

A concept for the NA62 trigger and data acquisition

M.S. Sozzi

November 1st, 2007

Abstract

In this document some requirements and a possible implementation for part of the trigger and data acquisition system (TDAQ) of the NA62 experiment are discussed, collecting input from discussions and work by many people in the collaboration, as an illustration of the present status and a basis for further discussion.

1 Introduction

The NA62 experiment at the CERN SPS [1] aims at a measurement of the ultra-rare decay mode $K^+ \rightarrow \pi^+ \nu \bar{\nu}$, to provide a stringent test of the Standard Model. The current schedule foresees the start of data-taking in 2011.

A high kaon flux is required for such measurement, and a high-performance trigger and data acquisition system is a mandatory requirement. The system must have reasonably low dead time and must be highly reliable in terms of data collection, since the vetoing capabilities of the sub-detectors are crucial to the experiment.

The existing NA48 trigger and data acquisition electronics, designed more than 10 years ago, is unsuited to the task and can no longer be maintained, therefore an entirely new system is required.

Considerations linked to reliability, simplicity and cost issues, led to the concept of a unified trigger and data acquisition system in which data is handled in the digital domain as soon as possible.

In the following some of the ideas and possible implementation solutions which emerged in discussions so far are described, as a first step towards the

collection of final requirements for the TDAQ system, the identification of viable (possibly already existing) solutions and the definition of a design.

2 Input figures

This section summarizes some information on the expected sub-detector rates and capabilities; this is not meant to be exhaustive nor definitive, but should be rather considered as a collection of figures useful to design the TDAQ system. Please note that several of the above figures are still guesstimates, and should be reviewed by sub-detector experts (this being actually one of the purposes of this note).

2.1 Rates

This section summarizes the relevant expected rates in the experiment according to the recent high-intensity K^+ beam layout [2].

A 75 GeV/c positive hadron beam produced from $3.3 \cdot 10^{12}$ protons per spill (4.8 s burst duration with a period of 16.8 s, the effective burst duration for computing instantaneous rates being rather around 3.0 s) will provide about $1.5 \cdot 10^8$ ($1.6 \cdot 10^9$) kaons (pions) per pulse at the exit of the final collimator. Such rates will be roughly 30 times higher than those due to the two simultaneous beams in the NA48/2 experiment (2003-04).

As for all SPS fixed-target experiments, an important difference with respect to LHC experiments is that the beam is not strictly time-bunched, and therefore there is no synchronous time signal correlated to the occurrence of interactions, which occur at a roughly constant rate during the spill, with no intrinsic time reference.

The rates of particles illuminating the different detectors were estimated using a fast MC particle-level simulation [3] (including beam transport, photon conversion, *bremsstrahlung*, etc.) of the 6 main kaon decay modes, occurring in a longitudinal region extending from the end of the final collimator (102.5 m) to just before the LKr calorimeter (240 m); the amount of simulated decays corresponds to about 200K events in the region upstream of the detectors, results being scaled to 150.5×10^6 K at the exit of the final collimator.

The simulated setup is a sketchy version of the one appearing in the P-326 proposal, with a double spectrometer, 12 large-angle veto rings and

MAMUD; a hodoscope plane in the MAMUD region is also included, whose acceptance matches that of the MAMUD except that it covers the coils of the former (leaving out only a central $30 \times 20 \text{ cm}^2$ hole). While an improved simulation should include the new design with a single spectrometer and a different muon veto detector, the present estimates are expected to be in first approximation indicative of those for the final setup, and the changes in the design are not too significant at the present rough level of study.

Figures 1 and 2 show the simulated kaon momentum and decay point distributions, respectively.




Figure 1: Momentum spectrum of simulated K decays.

It goes without saying that all the numbers in this study should be considered to be indicative only and definitely not precise estimates, which need to wait for a more accurate campaign of simulation.

Since only K decays in the detector region were simulated, in the case of muons two additional contributions (“halo”) were added “by hand” to




Figure 2: Horizontal projection of the K decay position.

the above estimated rates: muons originating from beam π decays occurring anywhere in the setup, and those due to K decays occurring upstream of the end of the final collimator [4] (muons from late pion decays can illuminate somewhat the downstream straw chambers); these figures were obtained by N. Doble with a different simulation and interpolated where required.

Table 2.1 lists the particle rates per spill and the corresponding in-spill average rates (over 4.8 s); design values including some arbitrarily chosen safety factors are also listed, both for average and instantaneous rates, the latter being expected to be larger by a factor $\sim 4.8/3 = 1.6$ due to the non-uniform time structure of the beam. The rates include the particle multiplicities per event.

Sub-detector	Hits/spill (10^6)							Rate (MHz)	Design (MHz)		Average
	K^+	π^\pm	μ^\pm	(+ halo)	γ	e^\pm	all	avg	avg	inst	mult.
CEDAR	150	–	–		–	–	150	31.3	60	100	1.00
Gigatracker	150	1600	20		–	1	2280	475	1000	1600	1.00
Anti 1	0	0	0	(+4.1)	0.04	0.003	4.1	1	3	4	1.00
Anti 4	0	0	0	(+4.4)	0.2	0.01	4.6	1	3	4	1.01
Anti 8	0	0	0.4	(+4.1)	0.3	0.02	4.8	1	3	4	1.01
Anti 12	0	0.2	1.2	(+3.0)	0.7	0.04	5.1	1	3	4	1.01
Anti total	0	0.3	4.6	(+36.9)	3.4	0.25	45.5	9.5	20	30	1.62
Anti OR	0	0.2	2.1	(+11.5)	3.4	0.25	17.5	3.6	10	15	1.11
Straw 1	0	6.1	11.7	(+16.3)	–	0.8	34.8	7.2	15	25	1.09
Straw 2	0	6.9	13.3	(+16.4)	–	0.9	37.4	7.8	15	25	1.09
Straw 3	0	7.4	14.1	(+16.9)	–	0.9	39.4	8.2	15	25	1.09
Straw 4	0	8.1	15.1	(+17.4)	–	0.9	41.5	8.6	15	25	1.10
Straw 5	0	8.6	15.9	(+17.6)	–	1.0	43.2	9.0	15	25	1.10
Straw 6	0	9.1	16.9	(+19.0)	–	1.0	46.0	9.6	15	25	1.10
Straw total	0	46.2	87.0	(+103.6)	–	5.5	242.3	50.5	100	150	5.48
Straw OR	0	9.7	17.4	(+19.0)	–	1.2	47.3	9.9	20	30	1.12
RICH	0	9.2	17.1	(+14.2)	–	1.1	41.5	8.6	15	25	1.10
IRC	0	0.9	1.9	(+8.7)	1.5	0.07	13.1	25	40	70	1.08
LKr	0	7.4	16.3	(+12.0)	13.9	0.8	50.4	10.5	20	30	1.51
N. hod	0	6.5 (*)	16.3	(+12.0)	14.0	0.8	49.5	10.3	20	30	1.30
MAMUD	0	–	14.9	(+29.2)(**)	0.1	0.0	44.2	9.2	25	40	1.00
Muon hod	0	–	17.8	(+29.8)(**)	0.1	0.0	47.7	9.9	25	40	1.00
SAC	0	1.7	2.7	(+1.3)	1.6	0.1	7.4	1.5	10	15	1.20

Table 1: Detector hit rates estimates; the figures include the average particle multiplicity and (for “Anti total” and “Straw total”) the average ring/chamber multiplicity, as listed in the last column. The muon ‘halo’ includes contributions due to pion decays and decays upstream of the final collimator. (*) Rough estimate assuming pion showering probability 0.7 and lateral sampling fraction 55%. (**) Rough estimate by area scaling from Straw 6.

