

DAPNIA/99-06

October 1999

ONLINE TRACK RECONSTRUCTION
AND LEVEL 2 TRIGGERING IN NA48

**S. Anvar, J. Cogan, P. Debu, A. Formica, H. Le Provost,
I. Mandjavidze, M. Mur, S. Schanne, B. Vallage**

DAPNIA

*Presented at the IEEE Nuclear Science Symposium and Medical Imaging
(NSS/MIC 98), Toronto (Canada), November 11-14, 1998*

Online Track Reconstruction and Level 2 Triggering in NA48

S. Anvar, J. Cogan, P. Debu, A. Formica,
H. Le Provost, I. Mandjavidze, M. Mur, S. Schanne and B. Vallage
DAPNIA, CEA Saclay, 91191 Gif-sur-Yvette, France

Abstract

The level 2 trigger on charged particles (L2C) of the NA48 experiment features powerful capabilities for implementing complex triggering algorithms in real-time software. These algorithms may involve track reconstruction, vertex, invariant mass or transverse momentum computation, etc. After a brief outline of the L2C architecture, this paper describes the salient features of the I/O and processing capabilities of the event processor farm, with examples of algorithms that have been effectively implemented within the strong timing constraints of the NA48 level 2 trigger architecture. It concludes with the upgrade possibilities of the system and the possible use of it in other environments with minor changes.

I. INTRODUCTION

NA48 is a particle physics experiment aiming at the measurement of direct CP violation in kaon to 2-pion decays [1]. A calorimeter for neutral decays and a spectrometer for charged ones constitute the two main detectors of the experiment. The spectrometer is made of four drift chambers (DCH) and a magnet. The data acquisition system of the spectrometer is triggered by the so-called "level 2 charged trigger" (L2C) [1]. The L2C is fed with drift chamber data (hits) selected by the level 1 trigger at a rate between 100 kHz and 200 kHz. As shown in figure 1, these hits are first associated into usable coordinates by specialized hardware based on FPGAs. Then, all coordinates of drift chambers 1, 2 and 4 are sent by fiber optics and through a custom event-building switch to a PowerPC-based farm of up to 16 event processors called "event workers" (EW). Each event is processed by one 200 MHz PowerPC running Lynx real-time OS. A custom FPGA-based PMC board, the PMC Event Worker Interface (PEWI), takes care of all the critical I/O of the EW. Through the PEWI, each EW receives the data corresponding to an event, sends out the "trigger word" summarizing the characteristics of the event to the experiment's trigger supervisor and sends its own ID to the switch in order to notify that it is again free to process another event.

II. MAIN CONSTRAINTS

Any algorithm can in principle be implemented in the EWs. However, the circular buffers that store the data coming from the front end electronics are limited in size, leading to a maximum persistence of the data of $\sim 200 \mu\text{s}$. Of this maximum latency, only $100 \mu\text{s}$ is dedicated to the trigger decision. As a consequence, any event computation that exceeds this time limit causes the loss of the event. The L2C

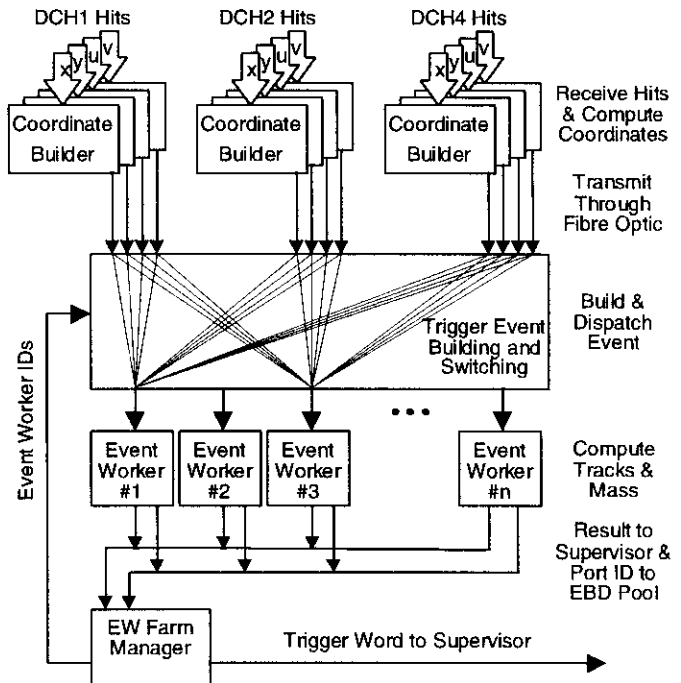


Figure 1: The L2 Charged Trigger Hardware Architecture.

