

GPUs for fast triggering and pattern matching at the CERN experiment NA62

Gianluca Lamanna, Gianmaria Collazuol and Marco Sozzi

Abstract—In high energy physics experiment the trigger system is crucial to reduce the quantity of data recorded on tape and the acquisition bandwidth requirements. This is particularly true in rare decays experiments. The NA62 experiment aims at measuring the Branching Ratio of $K^+ \rightarrow \pi^+ \nu \bar{\nu}$, predicted in the Standard Model (SM) at level of $\sim 10^{-10}$. In this paper we describe the idea to use the commercial video card processor (GPU) to construct a fast and effective trigger system, both in hardware and software level. Due to the use of off the shelf technology, in continuous development for other purposes, the architecture described would be easily exported to other experiments, to build a versatile and fully customizable trigger system.

Index Terms—Trigger, GPU, Real Time, rare decays, flavor physics

I. INTRODUCTION

THE search for physics beyond the SM effects is one of the main topics for the present high energy physics experiments. This search is carried out both with direct search for new particles in the big accelerators (energy frontier) and with rare decays studies on high luminosity beams (intensity frontier). The NA62 experiment, briefly described in the next session, will study an ultra-rare kaon decay in order to point out, eventually, the physics beyond the standard model contribution. The main requirements for such an experiment is to have a very good background rejection in order to decrease the systematics uncertainty and a very high beam intensity in order to collect enough statistics. In this concern the trigger system is a very crucial part for the experiment. The on-line reduction of interesting data has to be very effective in order to decrease the total bandwidth requirements for the data readout. At the same time the trigger and readout system have to be very reliable to decrease the losses due to inefficiency, dead time or malfunctioning at a very low level. In the “standard” version the trigger and readout system will be based on FPGAs, in the hardware levels, and on PCs in the software levels. In this work we propose to use video card processors (GPU) to implement real time trigger decisions in the first level of trigger and to increase the computing capability in the higher levels. As an illustration of the possibility given by the fast computations on GPUs we present an “on going” work to exploit the parallel structure of those chips to a pattern matching problems: find a ring in a Cherenkov detector. The aim of this application

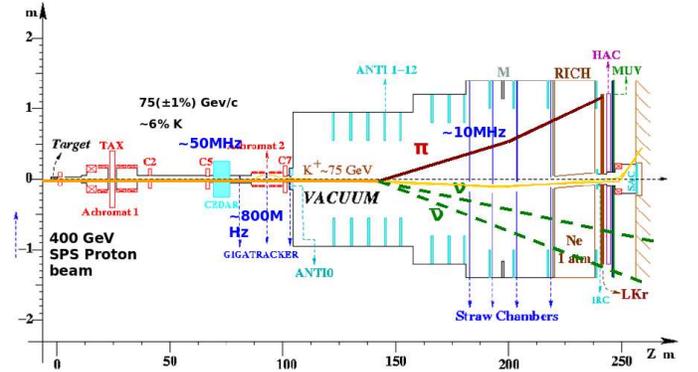


Fig. 1. Layout of the NA62 experiment.

is the definition of trigger primitives with high resolution in the hardware level of the where a fixed latency of 1 ms is required.

II. THE NA62 EXPERIMENT: DESCRIPTION AND PHYSICS GOAL

The NA62 experiment at the CERN SPS aims at measuring $O(100) K^+ \rightarrow \pi^+ \nu \bar{\nu}$ events in two years of data-taking. The theoretical cleanness of the Standard Model (SM) branching ratio (BR) predictions for this decay mode makes it very attractive both as a powerful test of the CKM paradigm and as a probe for new physics beyond the SM. Experimentally, the detection of this process is very difficult due to the smallness of the signal (in the SM the expected BR is at level of $\sim 8.5 \times 10^{-11}$) and the presence of a very sizable concurrent background, mainly from $K^+ \rightarrow \pi^+ \pi^0$ decays. The present measurement of this decay channel is based on 7 candidates collected by E949+E787 Brookhaven experiments [1] leading to a value of $BR = (1.47^{+1.30}_{-0.89} \times 10^{-10})$. NA62 is a fixed target experiment in which a charged hadrons beam, containing $\sim 6\%$ of kaons, will be produced from 400 GeV/c protons from the SPS accelerator. The kaon decays in flight will be studied in a fiducial region ~ 100 meters long, placed in vacuum in order to reduce secondary interactions. The decay’s products and the primary particle’s momentum will be measured with high resolution by STRAWS and GIGATRACKER spectrometers, in order to achieve good signal reconstruction and kinematic rejection. An efficient veto system for photons and charged particles (LAV, LK_r and SAC) and the PID system for primary particles and decay products (CEDAR, RICH and MUV), will guarantee the identification of decay modes not kinematically constrained. In fig.1 a schematic view of the experiment is shown. In order to collect the required number