2.2 Sub-detectors

In the expected high-rate environment the time resolution of sub-detectors (SDs) is the crucial parameter, the typical figure required being 100 ps. For some sub-detectors time (and channel, of course) is the only relevant information provided, and thus high-resolution TDCs are the only digitizing devices required, possibly providing two time measurements per signal to estimate the pulse height and perform a slewing correction. For other SDs accurate pulse height information is also required, and flash ADCs (FADCs) are the preferred choice, allowing to extract both pulse height and accurate timing information (with some computational effort). The above points to one main distinction between sub-detectors from the TDAQ point of view: TDC-equipped SDs are self-triggering and their raw (before any trigger) data rate is directly determined by the hit rate and channel occupancy per event; on the contrary, FADC-equipped SDs are continuously sampled and their raw data rate is determined by the sampling frequency for the total number of channels and the zero suppression scheme adopted.

Here some relevant design parameters of the NA62 sub-detectors are summarized, as far as they are known at this time. The overall total number of channels is estimated of the order of 170 thousand for the whole experiment, with a very low occupancy for interesting events.

- **CEDAR.** The CEDAR RICH detector, located about 200 m upstream of the main sub-detectors on the K/π beam, will only detect K^+ , firing a set of fast photo-multipliers (or similar), whose number is not expected to exceed 256. Fast and accurate time information is provided, and the high channel hit rate also requires a high double pulse resolution.

The use of the CEDAR information in the trigger would help in reducing the rate due to muons originating upstream the decay volume and beam pion decays, which amount to about 1/3 of the total track rate at the downstream detector location; as such figure is not very large the use of the CEDAR in the trigger is not mandatory but should be left as an option.

- **Gigatracker.** The silicon pixel tracker on the K/π beam will be located about 150 m upstream of the main sub-detectors and will detect all charged particles in the beam.

About 54 thousand detector channels are foreseen in total, using highly integrated electronics. Only time information will be recorded, possibly with two time values per hit to allow offline slewing correction.

An average occupancy of 2 pixels \times 3 stations per crossing particle can be

assumed. The gigatracker information will not be used at the lowest trigger level.

- **Large-angle vetos (ANTI 1-12).** 12 annular photon veto ring counters, spaced about 10 m apart from each other, will be read by about 2.5 thousand photo-multipliers in case of the lead glass solution (this number being smaller for other counter technology choices proposed). The particle rates increase moving from upstream to downstream rings (a factor 10 for electrons, 20 for photons and 100 for muons); figures 3, 4, 5 and 6 show the expected energy spectra for π^+ , μ^+ , e^+ and γ hitting any of the ANTIs and originating from K decays after the final collimator (muon halo not included).




Figure 3: Momentum spectrum of π^+ from K decays downstream of final collimator hitting the ANTI.

Being a vetoing sub-detector, time information and resolution is crucial for the ANTI; pulse-height information from this sub-detector is also desir-

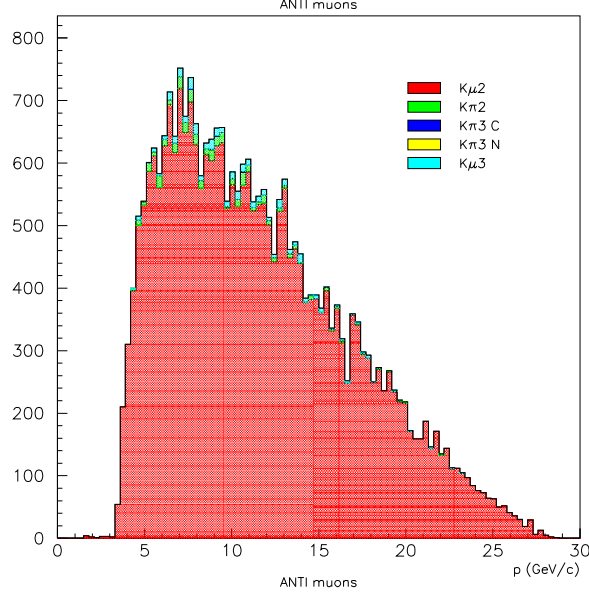


Figure 4: Momentum spectrum of μ^+ from K decays downstream of final collimator hitting the ANTIs (no muon halo).

able in order to distinguish showering particles from MIPs, but this is not strictly required to be available online.

The overall time resolution will be limited at least by the spatial extension of the individual counters, expected to be at least of the order of 40 cm, corresponding to a few ns at best if double-side readout is not implemented (of course such figure might be further degraded depending on the chosen read-out scheme); to keep random vetoing below $O(1\%)$ individual counter rates should not exceed $O(1 \text{ MHz})$.

By default ANTIs are not assumed to be used in the lowest level trigger, since their standalone reduction to the $(\pi^+\pi^0)$ background rate is estimated to be not too large at face value (about 30%, less than 20% on the total decay rate); however their effect on top of other trigger conditions can be larger and

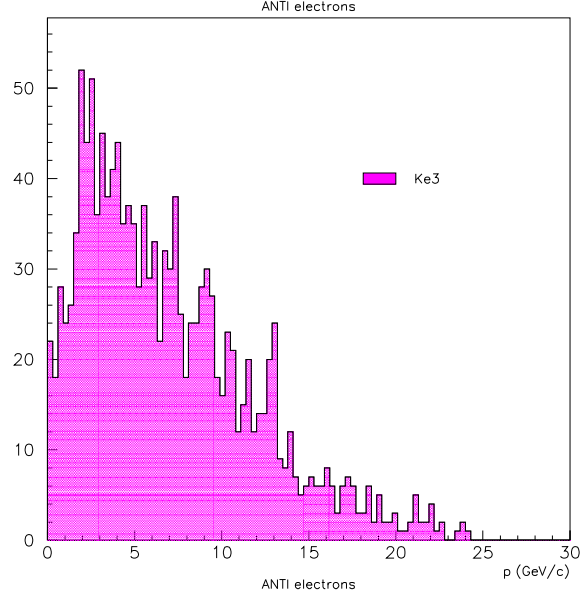


Figure 5: Momentum spectrum of e^+ from K decays downstream of final collimator hitting the ANTIs.

significant, and therefore inclusion of the ANTIs in the lowest level trigger should be kept as an option.

- **Straw tracker.** 6 stations of straw chambers, spaced about 10 m apart from each other, formed the double magnetic spectrometer in the original design¹, tracking charged particles outside the beam pipe region, for a total of about 12 thousand channels.

Only time information will be recorded. Due to the intrinsic drift time the signals are assumed to have an intrinsic time spread (channel by channel and event by event) of up to 50 ns.

¹The use of a single 3 or 4-chamber spectrometer is not expected to impact significantly on what is discussed here, except of course for the corresponding proportional reduction of the data rate from the straw tracker system itself.




Figure 6: Energy spectrum of γ from K decays downstream of final collimator hitting the ANTIs.

