



# The Trigger Supervisor of the NA48 experiment at CERN SPS

R. Arcidiacono\*, P.L. Barberis, F. Benotto, F. Bertolino, G. Govi<sup>1</sup>, E. Menichetti

*I.N.F.N. and Dipartimento di Fisica Sperimentale, Università di Torino, 10125 Torino, Italy*

Received 3 August 1999; accepted 20 September 1999

---

## Abstract

The NA48 experiment aims to measure direct CP violation in the  $K_L^0$  decays system with an accuracy of  $2 \times 10^{-4}$ . High performances are required to the trigger and acquisition systems. This paper describes the NA48 Trigger Supervisor, a 40 MHz pipelined hardware system which correlates and processes trigger informations from local trigger sources, searching for interesting patterns. The trigger packet include a *timestamp* information used by the readout systems to retrieve detector data. The design architecture and functionality during 98 data taking are described. © 2000 Elsevier Science B.V. All rights reserved.

*Keywords:* Trigger; Pipeline; High-energy physics; Data acquisition; NA48

---

## 1. Introduction

The NA48 experiment at the CERN SPS is a high precision experiment, searching for direct CP violation in the neutral kaon system through the measurement of

$$\begin{aligned} \text{Re}(\varepsilon'/\varepsilon) &= 1 - \frac{1}{6}R \\ &= 1 - \frac{1}{6} \frac{K_L \rightarrow \pi^0 \pi^0}{K_S \rightarrow \pi^0 \pi^0} \frac{K_S \rightarrow \pi^+ \pi^-}{K_L \rightarrow \pi^+ \pi^-} \end{aligned}$$

with an accuracy of  $2 \times 10^{-4}$  [1,2].

To achieve the required statistical accuracy and comparable systematic uncertainty, the experiment

is collecting a few millions  $K_L^0 \rightarrow \pi^0 \pi^0$  during three years of data taking at high beam intensity. The detector is exposed simultaneously to the  $K_L$  and  $K_S$  beams, measuring their charged and neutral two-pion decays. In order to handle large amount of data at high rate, very sophisticated trigger and data acquisition systems have been built.

The need of making the trigger and read-out systems essentially deadtimeless is a consequence of two independent requirements of the NA48 experimental technique. On the one hand, very high statistics is required to reach the goal sensitivity of  $2 \times 10^{-4}$ . Any sizeable loss of events due to dead time must be avoided. On the other, to prove with the required accuracy that losses and inefficiencies do not bias the result, they must be kept very small, as it is only possible when dead time is zero or negligible.

A key component of the trigger system architecture is the Trigger Supervisor (TS). The TS is a hardware system which correlates trigger information

---

\* Corresponding author. CERN, Div. EP, Bat. 32-1C-13, 1211 Geneva 23, Switzerland. Tel.: + 41-22-7678098.

*E-mail address:* roberta.arcidiacono@cern.ch  
(R. Arcidiacono).

<sup>1</sup> Now at CERN Geneve.

coming from different local trigger sources in order to produce a final decision to record event data. Once the decision is taken, the TS dispatches it to all front-end systems to initiate the readout sequence. The TS also performs trigger counting, downscaling and recording, dead-time control and trigger monitoring.

This paper describes the architecture of the NA48 Trigger Supervisor.

## 2. The NA48 trigger system

The NA48 detector main components are a charged particle spectrometer, used to reconstruct  $K^0$  decays to  $\pi^+\pi^-$ , and a highly segmented electromagnetic calorimeter, detecting  $2\pi^0$  decays. These same detectors also provides trigger information for the two modes.

The data acquisition system has been designed to have negligible dead time and to be highly efficient, to avoid possible biases between the charged and the neutral trigger branches. It is entirely organized as a data-driven system. Due to the very high rate to be handled, no handshake is provided between the different parts of the system, although several error detection mechanisms are implemented.

The SPS proton beam extraction for NA48 has a periodic structure of 14.4 s, with a flat top of  $\sim 2.5$  s during which protons are sent to the experiment. The interburst period of  $\sim 11.9$  s is used for data transfer to central data recording, setting-up, calibration and statistical calculations.

Data from all subdetectors are continuously digitized and stored into circular buffers 200  $\mu$ s deep, at a frequency of 40 MHz. The trigger system, composed by a Charged Chain, a Neutral Chain and the Trigger Supervisor, has to process, select and tag the events to be read out within the 200  $\mu$ s data persistence time.

