Level Zero Trigger processor for the ultra rare kaon decay experiment—NA62

This content has been downloaded from IOPscience. Please scroll down to see the full text.

2016 JINST 11 C02037

(<http://iopscience.iop.org/1748-0221/11/02/C02037>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 188.184.3.52

This content was downloaded on 16/03/2016 at 10:31

Please note that [terms and conditions apply](#).

RECEIVED: October 29, 2015

REVISED: December 17, 2015

ACCEPTED: January 4, 2016

PUBLISHED: February 12, 2016

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS 2015,
SEPTEMBER 28TH – OCTOBER 2ND, 2015
LISBON, PORTUGAL

Level Zero Trigger processor for the ultra rare kaon decay experiment — NA62

S. Chiozzi,^{a,b} E. Gamberini,^{a,b} A. Gianoli,^{a,b} G. Mila,^{c,d} I. Neri,^{a,b} F. Petrucci^{a,b} and D. Soldi^{c,d,e}

^aINFN Ferrara,

Via Saragat, 1 44122 Ferrara, Italy

^bUniversity of Ferrara,

Via Savonarola, 9, 44121 Ferrara, Italy

^cINFN Torino,

Via Pietro Giuria 1, 10137 Torino, Italy

^dUniversity of Turin,

Via Pietro Giuria 1, 10137 Torino, Italy

^eCERN,

Geneva, Switzerland

E-mail: dario.soldi@cern.ch

ABSTRACT: In the NA62 experiment at CERN-SPS the communication between detectors and the Lowest Level (L0) trigger processor is performed via Ethernet packets, using the UDP protocol. The L0 Trigger Processor handles the signals from sub-detectors that take part to the trigger generation. In order to choose the best solution for its realization, two different approaches have been implemented. The first approach is fully based on a FPGA device while the second one joins an off-the-shelf PC to the FPGA. The performance of the two systems will be discussed and compared.

KEYWORDS: Trigger concepts and systems (hardware and software); Data acquisition circuits; Trigger algorithms; Digital electronic circuits

ARXIV EPRINT: [1505.03058](https://arxiv.org/abs/1505.03058)

Contents

1	Introduction	1
2	Trigger and data acquisition system	1
3	Level Zero Trigger	3
3.1	L0 board	3
3.2	L0 trigger processor on FPGA	4
3.3	L0 trigger processor on PC	6
4	Conclusions	7

1 Introduction

The NA62 experiment is designed to measure the (ultra) rare decay [1] $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ branching ratio with a precision of $\sim 10\%$ at the CERN Super Proton Synchrotron (SPS). The standard model expected value for the branching ratio is

$$BR(K^+ \rightarrow \pi^+ \nu \bar{\nu}) = (7.81_{-0.71}^{+0.8} \pm 0.29) \times 10^{-11} \quad (1.1)$$

Kaons are produced by a 400 GeV/c high-intensity SPS proton beam impinging on a beryllium target. Secondary particles of 75 GeV/c momentum are selected in a non separated beam composed of 6% of kaons, with a total rate of ~ 750 MHz. Only about 10% of the kaons produced decay in flight along the 65 m of fiducial volume (see figure 1), allowing the detector to collect $\sim 4.5 \cdot 10^{12}$ decays per year.

2 Trigger and data acquisition system

NA62 requires high-performance trigger and data acquisition due to the intense flux of particles. For this reason a multi-level trigger system is developed to scale the 10 MHz input rate to the 20 kHz of data finally recorded.

The first level (Level 0, L0) is based on the input from a few sub-detectors; after a positive L0 is issued, data is readout from front-end electronics buffers to the PC-farm.

Two different approaches for the realization of the L0TP were developed [2]. The first solution is fully based on the FPGA, and the whole logic for trigger selection is hardware programmed, while the second adds a commodity PC to compute the selection of the data (see figure 3). The same FPGA board is used in both projects, but in the second solution it only receives the primitives and dispatches triggers to the detectors, while the trigger algorithm is totally software implemented.

The higher trigger levels (L1 and L2) are implemented via software on the PC farm for further reconstruction and event building. A schematic drawing of NA62 trigger system is shown in figure 2. The L0 trigger algorithm handles the information from a selection of sub-detectors (Charged

Figure 1. Schematic view of the NA62 experiment at CERN SPS.

Figure 2. Schematic drawing of the NA62 trigger and DAQ system.

Hodoscope (CHOD), muon veto (MUV3), RICH, liquid krypton electromagnetic calorimeter (LKr), hadron calorimeters (MUV1, MUV2) and large-angle vetoes (LAV)), with a maximum output trigger rate of 1 MHz and a maximum latency of 1 ms. L1 algorithms run on data of the individual detectors, with a maximum output rate of 100 kHz and a non-fixed total latency of about 1 s, while the L2 trigger has an output rate of the order of 20 kHz with a maximum total latency equal to the period of the SPS beam-delivery cycle (about 10 seconds).