Due to the complexity of a track reconstruction algorithm the straw tracker is not expected to be used in the lowest level trigger.

- **RICH.** The RICH detector [5] will detect all fast charged particles outside the beam pipe and will be positively used in the lowest level trigger to identify the presence of a fast charged particle (this function was performed by a charged hodoscope in an earlier design).

Figures 7, 8, 9 show the expected energy spectra for π^+ , μ^+ and e^+ hitting the RICH and originating from K decays after the final collimator (muon halo not included).

Light will be collected by an array of about 2100 photo-multipliers; The average number of PMTs hit per track was estimated to be around 30, the design value being 40 (channel occupancy $\simeq 0.02$). Only time information




Figure 7: Energy spectrum of π^+ from K decays downstream of final collimator hitting the RICH.

will be recorded, although to perform the required slewing correction some information on the pulse-height is required: the present scheme foresees the recording of both leading and trailing edge threshold-crossing times for this purpose (two times per hit PMT).

- **LKr calorimeter.** The Liquid Krypton (LKr) ionization chamber calorimeter must be used as a highly-efficient veto at the lowest trigger level to suppress $\pi^+\pi^0$ decays. Besides time resolution, also the energy and space resolutions of the calorimeter cannot be degraded significantly with respect to the NA48 setup, as it will be used in the analysis for further background suppression. Figures 10, 11, 12 and 13 show the expected energy spectra for π^+ , μ^+ , e^+ and γ hitting the LKr calorimeter and originating from K decays after the final collimator (muon halo not included).

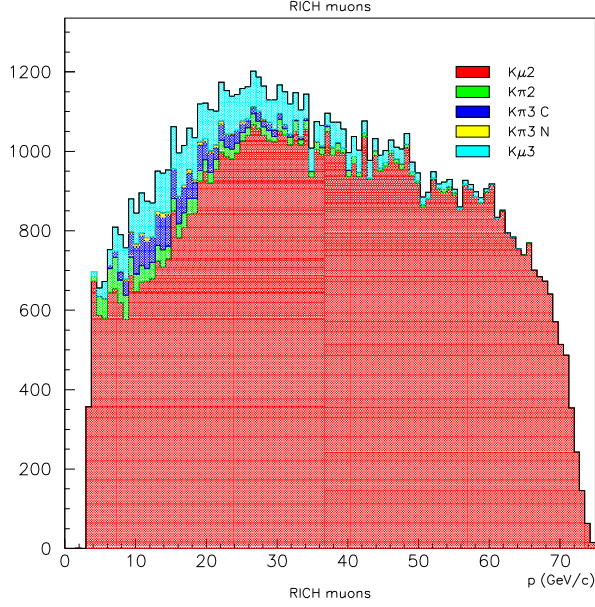


Figure 8: Energy spectrum of μ^+ from K decays downstream of final collimator hitting the RICH.

In NA48 the signals from the 13248 cells of the LKr calorimeter were shaped to 70 ns FWHM and individually flash digitized at 40 MHz; a variable gain shaper was used, resulting in an effective 14 bit resolution encoded into a 12 bit word including 2 gain-range bits. Efficient pulse reconstruction was obtained using 8 time samples (per cell per event), and the offline time resolution was of order 200 ps per cluster. The number of cells used in the offline analysis averaged at 100 per showering particle; however, the useful cells could not be identified *a priori* on the basis of their own pulse-height, but only in relation to the presence of a neighboring significant deposit of energy, since cells with very low energy deposits at the periphery of clusters are important to achieve the required resolution; this indicates that any online zero suppression scheme requires non-trivial electronics and communication




Figure 9: Energy spectrum of e^+ from K decays downstream of final collimator hitting the RICH.

among different cells and regions of the calorimeter.

Calorimetric information at the trigger level is useful to veto the abundant $\pi^+\pi^0$ background decays; while the number, energy and positions of clusters is required for an efficient vetoing, the precisely timed amount of energy deposited in large regions (*e.g.* quadrants) of the LKr can be used in a profitable way at the trigger level, so that the presence of the charged pion shower (in two thirds of the signal events) can be accounted for. The use of the calorimeter information in the lowest level trigger might be avoided in case the neutral hodoscope (see below) is used for such purpose.

- **Neutral hodoscope.** The scintillating fibre hodoscope embedded in the LKr calorimeter volume at the approximate depth of maximum EM shower development is sensitive to EM showering particles above some thresh-




Figure 10: Energy spectrum of π^+ from K decays downstream of final collimator hitting the LKr calorimeter.

old energy. It consists in 32 photo-multipliers connected to 120 cm long scintillating fibre bundles (whose intrinsic time jitter can thus be expected to be of order 10 ns). In NA48 this device (optimized for higher-energy photons) suffered from low efficiency for photons below about 20 GeV; other problems are the fact that a few channels of the detector are dead, and equalization of individual channel gains (and times) is not possible. The above issues can only be solved by opening the LKr calorimeter, a rather delicate and expensive operation (to say the least).

Only time information is required to be recorded.

This detector might offer – in principle – a possibility to suppress $\pi^+\pi^0$ decays at the lowest trigger level without requiring to handle the significantly larger amount of data from the LKr calorimeter at such an early stage.




Figure 11: Energy spectrum of μ^+ from K decays downstream of final collimator hitting the LKr calorimeter.

- **Muon veto.** In the original proposal, the muon veto system was based on planes of plastic scintillators with embedded wavelength shifting fibres, inserted in a magnetized volume and detecting all muons crossing the apparatus, except those staying in the beam region.

Figure 14 shows the expected energy spectra for μ^+ hitting the MAMUD and originating from K decays after the final collimator (muon halo not included).

The muon detector is the most crucial element which must be used at the lowest trigger level to reject the dominant $\mu^+\nu$ decays, and therefore a very good online time resolution is mandatory. Since this cannot be provided by large scintillator slabs, the use of a fast timing detector placed in the middle of the detector was foreseen; such detector should of course be placed deep




Figure 12: Energy spectrum of e^+ from K decays downstream of final collimator hitting the LKr calorimeter.

enough within the muon shielding to be reached only by muons and avoid significant random vetoing due to hadronic shower leakage.

The use of a device similar to the existing NA48 charged hodoscope, with appropriately improved acceptance and online time resolution, was initially suggested for this purpose. The existing NA48 charged hodoscope consists in two planes of 64 plastic scintillator bars each, of maximum length 120 cm, corresponding to an intrinsic minimum time jitter of about 6 ns full-width. A suitable logic with multiple delayed coincidences among individual counters could in principle reduce such jitter to the ns level.

Other more performant schemes were proposed, in which the time jitter can be further reduced by using scintillator counters of smaller size.

It should be noted, moreover, that in the standard MAMUD design




Figure 13: Energy spectrum of γ from K decays downstream of final collimator hitting the LKr calorimeter.

the presence of coils implies a rather large non-instrumented central region (roughly $70 \times 40 \text{ cm}^2$): as far as such acceptance hole does not match a similar one in the RICH detector, this significantly affects the trigger rate by limiting the muon veto capability of the muon detector. The fast timing plane used online should therefore match or exceed the RICH acceptance.

For the MAMUD detector itself about 2100 photo-multiplier channels (corresponding to 73 detector planes) were foreseen, and it is assumed that only time information needs to be recorded (possibly with a pulse-height correction). The fast timing plane itself might account for about 10% of the total number of channels.