The neutral trigger (NT) [3], used for the selection of  $K^0 \rightarrow 2\pi^0$  decay, is a 40 MHz deadtime free pipeline which provides trigger information based on a large number of digitised sum signals, obtained from 13 340 cells of the Liquid Krypton e.m. calorimeter. It reconstructs the number of showers seen in each projection of the calorimeter, the total energy, the energy centre of gravity and the kaon

lifetime. The full NT chain produces a trigger decision every 25 ns after a fixed delay of 3.2  $\mu$ s from the event time.

The charged trigger chain is a mixed hardware and software real-time system, partly pipelined, designed to detect  $K^0 \rightarrow \pi^+\pi^-$  decay using informations from the magnetic spectrometer. It significantly increases the signal to background ratio in the recorded data. The chain is composed by a fast first level trigger, followed by a computing engine (L2C) [4]. In order to reduce the input rate to L2C to less than 100 kHz from the total beam activity of  $\sim 1$  MHz, the first level logic performs loose consistency checks on signals from the neutral trigger, the hodoscopes, the veto systems and the first drift chamber. This is done inside the fully pipelined Level 1 Trigger Supervisor (L1TS). It continuously receives these signals, performs a time alignment and produces a 3-bit identification code to be transferred to the L2C upon finding the conditions satisfied. A 30-bit timestamp is also attached, carrying the event time information.

The L2C is an asynchronous queued system, partly pipelined, with massive parallel processing capability. For each L1 trigger, it computes hit coordinates, space points, tracks, comparing the 2-tracks invariant mass to a preselected window in order to tag good  $\pi^+\pi^-$  candidate. The computing time – although guaranteed not to exceed a fixed upper limit – varies on an event-by-event basis.

## 3. The Trigger Supervisor architecture and physical realization

The NA48 Trigger Supervisor is a fully pipelined 40 MHz digital system which correlates and processes the informations from the local trigger sources L2C, NT, L1TS.

It provides the final *trigger word*, including a *timestamp* indicating the event real time, whenever a required condition is fulfilled, and it broadcasts it to all Read Out Controllers (ROCs) within the data persistence ( $\sim 200$   $\mu$ s). The last stage is derandomized, in order to guarantee a minimum interval between two consecutive readout requests.

The TS has been implemented on 8 VME boards, housed in a 9U standard crate, connected by a




Fig. 1. Block diagram of the NA48 Trigger Supervisor.

private bus and monitored by a single board CPU.<sup>2</sup> The TS is hardware controlled by a finite state machine during the burst, and by the CPU in the interburst period. The CPU, running a real-time Unix-like operating system,<sup>3</sup> is the system bus master and is used to set up and interface the TS with the NA48 run control program.

Fig. 1 shows a block diagram of the TS, indicating the main components.

### 3.1. Input stage

The TS receives and correlates both synchronous and asynchronous trigger informations, the latter occasionally out of time order. The input stage is structured as four identical subdetector cards, each dedicated to a different trigger source (L1TS, L2C, NT and one used for miscellaneous signals). Each source provides the TS with up to 24 bits of data, synchronized with the system clock, together with a strobe used for data validation, at a maximum rate of 40 MHz. These signals have already been aligned in time among themselves at the source, therefore no additional time adjustment is required in the TS.

The trigger information is identified by a 30-bit timestamp indicating its 25 ns time slot within the burst. Asynchronous trigger sources, like L2C, provide their own timestamp together with the data, while for synchronous ones (L1TS, NT) the timestamp is derived by 40 MHz counters located on the TS.

Since the timestamp information is required to retrieve detector data from the circular buffers, it is very important that different trigger sources provide the same timestamp for a given event. The relative time alignment of the synchronous trigger sources is realized by using different presets for the 40 MHz counters.

Data from each trigger source are continuously stored into dual-ported, 8K deep, 56-bit wide fast static RAMs, addressed by the 13 low-order bits of the timestamp. Simultaneously with the writing, the memories are read out sequentially via the second port after a fixed (programmable) delay of  $\sim 100$   $\mu$ s. This delay is the maximum time budget given to the L2C for its computations.