Figure 3. The L0TP designs are shown: on the left side, the system is fully implemented on FPGA, on the right side, the system is using an off the shelf PC.

3 Level Zero Trigger

When pre-selected characteristics of an event are satisfied, the single sub-detector locally generates 64 bit words called primitives containing the time information (40 bits) plus all the event features (PrimitiveIDs, 16 bits): energy, multiplicity, position of hits ecc. The primitives, collected in a Multi Trigger Packet (MTP) are sent via Ethernet links, using the UDP format, to the L0TP which selects good events matching the primitives with a preset mask. The remaining 8 bits are left reserved for a possible future implementation.

The UDP standard protocol has been chosen considering the 270 m total length of the experiment. A switch could easily be plugged to extend in space any possible other connection with detectors more distant than 100 m from the L0TP board without any modification of the protocol. The L0TP maximum latency permits to absorb the delay introduced by the switch.

MTPs are sent periodically, even in case of zero primitives, the default period being $6.4 \mu\text{s}$. All primitives generated in $6.4 \mu\text{s}$ should be merged in the same packet. For this reason, the number of primitives in each MTP is variable and it depends on the rate. The theoretical input rate limit is related to the gigabit Ethernet link capacity - 14 MHz.

The time is fundamental to align primitives coming asynchronously. It is stored in a 32-bit unsigned integer (Timestamp), relative to which all individual channel times are to be interpreted and made available to each sub-system through the Timing Trigger and Control (TTC) [6] system by a dedicated receiver. The Timestamp is uniquely related to the event number within a burst, this latter being the particle spill from the SPS. The LSB equals the period of the master clock, 24,95 ns. A Fine Time 8 bit information is also stored in the packet. The lowest significant bit is $1/256$ of the main clock period (97.47 ps), allowing a more precise data alignment and a tighter trigger coincidence window. Nonetheless detectors' precision can be different from one another, and therefore Fine Time information should be used accordingly and in combination with meaningful matching time windows.

3.1 L0 board

The L0TP board is a commercial development board, Terasic DE4 (figure 4), mounting an Altera Stratix IV FPGA. The DE4 board features 4 gigabit Ethernet links as default plus 4 available on the High Speed Mezzanine Cards (HSMC). 7 links are used as detector inputs while the last is reserved

Figure 4. The Terasic DE4 board.

Figure 5. Logical scheme of L0 Trigger Processor.

to send trigger information to the PC farm. An auxiliary board mounting the TTC receiver module is used to convert optical signals coming from SPS into digital levels to transmit the clock, start of the spill and end of the spill signals.

3.2 L0 trigger processor on FPGA

The logical scheme of the L0TP is shown in figure 5. The L0TP receives UDP ethernet MTP filled by primitives, unpacks them and sends PrimitiveIDs to a Look Up table (LUT) which selects the good events according to pre-selected masks. The trigger time precision is 25 ns. No Fine Time is used at trigger level.

Two clock domains are implemented on the FPGA. The first one uses a 125 MHz clock generated by a PLL fed with an internal 50 MHz oscillator. It drives the Ethernet logic and all the operations to align data and generate triggers. The other one is the 40 MHz clock coming from the main clock of the experiment. It is used to send triggers to the detectors in phase with their own clock.

For each Ethernet link, primitives are stored in a RAM for the alignment using the information of Timestamp and Fine Time to generate the memory address. It allows to have the primitives related to the same time to be written in the same address. Thanks to this realignment it is possible to read the addresses of each RAM in parallel and pass the information of the event to the LUT.

The memory is circular, with a programmable granularity, depending on the number of used bits of the Fine Time. The best granularity is 3.125 ns, obtained using 11 LSBs of Timestamp plus 3 MSB of Fine Time. In this case the depth of the RAM is 50 us. The RAM contains the primitive IDs and the whole Fine Time of the primitive. Being a circular buffer never cleaned during the burst, also the MSBs of the Timestamp not used in the RAM address are stored in the data block to avoid re-process the same primitives of a previous cycle. In this way the entire information of the time is accessible, and it ensures to process only data belonging to the same time.

All 7 Ethernet links work in parallel and after 3 packets received they start to process the first one passing PrimitiveIDs to the LUT. After 6.4 μ s they will receive a new packet and pass the second one to the LUT, and so on. It allows having always 3 packets stored in the RAM of alignment, avoiding the edge effects that could affect the information stored in the Ethernet. In fact, the processing time of an event by a detector is related to the complexity of the event itself. Waiting 3 packets guarantees a maximum delay of 19.2 us between the sources.

If a trigger mask is satisfied, a programmable downscaling could be applied. It permits to regulate the rate of the control triggers.

It is also forbidden to have two consecutive triggers closer than 100 ns, due to the TTC receiver logic. When the trigger has passed all these selections, the characteristics of the event are stored in a FIFO to be sent to the farm, while a 8-bit trigger word is ready to be synchronously dispatched to multiple (NA62 version) LTU modules [3]. Each such module will drive one TTC transmitter module passing the information received from the LOTP.