being an asynchronous queued system, the trigger algorithm should, on the average, take much less than $100 \mu\text{s}$ in order to cope effectively with the fluctuations in the event rate [3]. The actual processing latency of each event is due to 1) its intrinsic complexity and 2) the time lost waiting in queues. The latency caused by event complexity is determined by the processing power of *one* EW, whereas the latency due to queuing is an increasing function of the ratio between the event *rate* and the processing power of the *whole* EW farm. The number of EWs in the farm is therefore determined by the tolerable loss of events at the maximum rate. Due to increasing needs for event statistics, the original event rate constraint on the L2C design [1] has been doubled to 200 kHz.

III. THE TRIGGER ALGORITHMS

In the 1998 NA48 run, the Event Workers were programmed to select $K \rightarrow \pi^+ \pi^-$ decays which are directly related to the main measurement of the experiment as well as 4-track events for specific rare decay studies such as $K \rightarrow \pi^+ \pi^- e^+ e^-$.

As shown in figure 2, each drift chamber includes 4 views, each view corresponding respectively to coordinates x , y , u , v , where $u = (x+y)/\sqrt{2}$ and $v = (y-x)/\sqrt{2}$. In all the algorithms implemented to this day, the 12 coordinate packets of an event




Figure 2: Tracks to Be Reconstructed Online.

(corresponding to 4 views in each drift chamber) are first converted into floating-point numbers. Then, (x,y) space-points are formed by using u and v to find out which y is associated with each x . Also, using again the redundancy introduced by coordinates u and v , possible missing x or y coordinates are reconstructed, thus making up for the small inefficiencies ($\sim 1\%$) in the drift chambers.

The algorithm stops if there are less than 2 space-points in either DCH1 or DCH2. If not, the next step – common to all algorithms – is to go through all possible pairs of track segments, a track segment being an association of one space-point in DCH1 and another in DCH2. If we have n (resp. p) space-points in DCH1 (resp. DCH2), the number of possible pairs of tracks is $n(n-1)p(p-1)/2$, which means that the combinatorial complexity of the algorithm increases like n^4 , where n is the number of particles per event. As a consequence, an upper limit of 8 is applied to the number of space points in each chamber. Any event exceeding this limit is flagged as “TOO COMPLEX”. Such events represent 0.025 % of the incoming data and the NA48 trigger policy is to read them out since they do not increase much the L2 trigger rate and may be of interest. A given pair of tracks is considered as a genuine decay if the closest distance of approach between them is smaller than 5 cm and situated in the decay region of the beam. If no vertex meeting these criteria is found, the pair of tracks is discarded and the algorithm proceeds with the next pair. If at least one pair of tracks does meet the criteria, the space-point computation algorithm is executed on DCH4 data.

1) The $K \rightarrow \pi^+ \pi^-$ Trigger Algorithm

For each pair of tracks that has survived the vertex cut corresponding to a $K \rightarrow \pi^+ \pi^-$ decay, a cut on the opening angle of the 2 tracks is performed to select decays with a high enough energy. Then, the pair of space-points corresponding to the linear extrapolation of the 2 tracks up to DCH4 is computed. Since the magnet deflects the trajectories of all

charged particles in the x direction, at least one space-points is expected to be found in a band of ± 10 cm around the y coordinate of each extrapolated track. If the deflected space-points are found, the principle of the algorithm is to assume that the pair of tracks is a $K \rightarrow \pi^+ \pi^-$ decay and compute the corresponding invariant mass and lifetime; if the mass and lifetime are compatible with a $K \rightarrow \pi^+ \pi^-$ decay in the region of interest,¹ then the event is flagged as “ $\pi^+ \pi^-$.” The resolution on mass computation achieved by this trigger algorithm is ~ 5 MeV.

2) The $K \rightarrow \pi^+ \pi^- e^+ e^-$ Trigger Algorithm

This algorithm is run whenever a 4-track event is encountered during the $K \rightarrow \pi^+ \pi^-$ analysis. For these events, the trigger looks for compatible vertices within ± 3 m of each other. Taking into account the topology of these events which often have 2 tracks close to each other (i.e. coming from a virtual photon), only 2 compatible vertices are required instead of the 6 that can logically be built having 4 tracks. To ascertain the reality of the 4 tracks, it is also required that DHC4 contains at least 3 space-points, allowing for one space-point miss (Cf. figure 3).