Manuscript received November 20, 2009.

G.Lamanna is with Scuola Normale Superiore and INFN Pisa, Piazza dei Cavalieri 7, 56126 Pisa Italy e-mail: gianluca.lamanna@pi.infn.it

G.Collazuol is with INFN Pisa, Largo Pontecorvo 3, 56127 Pisa Italy e-mail: gianmaria.collazuol@cern.ch

M.Sozzi is with University of Pisa Physics department, Largo Pontecorvo 3, 56127 Pisa Italy e-mail: marco.sozzi@cern.ch

of events in a reasonable amount of time, a very intense hadron beam will be employed (3×10^{12} proton per SPS pulse will produce $\sim 5 \times 10^{12} K^+$ per year).

An efficient on-line selection of candidates represents a very important item for this experiment, because of the large reduction to be applied on data before tape recording. On the other hand a loss less data acquisition system is mandatory to avoid adding artificial detector inefficiencies, e.g. when vetoing background particles.

III. THE NA62 TRIGGER SYSTEM

The event rate on the main detectors of the order of 10 MHz will be reduced to the order of 10 kHz through a three level trigger system. The first level (L0) will be completely hardware while the further levels will be software, with the difference that the L1 will be devoted to the selection using single-detector reconstructed quantities and the L2 will apply selection on the full reconstructed event with good resolution. The L0 is based on the TELL1 board[2] (developed by EPFL for the LHCb experiment), a general purpose acquisition board in which 5 FPGAs allow a fully customizable configuration. Thanks to a total RAM of 384 MB, the TELL1 can store data in a first buffer stage, waiting for the trigger decision delivered to the board through the TTC [3] interface. The board houses up to 4 daughter boards with 4 high resolution (100 ps) TDCs with 512 channels in total, for the digitization of the information coming from the detectors. The TELL1 allows to directly construct, employing the FPGAs logic capability, trigger primitives used in the L0 trigger processor in order to take the L0 trigger decision, to reduce the rate for a factor of ten. The events of interest will go into the L1 for a software selection based on information coming from a single detector. Due to large amount of data to be processed in a reasonable time, the number of PC cores at the L1 will be quite large. After this level of selection the data from all the detectors will arrive at level 2 through a Gigabit switch; the event will be reconstruct in order to apply a selection based on the full kinematics and topology. The possibility to have fast and powerful computation in the software levels is a key point to reduce the dimensions of the on-line PC farms and the time of processing. In the other hands the possibility to define complex quantities (like invariant masses, momenta, etc.) at the L0, could be very useful in order to construct a more efficient on-line selection and a more effective rejection.

In the next sessions will be described the idea to use the massive parallel computation offered by the video card processors (GPU) in the trigger system, both in the software and hardware level, with the two aims of (1) processing in parallel many events and (2) speeding-up reconstruction algorithms for exploiting them at the lowest trigger levels. More details on the NA62 trigger system architecture can be found in [4].

IV. THE GPU (GRAPHICS PROCESSOR UNIT)

In the past years a big effort has been dedicated in the development of powerful processors exclusively devoted at graphical applications. The peculiar typology of problems

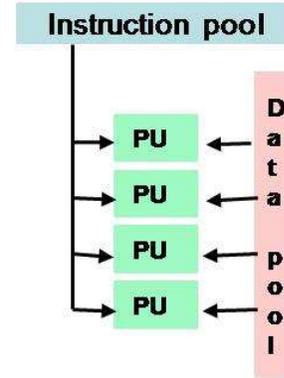


Fig. 2. In the SIMD architecture each Processing Unit (PU) runs the same program on different data set.

