a fundamental component that allows reducing the size of collected data to a manageable level for the data acquisition system.

In a typical multi-level trigger architecture, GPUs can be quite easily exploited in the higher (software) levels. The number of computing farm nodes can be reduced and the capability of the processing system can be improved without increasing the scale of the system itself, but at a cost of adapting the software framework and algorithms to a parallel architecture.

For example the ALICE collaboration is using GPUs in the high-level trigger system [1], while other CERN experiments like ATLAS [2], CMS [3] and LHCb [4] are investigating the advantage of using GPUs with respect to CPUs in order to cope with the increased computing power needs foreseen for the planned luminosity upgrades.

Low-level triggers (L0) could also take advantage from the use of GPUs but a careful assessment of their real-time performances is required. In fact, L0 trigger systems are designed to perform rough selections based on a sub-set of the available information, typically in a

* Corresponding author at: Università degli Studi di Ferrara and INFN Sezione di Ferrara, Ferrara, Italy.
E-mail address: fiorini@fe.infn.it (M. Fiorini).

http://dx.doi.org/10.1016/j.nima.2017.02.031
Received 18 November 2016; Received in revised form 30 January 2017; Accepted 12 February 2017

pipelined structure housed in custom electronics. In such systems, a very low latency is required due to the small buffer size in read-out electronics.

In the following we describe recent results from our project, which aims at using commercial GPUs for the production of trigger primitives reconstructing ring-shaped Cherenkov hit patterns, to be used at the lowest trigger level (L0) in the NA62 experiment.

3. Physics case: the NA62 RICH detector at CERN

The NA62 experiment at the CERN SPS aims at measuring the branching ratio of the ultra-rare decay of the charged kaon into a pion and a neutrino–antineutrino pair [5]. The expected value is $(8.4 \pm 1.0)\times10^{-11}$, therefore its measurement requires a kaon beam with very high intensity in order to collect sufficient statistics in a reasonable amount of time.

The NA62 experiment makes use of a three-level trigger system to reduce the 10 MHz incoming event rate of about three orders of magnitude before storage. The first trigger stage (L0) is a hardware synchronous system which achieves a rejection factor of 10 within the maximum latency of 1 ms. Two other trigger systems run at software level on a computer farm and perform further filtering and final event building.

The Ring Imaging Cherenkov (RICH) detector has been designed to separate pions from muons in the 15–35 GeV/c momentum range [6]. It consists of a 17 m long vessel, 3 m in diameter, filled with Ne at 1 atm, read out by 2000 photo-multiplier tubes (PMTs) [7]. According to the baseline design of the experiment, primitives from the RICH detector at L0 consist only of PMT multiplicities in a pre-defined time window.

A dedicated system for generating advanced trigger primitives for the RICH at L0 has been implemented on a commercial GPU and relies on the enhanced computation capabilities and high parallelisation available on such devices. A fast ring-fitting algorithm is fed with raw RICH data, with no track information from the spectrometer, while information on the particle speed and direction is provided for a more selective L0 trigger decision. This system was installed in parasitic mode during the 2015 NA62 experimental run and relies on direct GPU communication using a FPGA-based network interface card called NaNet that allows strong latency reduction. For the 2016 run the system has been upgraded with an improved version of NaNet equipped with a 10 Gigabit Ethernet link.

4. NaNet-1 and NaNet-10 boards

NaNet-1 is a low-latency 1 GbE NIC with GPUDirect capability developed at INFN Rome, integrated since the beginning of the project in the GPU-based low-level trigger system of the NA62 RICH detector [8,9]. Its design comes from the APEnet+ card logic [10] and the board supports a configurable number of different physical I/O links. The Distributed Network Processor (DNP) is the APEnet+ core logic, behaving as an off-loading engine for the computing node in performing inter-node communications [11]. NaNet-1 is able to exploit the GPUDirect peer-to-peer (P2P) capabilities of NVIDIA Fermi/Kepler GPUs equipping a hosting PC to directly inject into their memory an UDP input data stream from the detector front-end.