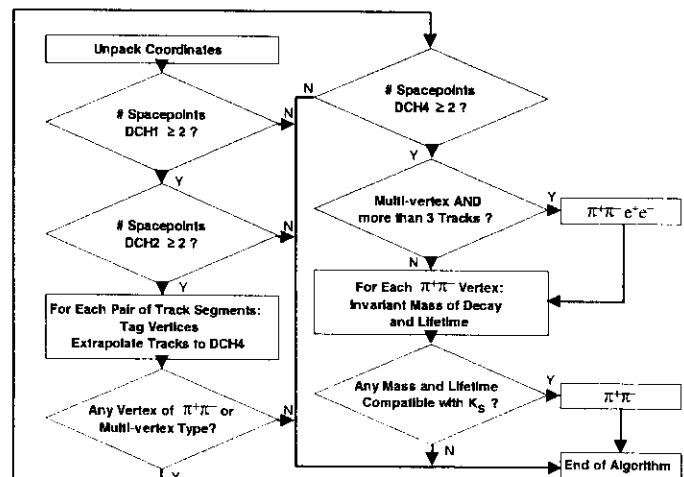


Figure 3: $K \rightarrow \pi^+ \pi^-$ and $K \rightarrow \pi^+ \pi^- e^+ e^-$ Trigger Algorithms.

¹ mass > 472 Mev and lifetime < $4.5 \times (K_S \text{ lifetime})$.

IV. THE EVENT WORKER FARM

A. Mono-processor Event Workers

The handling of a 200 kHz rate has become possible by today's off-the-shelf processors at an affordable cost. In fact it has even become possible to replace the former EWs based on 4-DSP clusters [1] [5] by mono-processor EWs based on PowerPC 604 VME single board computers from CES [4] which not only improve the performance (Cf. figure 4) of the system but also allow for a much easier maintenance of the trigger software since all the parallel code has been turned into a straightforward sequential program. A farm of 8 EWs, each based on 200 MHz PowerPC 604s, already copes with a 200 kHz rate with negligible event loss. The projected 1999 upgrade to 300 MHz processors will therefore not much improve the time performance of the nominal trigger but will allow for more complex algorithms especially those intended to improve triggering on rare decay events.




Figure 4: Average Performances of Successive Upgrades.

B. The EW I/O Interface

The data coming from the Event Builder & Dispatcher (EBD, see figure 1 and [2]) is fed to the EW memory through the so-called PMC Event Worker Interface (PEWI) (Cf. figure 5). It is a home-made FPGA-based card built after the PMC (PCI Mezzanine Card) standard that is plugged on the CES PowerPC board and is connected to the Event Dispatcher through the front panel. From one side, it receives the event data from the EBD and from the other it writes the data into the board memory through the PCI bus. It signals the processor through a mailbox when data is ready for processing. It has been measured that the fastest strategy to deliver data is PEWI as a PCI master and the host polling on a mailbox to receive the "data-ready" signal. The PEWI also takes care of the transmission of the trigger response to the

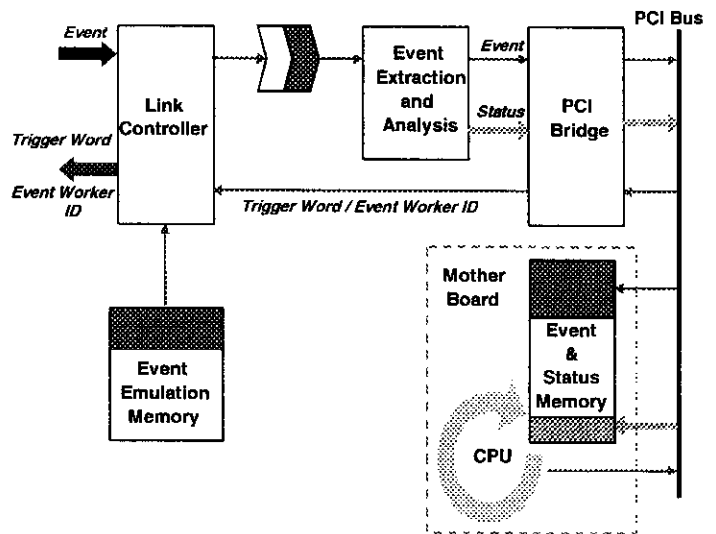


Figure 5: The PMC Event Worker Interface (PEWI)

EW Farm Manager and the EW Id to the EBD (see [2] for details). The PEWI card also runs some checks on the data such as checksum or time-stamp coherence (all 12 packets must have the same time stamp) and can generate errors though PCI interrupts.