- **IRC, SAC.** Small calorimeters will close the acceptance for photons at small angles; due to their small acceptance their use in the trigger is not




Figure 14: Momentum spectrum of μ^+ from K decays downstream of final collimator hitting the MAMUD.

very relevant. Both time and energy (pulse-height) information needs to be recorded.

Table 2.2 tries to summarize some figures related to the expected rates from each sub-detector.

Sub-detector	R-O type	Rate (MHz)	Total channels	Active channels	N. of bits		Input (Gb/s)	R-O win (ns)	Events in win	R-O (Gb/s)
CEDAR	F 1GHz	60	256	(32)	$256 \times 8/\text{ns}$		2000	50	(3)	100
Gigatracker	TDC	1000	$18\text{K} \times 3$	6	20	12	192	50	50	10
Antis	F 40 MHz	10	2500	(3.3)	$2500 \times 8/25 \text{ ns}$		800	200	(2)	160
Straws	TDC	20	$2\text{K} \times 6$	33	18	14	21	400	8	8.5
RICH	TDC	15	2K	40	21	11	19	50	1	1.3
IRC	F 40 MHz	40	48	(8)	$48 \times 10/25 \text{ ns}$		19	200	(8)	3.9
LKr	F 40 MHz	20	13.2K	(150)	$13.2\text{K} \times 10/25 \text{ ns}$		5280	200	(3)	1056
N. Hodoscope	TDC	20	32	4	21	5	2	50	1	0.1
MAMUD	TDC	15	2080	73	20	12	35	200	3	7.0
Mu Hodoscope	TDC	15	512	4	21	9	2	50	1	0.1
SAC	F 40 MHz	10	64	(8)	$64 \times 10/25 \text{ ns}$		26	200	(2)	5.1

Table 2: Detector data bandwidth estimates. Here “Design rate” is an event (rather than hit) rate. “R-O type” indicates the type of read-out: TDC or FADC with indicated sampling frequency; “Active channels” indicates the channel occupancy per event; “In rate” indicates the raw data rate from the detector digitizers, while “L0 rate” indicates the read-out data rate after L0 trigger, for the listed read-out time window size. The number of bits for time and channel information are estimated assuming that information is transmitted in multiples of 1 byte. Numbers in parenthesis are not relevant for the rate estimate of sub-detectors read with FADCs.

3 General overview

Some general concepts drove the design of the NA62 trigger and data acquisition system, which should:

- have an adequately high data bandwidth to cope with high kaon rates;
- not introduce significant degradation in the sub-detectors' time resolution;
- have a very high data acquisition reliability for all parts; this is arguably the most important (and less common) requirement: the probability that part of the system fails to deliver its data *without such fact being recognized by the system* should be kept below 10^{-8} for each sub-detector;
- record any data time-out and DAQ inefficiencies: while automatic error correction and data re-transmission is not mandatory, error checking and logging definitely is;
- provide a trigger based on few simple and reproducible cuts which can be shadowed in the offline analysis; such cuts should introduce as little correlation between different sub-detectors as possible;
- have a high trigger efficiency (say, above 95%) for the channel of interest (one single track and nothing else in time coincidence);
- maintain a reasonably low random vetoing (say, below 5%), and therefore have a very high online time resolution and double pulse resolution;
- provide ancillary triggers (normalization, minimum bias, accidental, calibration, debugging, monitoring) plus other (possibly downscaled) physics triggers;
- be able to record and store most of the sub-detector information for the few events of interest with minimal (possibly no) online zero suppression;
- be able to record and store most of the information used in the formation of the trigger;
- be flexible enough to allow for emerging requirements in later stages of the experiment;
- possess some significant scalability in terms of bandwidth, to be able to cope with increased beam intensities;
- exploit as much as possible existing solutions developed or under development for existing or future HEP experiments;

- be as uniform as possible for the largest possible number of sub-detectors, in order to minimize design and maintenance efforts.

The general outcome of the discussions of the above requirements led to the concept of an integrated fully digital trigger/data acquisition system in which separated digitization and data paths for the trigger branch are avoided, the trigger being based on the main digitized data stream and performed as much as possible in a farm of general purpose processors interconnected through high speed links (TDAQ farm). A completely digital implementation of the trigger has the distinct advantages of working on the same data which is used for event reconstruction, eliminating the need for a separate trigger branch, and being fully reproducible and monitorable offline.

The request to avoid or limit online zero suppression entirely poses of course a significantly high strain on the DAQ system: this requirement follows from the very nature of the experiment, which is based in an essential way on its high vetoing capability, so that data corruption and undetected read-out failures should be avoided as much as possible, this being more easily done if all the sub-detector data for an interesting event is available at the offline level.

The difficulty of implementing online a zero suppression scheme (mostly for large FADC-type sub-detectors such as the LKr calorimeter) capable of reducing in a significant way the data load while not degrading the sub-detector performance and – most importantly – not compromising in any way its vetoing capability suggests to avoid such approach entirely. A design with no zero suppression would allow the (few) final candidate events selected in the offline analysis to be scrutinized in every detail, reducing the possibility that small signals from additional particles get accidentally erased during the data collection stage by any malfunctioning. On the other hand any final event “suspected” of some malfunctioning can be discarded without significant penalty to the experiment (provided the fraction of such events is relatively low, of course).

It is worth stressing the fact that, wherever no viable existing solution for a specific sub-system can be identified which can satisfy the requirements of the experiment, synergy with developments being carried out in the HEP community for next-generation experiments should be pursued. While it is clear that dedicated solutions will be required in some case, avoiding unnecessary duplication of efforts in the development of parts of the TDAQ system which can find more general application in other experiments should

be considered as a key point.

4 Trigger concepts

A “triggerless” system is one in which all data produced by the front-end digitizers is passed to a processor farm (where of course some trigger selection is performed to reduce the bandwidth to a level suitable for tape logging); the advantages of such a system are that all data reduction is performed in software, with the related benefits in control and configurability.

A fully triggerless solution was considered for NA62, but it is at odds with the above discussed requirement of avoiding any simple-minded zero suppression algorithm at the hardware level: considering the LKr calorimeter (as the largest FADC-type data producer, for which zero suppression is most challenging), the raw data rate for 40 MHz sampling (and 8-bit quantization) is 530 GB/s (most of it being pedestals, admittedly), and any triggerless solution would indeed require zero suppression (or the merging of cell information) to be able to deal with a manageable transfer rate to a PC farm.

For the above reason a system with one (L0) hardware trigger stage before read-out was chosen instead ². It is assumed that on each positive L0 trigger all sub-detectors transfer their data to the TDAQ farm.