The 8K memory space is scanned every 204.8  $\mu$ s, being rewritten many thousand of times in a burst. To recognize and discard data referring to “old” triggers, without having to clear the memories themselves, a special technique is adopted. The 17 high-order bits of the timestamp are stored into the memories together with each trigger word; during the reading process, these bits are matched with

<sup>2</sup> FIC 8234 from Creative Electronic Systems, Geneva, CH.

<sup>3</sup> OS-9 from Microware Systems Co., Des Moines, IA, USA.

a local timestamp counter and data are passed to the following stage only if there is agreement.

The extracted informations are also sent for monitoring purposes to external acquisition units, read out at every delivered trigger.

A general problem arising in a sampled system, superimposing a fixed time binning on a continuous event flow, is the possibility that an event occurring close to a time slice boundary is assigned to different time bins by independent trigger systems. This effect, always present in a sampled system, becomes more and more relevant as the sampling clock period becomes comparable to the relative time jitter of the independent trigger sources ( $\sim$  few ns for NA48).

For this reason, each local trigger source takes care of checking information in neighboring time slices before issuing its decision.

To overcome this potential source of inefficiency, the following scheme has been implemented. Whenever a coincidence between different conditions is required, the signal used as time reference is left unchanged while all the others are widened by two time slots, one preceding and the other following the actual position. In this way, signals displaced by one time slot in either direction are still effective in the trigger decision, allowing for greater efficiency.

### 3.2. Trigger word formation

For maximal flexibility, any subset of the 96 input bits can be used by the decision logic. This is achieved by a routing and combining network, implemented in FPGAs and RAMs.

The 96 bits coming from the 4 subdetector input cards are fed into a first layer of FPGAs<sup>4</sup> devoted to routing. A first reduction of the number of signals used in the decision is obtained. A second layer provides some basic logic functions, reducing the output bits to 48. A third layer, made by 4 RAMs, is used to encode the final 16-bit trigger word. In this last stage some timing signals allow for trigger enable inside or outside the SPS burst. All FPGAs

and RAMs are loadable by VMEbus, allowing for easy change of the trigger configuration by software.

### 3.3. Counting and downscaling

Each bit of the final trigger word can be individually downscaled by a factor up to 65 535, using programmable synchronous downscalers implemented in FPGAs.<sup>5</sup> The extensive use of downscaled trigger bits, with looser requirements, is mandatory in order to get accurate measurement of the main triggers' efficiency. Each bit is counted along the burst by a 24-bit counter, also implemented in a FPGA, whose output is available for reading during the interburst period.

The 16 trigger bits are used to generate a validation bit, called strobe. This process is implemented using a  $16\text{K} \times 1$  RAM driven by the trigger word itself, where every possible combination can be selected to generate the strobe. This scheme allows to temporarily disable any individual trigger type already produced by the decision logic.

### 3.4. Output stage

Whenever the strobe signal is present, the trigger word and a locally generated 30-bit timestamp provided by a master clock counter, are written into the output Trigger Queue Buffer (TQB). The purpose of the TQB is to buffer fast sequences of triggers and to allow for a fixed minimum time interval  $\Delta t$  between the broadcast of two consecutive triggers. The TQB has a FIFO-like structure, whose depth  $N$  is tuned according to the subdetectors memory depth (204.8  $\mu\text{s}$ ), the ROCs readout time  $\Delta t$  and the maximum L2C computation time ( $T_{L2C}$ ). These quantities are constrained by the inequality

$$5 \mu\text{s} + T_{L2C} + N \cdot \Delta t < 204.8 \mu\text{s} - \Delta t$$

where 5  $\mu\text{s}$  is an overestimate of the time spent through the L1 stage. Both the FIFO depth  $N$  and the  $\Delta t$  are programmable.

<sup>4</sup>XC3195AP175/223, made by XILINX Corp., San Jose, CA, USA.

<sup>5</sup>pASIC1, made by QUICKLOGIC Corp., Sunnyvale, CA, USA.

While the TQB has the functionality of a FIFO, due to the requirement of a programmable depth, it cannot be realized by standard FIFO chips. Instead, a fast dual-ported SRAM, storing trigger words and timestamps, is continuously addressed by two external counters providing the read and write pointers. A programmable down counter generate the  $\Delta t$  interval. An additional RAM, working as a finite state machine, controls the read and write operations to the SRAM, guaranteeing the selected FIFO depth.