A built-in debugging feature of the PEWI, the Event Emulation Memory, allows for standalone tests that do not need the presence of an EBD board. This feature proves also useful for the study of trigger algorithms by eliminating the need for a complete system.

Due to historical reasons [2], the protocol used between the PEWI and the EBD is the same as I/O links of the TMS320C40 DSP of Texas Instruments [5]. This protocol accounts for the relatively slow data rate between the EBD and the PEWI card (~ 8 Mbytes/s). In 300 MHz EWs, the bottleneck that hinders performance is the slow I/O capability. Since there are no more DSP-based Event Workers in the system, the PEWI and EBD FPGAs will be reprogrammed to implement a faster protocol (~ 20 Mbytes/s).

V. PERFORMANCES OF THE EW FARM

A. The Timing Distribution

The processing time distribution for nominal events inside a PowerPC-based EW is given in figure 6. The four peaks on the distribution correspond to respective abundance of events that survive the different cuts of the algorithm (Cf. figure 3). The first peak is populated by "empty events", that is events that contain no space-points in DCH1. The second peak corresponds to events that have enough space-points in DCH1 but none in DCH2. The third corresponds to events that have enough space-points (≥ 2) in DCH1 and DCH2 but no valid vertex that would indicate the presence of a genuine decay. The fourth peak is populated by events that have survived the vertex cut and for which kinematics computations have been made.

Event Processing Time Distribution
(PowerPC 604 @ 300 MHz)




Figure 6: EW Processing Time Distribution.

B. Performance Analysis

In spite of the complexity of the reconstruction algorithm, the average computing time on a 300 MHz PowerPC 604 is an impressive 16.3 μ s. This can be explained by two facts: 1) the event data being limited (\sim 100 bytes for a nominal event), all of the data remains in the processor's 32 kBytes level 1 cache and is thus processed at full speed, and 2) the 32 kBytes instruction cache can contain the whole algorithm code. The loading of the code into the instruction cache can take some precious time measured in milliseconds. For this reason, before entering a high-priority state the algorithm is run a first time on a pre-loaded complex enough event so that almost all of the different parts of the code are run and therefore loaded into the instruction cache.

VI. THE EW FARM AND TRIGGERING IN GENERAL

The use of a general-purpose processor such as the PowerPC in the EW farm has made the idea of a software-based, versatile and powerful level 2 trigger come true. The behavior of such a system is deterministic enough (up to the microsecond level) to allow for tough real-time constraints such as the 100 μ s latency in the L2C². At the same time, the trigger algorithms can be worked on by simple classical programming in a high-level language such as C or C++. This is important if physicists who are not computer experts are to work on the trigger without having to take care of communications or the hassles of parallel programming. The core of the trigger code thus becomes more readable and more comprehensible, allowing for a better understanding of its behavior and, therefore, a greater control over its effects on the physics off-line analysis. In the case of the L2C, the same source code is used both for the real-time trigger and for the off-line analysis of the trigger. In other words, simulating the

trigger is no more a challenge since you can just use the same code.

The versatility of the system is total as long as the trigger code fits into the instruction cache of the processor. If bigger codes are needed, then the constraints on latency must be relaxed, which is possible only if the size of the front-end ring buffers is increased. Extending the size of the event farm allows, in principle, to cope with higher event rates.

The EW farm design can be used for other triggers, provided that the data input protocol is adapted to the device that will feed it. This device will necessarily be an event builder on top of a switch: the data related to one event must be collected from different distant detectors and then sent to one of the EWs of the farm. The event builder's job is to build the event, to obtain a free EW Id from the farm manager and to send the event data to the EW through the switch. The switch can be either a home-made one as in the EBD, or a commercial one based on a standard protocol such as ATM or Ethernet.

VII. REFERENCES

- [1] "Proposal for a Precision Measurement of ϵ'/ϵ in CP Violating $K_0 \rightarrow 2\pi$ Decays," CERN/SPS/90-22 SPSC/P253, 1990.
- [2] S. Anvar, F. Bugeon, P. Debu, J.L. Fallou, H.Le Provost, F. Louis, M. Mur, S. Schanne, G. Tarte and B. Vallage, "The Charged Trigger System of NA48 at CERN," *IEEE Transactions on Nuclear Sciences*, vol. 45, pp. 1776-1781, August 1998.
- [3] James Martin, System Analysis for Data Transmission, Englewood Cliffs, New Jersey: Prentice-Hall Inc., 1972.
- [4] Creative Electronic Systems S.A., <http://www.ces.ch>.
- [5] MIZAR, 2410 Luna Road Carrollton, TX 75006, USA.

² Provided that a real-time OS is used.