It should be mentioned that a different approach could be taken, in which some sub-detectors do not read-out their data to PCs on each L0 request but rather do it only later, say on a subsequent L1 request (at reduced rate). This would clearly further reduce the readout bandwidth requirements for sub-detectors with large data load (*e.g.* the LKr), at the expense of some complications: first of all, the sub-detectors with “delayed” read-out would not be usable in some later trigger stages performed in PCs (*e.g.* L1); the local buffer space would have to be increased according to the (longer) expected latency of the higher level trigger, and – most importantly – an absolute maximum for such latency should be defined, which might turn out to be not easy for the intrinsically asynchronous environment of a switched TDAQ farm, for which average processing time constraints can be easily met but the avoidance of timeouts is a harder issue. Finally, the amount of

²It should be remarked in this respect that with the flexibility provided by today’s field-programmable components the distinction between “software” and “hardware” is pretty much reduced, and a huge flexibility is present also in the “hardware” domain, thus allowing no hard commitment to specific trigger algorithms from the beginning.

rate reduction in trigger levels higher than L0 should be precisely estimated in order to correctly design the front-end system, thus requiring an earlier commitment to the figures obtained. As a baseline option it is therefore assumed that all sub-detectors are read-out into PCs at the L0 rate.

We define three abstract trigger levels:

- **L0**: the first trigger level (before read-out to the TDAQ farm) is the only one which might be implemented in hardware, and defines the read-out rate of all sub-detectors to PCs. Data will be stored in front-end buffers (which might even actually be PCs as well, of course) during the evaluation of the L0 trigger conditions, and is expected to enter the TDAQ farm on a positive L0 request. Sub-detectors involved in L0 trigger formation will locally elaborate their data and send to a central L0 processor the information on time values around which their trigger requirements were satisfied, possibly together with a small amount of additional information (primitives) concerning it (such as multiplicity, energy, etc.). The central L0 processor will match such partial trigger time sequences among all the involved sub-detectors in appropriately chosen time windows, and broadcast the information on the times at which the overall L0 trigger conditions are satisfied back to the front end systems, to initiate readout.

The input rate to the local L0 trigger processors dealing with a single sub-detector is the raw event rate, in the 10 MHz range, its output rate from the central L0 trigger processor is the experiment read-out rate to the TDAQ farm.

Both the transmission of L0 information from the front-end systems to the L0 central processor, and that of L0 trigger requests from the latter to the former could be either synchronous or asynchronous with respect to the events: this is of course a relevant issue for the design of the front-end systems. For the expected low-occupancy conditions, the elaboration of L0 information is more conveniently performed in a highly multiplexed way, leading naturally to an asynchronous system; asynchronous transmission of such information to the L0 central processor is therefore most natural (and would also provide a higher transmission efficiency in case a packet-based transmission protocol would be used for trigger primitive transmission).

Each and every sub-detector should receive L0 triggers, and is required to react on them by sending at least some information (which could

be an empty frame, or an error frame) to the TDAQ farm. For what concerns the delivery of such L0 trigger requests, synchronous transmission would significantly reduce the bandwidth requirement for this signal path (by eliminating the need for trigger time information, which would be intrinsically defined by the trigger delivery time) and simplify the storage and extraction of the data from front-end buffers, at the price of requiring re-synchronization in the L0 central processor. Asynchronous transmission would have the advantage of allowing an intrinsic time consistency check already at the level of front-end read-out: this is not possible (nor really required) in a synchronous L0 dispatch scheme. Qualified L0 triggers would allow a selective readout of different parts of sub-detectors for different trigger types, at the expense of transmitting more information to the front-end systems; if zero suppression is not performed (at least not in the front-end systems) the identification of different trigger types at L0 is not strictly mandatory, but remains a highly desirable feature.

- **L1:** all trigger levels after L0 are assumed to be performed in software within commercial processors: we logically distinguish L1 as that part of trigger which is evaluated by considering separately the data from each (or some) sub-detector, before event-building has taken place, and eventually checking the match of a small set of single-sub-detector conditions in a L1 central processor, leading to either the release of data buffers or to their transmission to event-building nodes of the TDAQ farm.

The L1 trigger is inherently asynchronous and qualified, and each sub-detector system is allowed to respond differently to different L1 trigger types, transmitting different amount of information to later stages; this means that the trigger decoding logic is allowed in principle to discard parts of the data in response to L1 requests. As for L0, in no circumstance a sub-detector is allowed to completely ignore a trigger request nor to avoid sending at least an empty event frame (possibly including just an error word) to later stages in response to a L1 trigger. Delivery of L1 trigger information can occur via standard network links.

- **L2:** the L2 trigger is that part of the software trigger performed on the TDAQ farm on complete events. Note that for performance reasons this level could be actually split in a hierarchical cascade of algorithms in which only some sub-detector data is actually decoded and used,

or in which progressively more refined reconstruction procedures are used: this is an implementation detail which makes no difference in the logical scheme. The L2 trigger is also asynchronous and qualified, allowing for selective readout as L1, with the same requirements and transmission mode.

The output rate of the L2 trigger should match the data logging rate, in the tens of kHz range.

The main sub-detectors involved in the formation of the L0 trigger will be the RICH (positive identification of a charged particle in the acceptance), the LKr calorimeter and/or the neutral hodoscope ($K_{\pi 2}$ vetoing) and the muon detector ($K_{\mu 2}$ vetoing).

In the P-326 proposal positive identification of the track was originally intended to be performed by a charged hodoscope, reducing the rate due to downstream K decays³ to 70%. Due to its similar acceptance, the RICH itself also reduces such rate to about 65% of the total.

Figure 15 shows the effect of the RICH and a charged hodoscope on the raw rate; despite the limited effect on the rate reduction, this contribution to the L0 trigger is of course essential to provide the time definition for the event against which veto cuts are applied.

Particle multiplicity information from the RICH might be used in principle to further reduce the trigger rate, but more than 90% of the events hitting the RICH or hodoscope have only one hit within acceptance, so this appears to be not too useful in presence of a good online time resolution.

Since muons are the highest source of rate in the detector, active $\pi - \mu$ discrimination in the RICH would be helpful at the trigger level, as the number of μ hitting the RICH is at least 3 times that of π . Such discrimination might be based on the measurement of the Čerenkov ring radius for a particle of known momentum: with a maximum $\pi - \mu$ radius separation of 4 cm and a 1 cm (offline) fit resolution one could imagine using a 1-bit map of hit PMs, continuously sampled at 500 MHz, to be transformed with a fast algorithm into an accept/reject verdict; unfortunately the momentum and angle distributions of pion and muon tracks are rather similar, thus requiring spectrometer momentum information to obtain particle identification. While this would help in reducing the muon background at trigger level, the complexity of an online ring fitting algorithm [7] and – even more important – the

³All the following estimates are based on the rates due to K decays downstream of the final collimator, thus not including the muon “halo” from π and upstream K decays.




Figure 15: Background rate reduction obtained by requiring a RICH hit (dashed line) and a charged hodoscope hit (dotted line); muon “halo” is not included.

requirement of synchronized online data communication between RICH and magnetic spectrometer (and online reconstruction of the latter) suggested not to pursue this approach further.

The suppression of $K_{\pi 2}$ trigger rate relies on photon detection: this is mostly achieved in the LKr calorimeter, as shown by the comparison of the photon multiplicities in different sub-detectors in fig. 16.

The charged pion is expected to produce a shower in the LKr calorimeter in a significant fraction of the signal events, thus background suppression can be based on the cluster multiplicity; more than 25% of all background events have more than one particle within the LKr acceptance.