A trigger word is extracted from the queue whenever the minimum time interval from the previous dispatching has elapsed, if a XOFF condition (see below) is not active. A sequential, burst-based event number (16 bit) is attached to the data, and the information is passed to the transmission stage.

### 3.5. Transmission stage

The resulting 64-bit trigger packet consists of the *event number* (16 bit), the *trigger word* (16 bit) and the *timestamp* (30 bit plus 2 spare). The sequential numbering allows both for transmission checks by the ROCs, comparing it with a locally generated event number, and for event building by the acquisition system. The trigger packet is sent to 10 destinations simultaneously, over fast dedicated serial links. Serialization, transmission and deserialization are performed by commercial VLSI interfaces.<sup>6</sup>

For the special case of a ROC sitting 200 m away from the TS electronics, an optical fiber channel has been developed, adapting the transmission and reception serial interfaces to the optical path. Transmission overhead is around 900 ns for 64-bit packet.

### 3.6. Dead-time control

Two basic dead-time sources exist in the system, both centrally controlled and monitored by the TS.

#### 3.6.1. XOFF

As already mentioned, the most important feature of the NA48 trigger system is the almost total dead timeless architecture. Nevertheless, since the intrinsic characteristics of some subsystems can still introduce dead time, special care has been made to have it monitored. This is necessary in order to carefully control possible, dead time introduced biases among the four decay modes. A simple XON/XOFF protocol is used to pause trigger dispatching by the TS whenever a ROC is unable to cope with the r/o request rate. This scheme is not intended for rate control, but rather for anomalous condition handling.

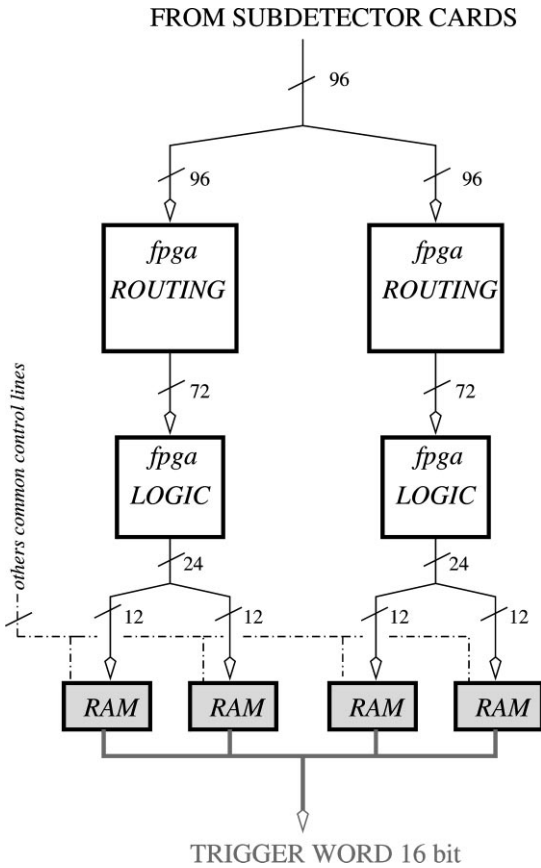
Each ROC system asserts its XOFF line whenever the amount of data in its output buffers exceeds a predefined upper limit. This limit is such that the ROC will still be able to successfully accept triggers in the following 4  $\mu$ s, the overestimated latency time for the propagation of XOFF condition. Due to the finite data persistence in the front-end memories (which are continuously written), when the TS receives an XOFF signal it cannot be simply frozen in the current status. Indeed, this could make the pending triggers too old at the time they are actually sent. The TS therefore responds to XOFF by simply blocking the transmission stage, while trigger data are still being received, processed, stored and extracted from the TQB. When all XOFF signals become unasserted, i.e. the amount of data in each ROCs is under some “safe” level, trigger transmission is resumed with the triggers which happen to be stored inside the TQB at that time. The XOFF is the only handshake protocol between the TS and the ROCs.

The XOFF dead time is monitored by counting both the number of 25 ns time slices during which the condition has lasted, and also by counting the number of valid triggers not dispatched.

#### 3.6.2. TQB full

The second dead-time source in the system is due to the limited FIFO depth in the TQB. Trigger rate fluctuations may result in the TQB being filled up, so that a new valid trigger cannot be stored. It must be recalled that the FIFO depth cannot be increased at will, just because of the limited data persistence in the detector circular buffers. As for

<sup>6</sup> AMD TAXIchip Tx/Rx set, handling a rate of 80 Mbit/s on 8-bit frames, using 4b/5b encoding.