related to 3D, rendering and, in general, management of images, drove the development of architecture to the parallel SIMD (Single Instruction Multiple Data) scheme (fig.2), where a big number of transistors are employed for calculation and a relatively small for flow control and caching (fig.3). In the last years, some interest rises for the use of GPU in general purpose applications (GPGPU)[6] outside the field of imaging processing, with particular attention to the high performance computing for scientific issues. For vectorizable problems, in order to have the best performance from the SIMD architecture, the last generations of video cards provide a computing power exceeding the Teraflops. In particular NVIDIA [5] turns proposes a widely comprehensive and consistent approach to general purpose scientific computation. The NVIDIA Tesla C1060 Video Card, employed in the work described in this paper, houses one GPU GT200 with 240 computing's cores, 4 GB DDR3 memory, with 800 MHz speed and with a bandwidth of $\sim 102GB/s$.

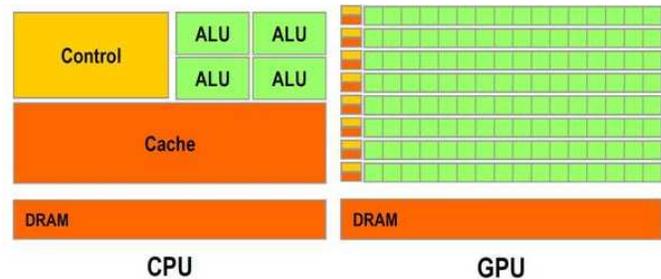


Fig. 3. CPU and GPU: different distributions of resources.

V. GPU USED FOR FAST PATTERN RECOGNITION IN THE NA62 RICH DETECTOR

In many cases the definition of trigger's primitives can be reduced to pattern recognition issues. This is the case for charged particle's track identification in magnetic spectrometers, trajectory in silicon detectors or photons rings in Cherenkov detectors. The RICH counter in the NA62 experiment falls in this last category. The main purpose of this detector is to distinguish between π and μ to contribute to the effective rejection of the $K^+ \rightarrow \mu^+ \nu$ decay, the main

TABLE I
RESULTS

Algorithm	Time (μs)
GHT	85
OPMH	139
ODMH	12.4

background for the signal $K^+ \rightarrow \pi^+ \nu \bar{\nu}$. The RICH will provide a $\pi-\mu$ separation at level of $5 \cdot 10^{-3}$ in the 15–35 GeV range, for decay products in the acceptance of the spectrometer (STRAWs). The RICH is a 17 m long 3 m in diameter tube filled by atmospheric pressure Neon. The Cherenkov light, reflected by a mirror of 17 m of focal length, will be collected in 2 spots of ~ 1000 PMTs each located in an hexagonal lattice with a side of 18mm. The fill factor is $\sim 90\%$, due to the use of reflective cones. Each PMT has a single photon resolution of $\sim 250ps$. The position of the rings center and radius are respectively related to the angle of the particle and its momentum. Having this information at trigger level is quite important for increasing the purity of many triggers of interest. The power and the speed of GPUs seems to be an interesting solution for this problem.

In order to test the feasibility and the performances, we have implemented three algorithms for ring finding in a sparse matrix of 1000 PMT, with ~ 20 fired PMT (“hits”) and noise; for better reproducing the conditions of interests we used the Geant4 based official Montecarlo for the NA62 experiment. The first algorithm we tested is based on the Generalized Hough Transform (GHT) where each hit is the center of a test circle; if several test circles, for a given radius, match the same point then this point is the center of the Cherenkov ring.

In the other two algorithms each PMT in the matrix is considered as candidate center. For each center the distance with all the hits is computed and the Cherenkov ring center is defined as the center with more hits at the same distance. The difference between the two algorithms, called respectively Optimized-for-Problem Multi-histograms approach (OPMH) and Optimized-for-Device Multi-histograms approach (ODMH), is in the management of the parallelization in the GPU structure. In the former case each cores have to make very simple operations (distance calculation), but the whole processor is employed for the same event, in the second case the single core has to perform more heavy duty in order to optimize the number of concurrent processes with the read and write procedure in the fast memory and to allow processing multiple events simultaneously in the same chip. In Tab.I we show the encouraging results obtained in terms of time processing for single event. The ODMH shows the best performances because is the only one where there is no conflict of writing and reading memory and allows simultaneous multi-events processing.