In order to avoid the implementation of relatively complex cluster counting algorithms at L0, the information on the energy deposit in large areas of the sub-detector can be used. By considering the energy release in the four quadrants of the LKr it is possible to achieve a significant suppression of the main $K_{\pi 2}$ background. The LKr quadrant multiplicity for such events

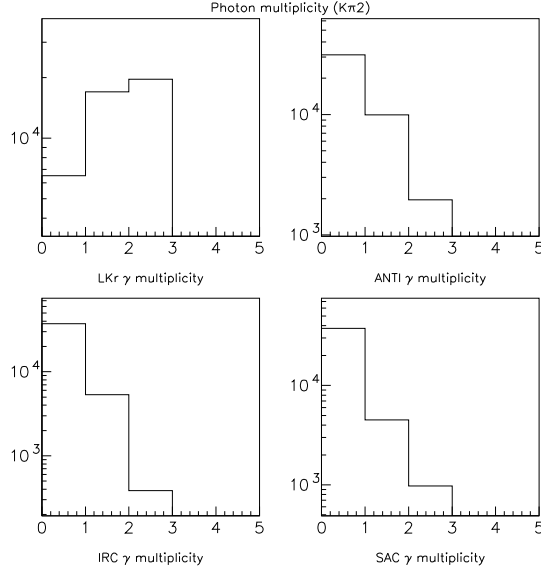


Figure 16: Photon multiplicities in the LKr calorimeter, large angle ANTI, IRC and SAC for $K_{\pi 2}$ decays.

(with the pion within the RICH acceptance) is shown in the top half of figure 17; the naive estimation is that the requirement of having not more than a LKr quadrant with energy deposition would reduce the $K_{\pi 2}$ background by a factor 3 (a factor 15 when on top of the RICH condition).

The above however does not take into account the lateral extension of the showers, which can increase the quadrant multiplicity for the background (the bottom part of fig. 17 shows the background quadrant multiplicity for a shower radius of 10 cm, for which the $K_{\pi 2}$ rejection factor rises to 5 (30 on top of the RICH) but also the signal efficiency decreases.

A safer minimal requirement might be the absence of significant energy release in two opposite quadrants of the calorimeter; due to the presence of the beam pipe the signal efficiency of such a condition is expected to be high ⁴. Figure 18 compares the rate reduction obtained by imposing the above conditions to the one for the quadrant multiplicity discussed above,

⁴Of course if the fine-grained LKr calorimeter data is used much more elaborate schemes can be foreseen.

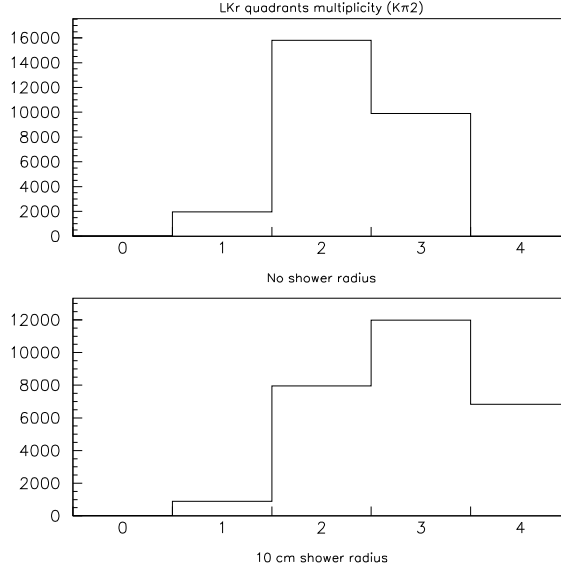


Figure 17: LKr calorimeter quadrant multiplicities for $K_{\pi 2}$ decays with the charged pion within the RICH acceptance. Top: zero shower radius; bottom: 10 cm shower radius.

showing their comparable rejection power: the opposite quadrant condition (zero shower radius) reduces the $K_{\pi 2}$ rate by a factor 2.8 (9.6 on top of the RICH), corresponding to a factor 1.3 (1.4 on top of the RICH) overall, and is considered as the default choice in the following.

Note that the impact of the LKr veto condition on the overall data rate (on top of the RICH condition) is however rather limited, reducing it by a factor 0.7 for the safer opposite-quadrant condition (the corresponding figure for the quadrant multiplicity cut being 0.5).

In principle the same information on energy deposition in quadrants could be obtained from the neutral hodoscope: the great simplification due to the analog summation and limited number of channels are partially offset by the a reduced vetoing power due to the fact that the energy threshold would have to be significantly higher. With a 15 GeV energy deposition threshold per quadrant the overall effective vetoing power of the condition implemented with the neutral hodoscope on $K_{\pi 2}$ events is about a factor 3 lower with




Figure 18: Background reduction obtained with requirements on LKr calorimeter quadrants energy deposit (one charged particle within the RICH acceptance). Dashed lines: no energy in opposite quadrants; dotted line: no energy in more than one quadrant. Case of zero shower radius.

respect to when the (very low threshold) LKr is used (a factor 4 reduction on top of the RICH).

Of course also for the LKr the online energy threshold for quadrants cannot be expected to be as low as the detector capabilities would allow, and this would of course reduce its vetoing power: while the above figures neglected this, the LKr vetoing power on $K_{\pi 2}$ (on top of the RICH) can be reduced by 25% for a 2 GeV threshold.

A better estimation of all the above rough figures requires an improved simulation.

The baseline L0 condition thus requires the presence of an active RICH signal with tight time vetoing of the muon detector and the opposite-quadrant LKr signals: the rate reduction depends significantly on the muon veto system performance and acceptance match to the RICH: in the following $\bar{\mu}_M$ will indicate muon vetoing from the MAMUD acceptance and $\bar{\mu}_H$ muon vetoing

from a detector with a $30 \times 20 \text{ cm}^2$ central hole. Neglecting time resolution effects and assuming perfect muon vetoing the maximum (geometrical) rate reduction factor from the above L0 condition is about 20 (11) for $\bar{\mu}_H$ ($\bar{\mu}_M$), corresponding to an average on-spill trigger rate of 340 (630) kHz.

The overall rate reduction factors which can be expected by adding to the above condition other requirements can be read in table 4.

Condition	Average rate	Condition	Average rate
$HODO * LKrQX * \bar{\mu}_M$	710 kHz	$HODO * LKrQX * \bar{\mu}_H$	200 kHz
$(1M) = RICH * LKrQX * \bar{\mu}_M$	630 kHz	$(1H) = RICH * LKrQX * \bar{\mu}_H$	340 kHz
$(1M) * HODO$	480 kHz	$(1H) * HODO$	180 kHz
$(1M) * \overline{ANTI}$	410 kHz	$(1H) * \overline{ANTI}$	120 kHz
$(1M) * \overline{SAC}$	610 kHz	$(1H) * \overline{SAC}$	310 kHz
$(1M) * \overline{IRC}$	600 kHz	$(1H) * \overline{IRC}$	300 kHz
$(1M) * \overline{RICH > 1}$	620 kHz	$(1H) * \overline{RICH > 1}$	320 kHz
$(1M) * E(LKr) < 40GeV$	530 kHz	$(1H) * E(LKr) < 40GeV$	230 kHz
$(1M) * STRAW1 > 0$	430 kHz	$(1H) * STRAW1 > 0$	240 kHz
$(1M) * \overline{ANTI} * STRAW1 > 0$	220 kHz	$(1H) * \overline{ANTI} * STRAW1 > 0$	34 kHz
$(2M) = RICH * NHODQX * \bar{\mu}_M$	1.0 MHz	$(2H) = RICH * NHODQX * \bar{\mu}_H$	720 kHz
$(2M) * \overline{ANTI}$	670 kHz	$(2H) * \overline{ANTI}$	350 kHz
$(2M) * E(LKr) < 40GeV$	630 kHz	$(2H) * E(LKr) < 40GeV$	310 kHz

Table 3: L0 trigger rates estimate for perfect vetoing. $HODO$ = at least 1 hit in the hodoscope vertical plane; $RICH$ = at least 1 hit in the RICH (12-140 cm); $LKrQX$ = energy ($>MIP$) in opposite quadrants of the LKr; $NHODQX$ = energy (> 15 GeV) in opposite quadrants of the neutral hodoscope; $\bar{\mu}_{M,H}$ (perfect) muon veto from MAMUD or muon hodoscope plane.