In order to overcome the rate limitation of NaNet-1 and collect data from 4 TEL62 boards [12,13], another generation of FPGA-based NIC has been developed at INFN Rome [14]. NaNet-10, which supports 10 GbE I/O channels and PCIe Gen3 (8 GB/s), was implemented on a Terasic board equipped with an Altera Stratix-V GX FPGA featuring four SFP+ cages. NaNet-10 uses a UDP transport protocol as the previous version, and is installed in the NA62 setup since 2016.

In Fig. 1, NaNet-10 and NaNet-1 latencies are compared within UDP datagram size range. NaNet-10 guarantees sub-μs hardware latency for buffers up to ~1 kByte in GPU/CPU and it reaches its

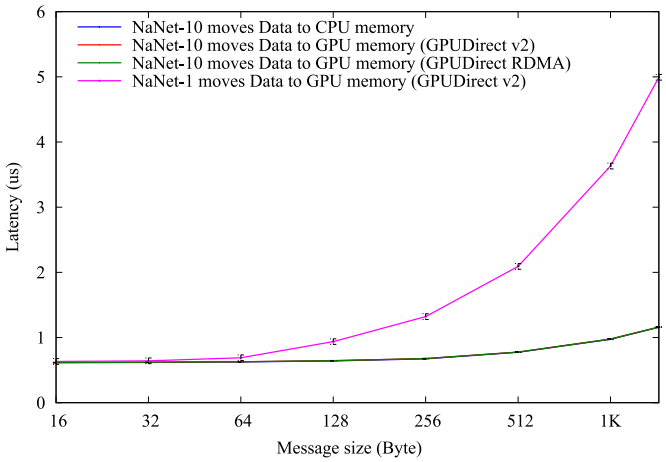

Fig. 1. NaNet-10 vs. NaNet-1 hardware latency. All curves corresponding to NaNet-10 overlap in this scale and only the green one is visible. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.).

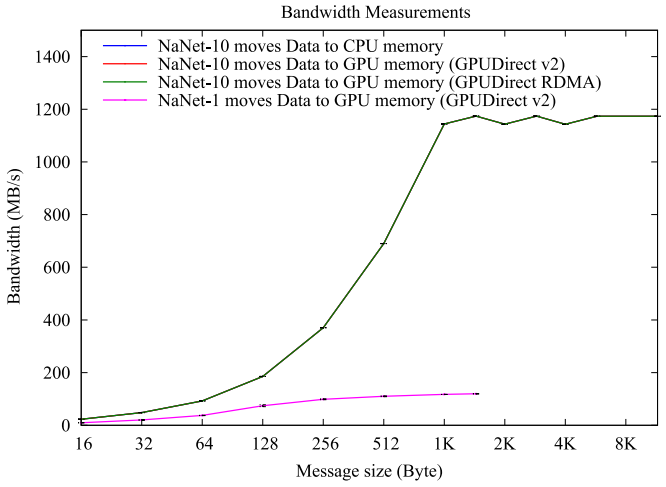

Fig. 2. NaNet-10 vs. NaNet-1 bandwidth. All curves corresponding to NaNet-10 overlap in this scale and only the green one is visible. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.).

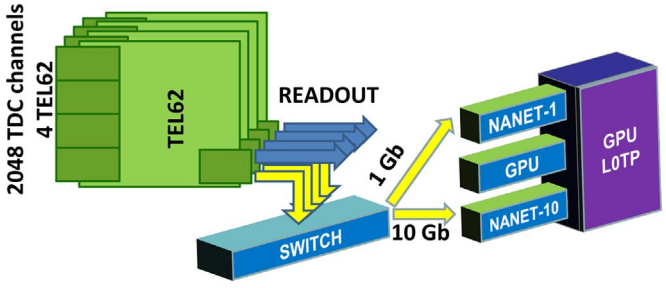

Fig. 3. Pictorial view of GPU-based trigger.

10 Gbps bandwidth peak capability, corresponding to about 1200 MB/s in Fig. 2, already at ~1 kByte size.

5. GPU-based L0 trigger system at the NA62 experiment

A GPU-based L0 trigger system is currently installed in the experimental setup of the RICH detector at NA62 experiment at CERN. The system consists of 4 GbE links to transfer data from the readout boards to the GPU (see Fig. 3). Each link gathers the data coming from ~500 PMTs. The main requirement is that the entire response latency comprising both communication and computation tasks must be less than 1 ms. Refined primitives coming from the GPU-