Thanks to the high bandwidth provided by the PCI-E bus (4 GB/s), the data transfer from an acquisition card (such a Gigabit card) to the GPU through the internal PC BUS, should not be an issue. In addition the processes of data transfer and computation in the GPU are managed concurrently: there is not time loss neither in data preparation for the processing nor

for transferring back the results from the video card to the PC.

VI. INTEGRATION IN THE NA62 TRIGGER SYSTEM

The easier way to exploit the computing capabilities offered by the GPUs in a multilevel trigger system, is to include them in the software level, as support for the calculations in the PC farm usually based on standard processors. In the NA62 trigger architecture the GPUs could be easily integrated in the L1 and L2 software levels, devoted, respectively, to the processing of data coming from the single detector and from the entire experiment. The rate in the software levels is lower than 1 MHz (maximum rate in input at L1) and the latency is not an issue: the main advantage of the GPU, in this case, is to reduce the cost and the dimension of the on-line farm.

Even more interesting is the possibility of using a system based on GPU directly to L0. In this case the advantages would be manifold: trigger primitives could be computed with high resolution (comparable to off-line), triggers with higher purity and efficiency could be defined at the lowest levels allowing the occupation of the readout bandwidth with additional physics or control triggers of interest. Use the GPU in the on-line trigger L0 is very challenging because of the needed for a answers with small and defined. The time for ring definition, as in the example above, should be 100 ns, being 10 MHz the rate at L0. The best preliminary result shown here ($12\mu s$) is still too far from this goal, but the following consideration should be taken into account:

- There is still room for improvement in the presented algorithms and new algorithms are already under test
- The system employed is based on a single GPU. The total load in the real environment could be subdivided between several GPUs.
- The next generation of graphics processors (already in the market) provides better performances at least a factor of two.

Given these considerations, the possibility to use the GPUs directly in the hardware level of the trigger is still very attractive.

VII. CONCLUSIONS

The use of commercial video cards in high energy physics trigger system is a very interesting possibility for several reasons. The power of the GPUs, for a defined class of algorithms, exceeds the CPUs of order of magnitude, allowing

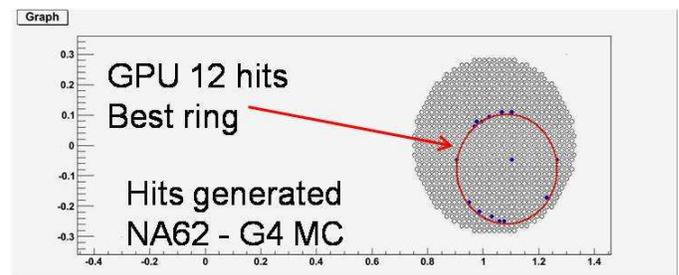


Fig. 4. Example of a ring reconstructed using the GPU (12 hits in a matrix of 1000 PMT).

to have very versatile and compact computing units to address problems of on-line event selection. Thanks to the use of commercial components, designed for sectors with very large market, this solution appears to be very cheap with respect to other solutions based on specialized hardware. In addition a system based on GPUs benefits directly from the continuous technological progress required from the video games and image processing industry. The use of GPUs both in hardware and software level allows to define a new architecture for the trigger system, where the hardware customized part is reduced to digitization and buffering while the whole logic, based on digital information, is performed in software. Such scheme can be easily adapted to any high energy physics experiment to increase the on-line selection power and to decrease the total cost.

REFERENCES

- [1] S. Adler *et al.* [The E949 Collaboration and E787 Collaboration], Phys. Rev. D **77** (2008) 052003 [arXiv:0709.1000 [hep-ex]].
- [2] G. Haefeli, A. Bay, A. Gong, H. Gong, M. Muecke, N. Neufeld and O. Schneider, Nucl. Instrum. Meth. A **560** (2006) 494.
- [3] B. G. Taylor, *Prepared for 8th Workshop on Electronics for LHC Experiments, Colmar, France, 9-13 Sep 2002*
- [4] G. Collazuol, In the same Conference IEEE NSS-MIC 2009
- [5] <http://www.nvidia.com/cuda/>
- [6] <http://www.gpgpu.org/>