On top of the LKr contribution, the large-angle photon vetos would help to reduce further the $K_{\pi 2}$ background rate by a factor 8 (a factor 3 on the overall L0 rate); their inclusion in the trigger is maybe not mandatory but useful, and should be kept as an option (which has of course implications for their readout scheme).

All the above assumed perfect online vetoing of muons, which is clearly unrealistic in view of the sub-detector time resolution and the presence of tails. This is an important issue, since with an effective rate of order 16 MHz on the muon detector the vetoing time window should be kept below 3 ns to limit random vetoing to 5%, while on the other hand any vetoing inefficiency for such a window results in a dramatic rate increase. The dependence of the L0 rate on the vetoing efficiency can be estimated on the simulated K decays and scaled up to account for the muon “halo” (assuming a ratio for the two contribution as listed in table 2.1) obtaining as a rough scaling $R(\epsilon) = R(0)[1 + 21\epsilon]$ for a muon veto system with inefficiency ϵ .

The L0 trigger latency defines the minimum size of the L0 buffers (and is thus limited in practice by the overall amount of memory in the front-end of the LKr calorimeter read-out system) and the time required for the (multi-tiered) time-matching algorithm of the L0 trigger. While a L0 trigger might be formed in some tens of μ s on average, the actual processing time to generate it is only limited by the L0 trigger latency in case the L0 processing would be shared between different trigger processors.

It should also be remarked that the vetoing power of the ANTIIs on $K_{\pi 2}$ decays is much higher for upstream decays, and thus it would be nicely complemented in the trigger by some minimal multiplicity requirement on the upstream straw chambers which, not requiring any track reconstruction, might be implemented in L0 more easily.

Given the large uncertainties in the above estimates, it should be concluded that some further work is required to ensure that a L0 trigger rate below 1 MHz can be achieved with some significant safety margin.

5 Data acquisition concepts

All the individual front-end systems shall be time-synchronized by the use of the same experiment-wide clock signal. The uniformity of the clock in each system is clearly of paramount importance to guarantee consistent time

measurements, and clock jitter limitation is a crucial issue to achieve the required time resolution.

Each event will be uniquely identified both by a progressive event number within the spill and by a timestamp; individual channel data is only required to provide time information relative to such (common) global time reference. One important point should be repeated, namely the fact that the time structure of events is basically uniform, rather than intrinsically periodic and linked to a clock as it happens in LHC experiments.

Event vetoing failure due to data transmission errors must be avoided to a very high level. This might be expected to occur through two basic error mechanisms: (a) failure to deliver data from some (part of a) sub-detector going unnoticed or (b) time mis-alignment between the data sent by different sub-detectors (or parts of them). Once the data is within the processor farm the error checking mechanisms of the networking infrastructure can be exploited to limit the rate of occurrence of such errors, but particular care has to be taken in the first part of the data path, where custom electronics is used and data transfer occurs between modules of different kinds.

Errors of type (a) are controlled by requiring that all sub-detectors (and all modules within a sub-detector read-out system) always actively respond to read-out requests (*i.e.* L0 triggers), even if they have no useful data for the corresponding event. Periodic DAQ integrity checks should also be performed asynchronously in an automatic way as a way of monitoring the live status of the readout system (and sections of it), and this information should be propagated downstream to tape together with event data.

Errors of type (b) are controlled by continuous burst-level clock alignment checking and error-detecting coding of timestamp information; this is important since data corruption in the trigger time information represents a single point of failure for what concerns vetoing (corruption in other parts of the data might not necessarily result in a vetoing failure).

Data from each sub-detector will be transmitted on high-speed links from the corresponding front-end systems to dedicated PCs (or PC clusters) which are part of the TDAQ farm (sub-detector Farm Entry Points, FEP). Re-formatting and L1 trigger computation will take place in such sub-detector specific PCs.

Downstream of the FEPs the data will be handled inside the global TDAQ switched farm, where event building, L2 trigger, monitoring and data transmission for tape logging will be performed.

6 A possible implementation

In this section a possible implementation of the TDAQ system for NA62 is discussed. Following a general principle discussed earlier, such a scheme is based on the use of an existing general purpose read-out board (designed for LHCb) for several sub-detectors. This is not meant to be a complete design but rather a realistic scenario discussed in some detail, to help identify critical points and simulate discussion.

6.1 Design parameters

Here a few definitions and values for some system-wide parameters are listed.

- **Burst:** the unit of data-taking period is the SPS burst (spill), whose duration is not specified and can usually vary in the range 1-20 s, but is (roughly) constant during each run. Event numbering and timestamping are relative to a burst.
- **Run:** runs are just a convenient way of grouping bursts with uniform data-collection conditions, and have no other particular meaning in identifying the data; runs cannot overlap in time.
- **System clock:** an experiment-wide synchronized “40 MHz” clock will be used and distributed to each system. While such frequency matches the one used both by NA48 and LHC experiments, its exact value can be chosen arbitrarily. All systems should derive their internal clocks from this signal; they should count the number of clock pulses received during each spill, between precisely timed start and stop signals, and such count should be appended to the data to monitor and identify possible synchronization failures which might led to rejecting some bursts in the offline analysis. As mentioned, everywhere in the experiment the average frequency of the clock signal should be highly constant and uniform, and the time jitter should be limited to cope with the time resolution requirements. Local regeneration of the clock will have to be used to address the latter issue.
- **Time information:** the least significant bit for timing information is chosen to be $1/256$ of the main clock period ($97.65625 \text{ ps} = 1/10.24 \text{ GHz}$ in case of a 40.0 MHz clock). Individual channel time information is relative to the event timestamp and is encoded into a signed 16 bit

word, thus allowing for a maximum read-out window of $\pm 3.2 \mu\text{s}$ with respect to the trigger timestamp.