The latency is obtained with a dual port circular buffer 1.6 ms long. It is fed by the output of the LUT 125-MHz-clock driven. The reading process starts from the first address after the latency, calculated from the start of the spill. Then each position is read clock by clock at 40 MHz allowing the synchronization of the output triggers with the global clock of the experiment. Triggers are sorted in this RAM using the 14 LSBs of the Timestamp. This mechanism permits to read at the time $X + \text{latency}$ the trigger word generated at the Timestamp X .

As shown in the figure 5, a dedicated logic for special triggers has been developed. The L0 Trigger Processor will continuously monitor choke/error lines coming from each sub-detector. In response to an unmasked choke (error) signal, the L0 Trigger Processor will send a special choke-on (error-on) trigger without latency, and cease to dispatch triggers. This special trigger is expected to be serviced as any other by the sub-detectors: the corresponding data frames will provide a means to check that no data were lost prior to the signal. When all choke lines go low, the L0 Trigger Processor will issue a special choke-off (error-off) trigger to resume normal operations.

The pulser input receives signals in NIM logic, and it is connected to the LKr calibration input. When a calibration signal is generated in the LKr, the input becomes high and a special trigger is stored in the final circular buffer and delivered after the programmed latency.

Figure 6. Test setup for the PC-based L0TP during the 2015 data taking.

3.3 L0 trigger processor on PC

A development of the trigger logic using a common PC has been investigated to upgrade the L0TP[2]. The DE4 board also features a 8x PCI Express Gen. 2 link which in this case is used to deliver data directly to the RAM memory of the PC. The PC takes advantage of an Intel Core i7-4930K CPU @ 3.40GHz with Scientific Linux 6 kernel. This system can exploit the entire memory of the PC without any problem of synchronization of circular buffers as in the FPGA-based solution, since all primitive information is stored for the length of the spill. On the firmware side, this project features similar DE4 remote control, primitives reception, TTC signals management and trigger shipping to detector at fixed latency with the addition of PCIe access to the PC memory. The connection with LTU to dispatch triggers has been done using the same daughter card used by the FPGA-based system. On the software side the CPU directly reads primitives from memory and, using a single-threaded algorithm, aligns them looking for overlapping time windows in real-time. Since all primitive information is accessible by the matching algorithm it is possible to exploit all bits of Timestamp and Finetime. The best time granularity in this way is obtained using all bits of the finetime, 97 ps precision.

A great effort has been spent to understand if a PC-based L0TP could fit the timing constraints, which are not fixed as in a FPGA algorithm. It has been tested in Ferrara by using an Altera DE4 MTP generator. It continuously worked over 1500 bursts. At 13 MHz of input rate, a latency lower than 100 us was stably obtained. The communication between detectors, run control, and PC-based L0TP has been tested successfully at CERN. The system was tested parasitically receiving primitives from CHOD and MUV3 detectors without sending triggers to LTU, but dumping the event data into a file. This was possible using mirroring-port switches before the FPGA-based L0TP, with ability to copy the traffic from a single port to a mirror port 6. This connection didn't compromise the normal data-taking. The data consistency of the PC-based L0TP has been monitored. Further tests are planned at CERN after the 2015 data taking, to check the data quality and the efficiency of the system.

4 Conclusions

The NA62 experiment has been fully commissioned in 2014 during its pilot run [4], while from June 2015 the first physics data taking started with an intensity of the beam equal to 10% of the nominal. The FPGA-based L0TP was installed in September 2014 and it is working as expected with an efficiency of the electronic logic of 99.96 %, sending 100 KHz of triggers per spill with a latency of 100 μ s. The PC-based L0TP feasibility tests have been done and a comparison between the two systems is ongoing.

References

- [1] G. Anelli et al., *Proposal to measure the rare decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ at the CERN SPS*, CERN-SPSC-2005-013, CERN, Geneva Switzerland (2005).
- [2] E. Gamberini et al., *The Level 0 trigger processor for the NA62 experiment*, *Nucl. Instrum. Meth. A* (2015) in press.
- [3] M. Krivda, *Information on the ALICE trigger module*, <http://epweb2.ph.bham.ac.uk/user/krivda/alice/>.
- [4] NA62 collaboration, D. Soldi, *Results from the NA62 2014 commissioning run*, [arXiv:1505.03058](https://arxiv.org/abs/1505.03058).
- [5] M. Pivanti et al., *Implementation of a PC-based Level 0 trigger processor for the NA62 experiment*, *J. Phys. Conf. Ser.* **513** (2014) 012008.
- [6] B.G. Taylor, *TTC laser transmitter (TTCex, TTCtx, TTCmx)*, <http://ttc.web.cern.ch/TTC/TTCtxManual.pdf>.