#### 4. Performances

The NA48 Trigger Supervisor has been successfully used by the experiment during the 97 and 98 runs (Fig. 2).

The system has been very stable and reliable. In standard running conditions it provides in average 17K triggers/SPS burst at 6–7 kHz. In order to minimize the final error on the double ratio, also keeping some fraction of the total rate for rare decays studies, a 65–35% sharing between charged and neutral triggers is required. A further 65–35% sharing between main physics and efficiency measurement triggers is required for charged, while 80–20% is set for neutral. For this purpose, we fully exploit the powerful capability of the system of providing a large number of different trigger bits in parallel, individually down-scalable.

The measured L2C dead time is  $< 2\%$ , with a total charged efficiency  $> 97.5\%$ . The neutral efficiency is  $> 99.7\%$ , with no dead time. XOFF dead time measured by the Trigger Supervisor is  $< 1\%$  for 95% of the bursts.

In agreement with expectations from queuing theory [5], which foresees, for working values of  $\Delta t = 20 \mu s$ , TQB depth of 3, and a Poissonian input rate of 7 kHz, a dead-time fraction of  $2.4 \times 10^{-5}$ , the measured TQB full dead time is  $< 10^{-4}$ . Fig. 3 shows the measured FIFO occupancy (defined as the probability of occupying the FIFO in position  $n$ ), together with the prediction for 5,7,10 kHz input rates. The good agreement

Fig. 2. Look Up Table logic as implemented in the NA48 Trigger Supervisor.

the XOFF, this FIFO full condition is monitored by counting both the number of time slices during which it lasts and the number of valid triggers lost.

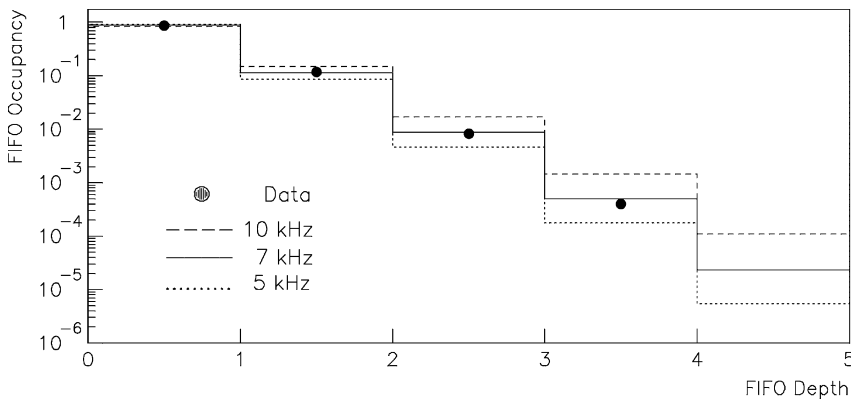


Fig. 3. Comparison between theoretical occupancy probability for different input rates and the measured values.

between the experimental points and the 7 kHz theoretical points demonstrates the correct functioning of the derandomizer.

## 5. Summary

The Trigger Supervisor of the NA48 experiment at the CERN SPS, designed to have negligible dead time, high-rate capability and easy configuration, has been successfully built and integrated into the innovative data acquisition system of the experiment. During the first two years of operation at high rate, excellent performance has been obtained, as needed to keep statistical and systematic uncertainties on the double ratio  $R$  below  $1 \times 10^{-3}$ .

## Acknowledgements

We gratefully acknowledge our founding Agency, Istituto Nazionale di Fisica Nucleare, for financial and technical support, and our colleagues in NA48 for many useful and clarifying discussions and suggestions.

## References

- [1] The NA48 collaboration, Proposal for a precision measurement of  $\varepsilon'/\varepsilon$  in CP violating  $K^0 \rightarrow 2\pi$  decays, CERN/SPSC/90-22, 20 July 1990.
- [2] V. Fanti et al., Nucl. Instr. and Meth. A, in preparation.
- [3] B. Gorini et al., IEEE Trans. Nucl. Sci. NS-45 (1998) 1771.
- [4] S. Anvar et al., Nucl. Instr. and Meth. A 419 (1998) 686.
- [5] G.P. Heath, Nucl. Instr. and Meth. A 278 (1989) 431.