- **Timestamp:** The event (trigger) timestamp is an unsigned 32 bit integer, with a LSB equal to the period of the master clock (close to 25 ns) and the two MSB reserved, thus corresponding to a time range slightly exceeding 26.8 s. This information might be encoded in 40 bit (5 bytes) with 32 bit payload and a CRC-8 error-correcting word. The master timestamp associated to each event is defined by the central L0 trigger processor, and for each burst is in unique relationship with the event number. For the purpose of data extraction in response to a valid L0 trigger (trigger matching) each sub-detector is free ignore some lower bits of the timestamp. The timestamp is part of the event structure at all levels of data transport. In a scenario of synchronous L0 trigger broadcast, the local trigger timestamp information is locally generated at every sub-detector front-end system, rather than being transmitted together with a L0 trigger packet.
- **Event number:** this is actually the L0 trigger number, defined to be 24 bit wide, corresponding to more than 16 s at 1 MHz rate, plainly binary encoded (the first valid event number being 1). The event number is the key used for event building and (together with the burst ID or equivalently the timestamp) to uniquely identify an event within the entire duration of the experiment. The event number is not transmitted by the central L0 processor, and each sub-detector is responsible for inserting its locally computed event number into its data stream at the time of read-out (the same is true of the timestamp in case of a synchronous L0 dispatch scheme): any difference between sub-detectors and L0 central processor in the event-number to timestamp matching is an error condition.
- **Burst ID:** this is a number identifying uniquely the SPS burst to which an event belongs. The UNIX time of a conveniently chosen instant of the burst, encoded as a signed 32 bit number ⁵, is used as the burst ID. The burst ID is assigned centrally and broadcast to the whole TDAQ farm on the network; the details should be defined according to the expected working mode of the farm, which at any given time might be processing data belonging to different bursts.

⁵Despite the long NA48 tradition of intense and diverse physics program, the experiment is expected to last less than 136 years.

- **Run number:** this is a 32 bit unsigned integer only used to identify each continuous data-taking period with homogeneous conditions for practical purposes. The only requirement is that a higher-numbered run shall not contain bursts with lower burst ID than a lower-numbered one.
- **Trigger type:** the overall trigger type is coded into a 32 bit unsigned integer, of which the lowest 8 are reserved for the L0 trigger type (and are the only ones present in a L0 trigger broadcast), the next 8 to L1 and the rest for L2.
- **Channel numbering:** from the DAQ point of view each sub-detector is defined to have a maximum of 65535 channels (sub-detectors with more channels can be identified as different DAQ sub-detectors), so that channel number is never longer than 16 bit (it can be shorter, of course).
- **L0 trigger rate:** the (average, in spill) output of the L0 trigger should have a rate not exceeding 1 MHz, as a compromise between the read-out data rate and the complexity of the electronics: all sub-detectors are expected to transmit data to their dedicated FEP PCs at such an (average) rate. No minimum allowed timestamp separation between valid L0 triggers is defined (thus being 25 ns).
- **L0 trigger latency:** the L0 trigger decision shall be delivered by the L0 trigger processor to the appropriate receiver of the front-end systems not later than 1 ms after the time of passage of a particle in the RICH. The delivery of L0 triggers is guaranteed to be in proper time order.
- **L0 trigger primitives latency:** the sub-detectors providing data required to form the L0 trigger decision shall make such data available at the central L0 trigger processor not later than 100 μ s after the time of passage of a particle in the RICH. The arrival of L0 trigger primitives information to the central L0 trigger processor is not required to be time ordered, but all information related to a given timestamp should be delivered contiguously in time.
- **Guaranteed L0 trigger response:** all the active sub-detectors are mandatorily required to transmit a (possibly empty) data frame to their FEP for each L0 trigger request.
- **Rate choking:** the only kind of flow control foreseen is the TDAQ farm being capable of signaling an anomalous data congestion state to the central trigger processors, which can therefore momentarily suspend

the generation and dispatch of triggers, while also indicating this error condition by inserting appropriate markers into its own data stream for monitoring. The front-end systems should never lose data due to internal buffer overflowing, but in case this occurs due to some malfunctioning they should take care of recording the time of occurrence of such event and insert the corresponding information into their data stream, to be sent together with the data for the next valid L0 trigger, so that the failure can be recognized and used to mark a range of following events as possibly suspect of corruption or lack of data.

- **L1, L2 trigger rate:** not specified. The combined L1 and L2 rate suppression can be expected to be around 50, if a tape logging rate comparable to that of NA48 has to be achieved.
- **L1, L2 trigger latency:** not specified. This needs not be fixed, actually, as long as the distributed buffering capacity of the TDAQ farm is enough to cope with the fluctuations in processing times.

6.2 The TTC system

The TTC system [6] developed at CERN and used by all LHC experiments transmits the “40 MHz” global clock to the entire system; low jitter (< 50 ps) is expected to be recovered locally by the use of PLLs and uniform quartz crystals. Each front-end system is required to use this clock as a reference signal for time measurements.

The TTC system can also encode synchronous and asynchronous signals over the same link, for a total bandwidth of about 15 Mb/s. Start and end of burst signals will be encoded in this way and broadcast by the TTC system respectively some time before the beginning and some time after the end of each spill: this time interval will only be approximately constant for each spill (to the level of some ms), and might change in a significant way in different data-taking periods.

The clock will be transmitted optically via a TTC distribution tree: each sub-detector must receive and decode such signal, which can be conveniently done using TTCrx receiver chips.

Each sub-detector shall count the number of clock cycles received between a “beginning of burst” and an “end of burst” signal; the end-of-burst final count value should be latched, saved and transmitted into the data stream upon receipt of a special command on the trigger link (the TTC again), to

allow an experiment-wide synchronization check.

The clock signal is not guaranteed to be time-synchronous over more than a single burst; it will be available both in and out of spill, and usually also while data-taking is not taking place, but each sub-detector should foresee ways of producing its own “local” clock, to allow debugging while the global clock system is not available.

The exact frequency (close to 40 MHz) will be defined by the clock source. The clock is not required to be synchronous with the SPS operation, although the use of the ramping SPS clock (SPS RF frequency divided by 5) as a source might be a convenient choice. In any case if an external (*i.e.* one not under direct control of the collaboration) signal is used the switching between such a source and a free running experiment-produced one must be foreseen for testing purposes.

The data transmission capabilities of the TTC are also attractive for the transmission of L0 trigger information, avoiding in this way the need for a second distribution network besides the (mandatory) one for the clock. Since the data bandwidth of the TTC for user data is limited, this can only work at the proposed rates if using a synchronous L0 trigger dispatch scheme (*i.e.* triggers sent with a fixed time delay with respect to the event), thus avoiding the transmission of timestamp information. A limited amount of (slightly delayed) trigger type information can be sent for each L0 trigger by using a suitable driver ⁶.

In case of asynchronous transmission of trigger primitive information, the required bandwidth would be estimated to be of the order of 10 MHz (average event rate) \times 6 bytes (5 bytes timestamp + 1 byte trigger information) = 60 MB/s from each sub-detector involved, with the bandwidth for L0 triggers 1 order of magnitude lower. For synchronous transmission the corresponding figures drop to 10 MB/s (and 1 MB/s); of course nothing prevents from having one of the two signal paths working synchronously and the other one asynchronously.

In the proposed scheme L0 trigger primitives are actually sent asynchronously as timestamped data on dedicated links, while L0 triggers are broadcast on the TTC link to all clock-receiving devices, as single-cycle synchronous pulses, and followed by 8 bit of trigger type information.

Readout controllers are responsible for latching local timestamp information upon receipt of a L0 trigger (to be included in the event data stream)

⁶Such as the TTC-ci developed for CMS.

and for initiating the transfer of the relevant data around such trigger time to the sub-detector FEPs. L0 triggers cannot be ignored by sub-detectors participating to the TDAQ system, although different response depending on the trigger type is allowed.

A simple way of monitoring the live status of most of the entire TDAQ system can be envisaged by the use of specially flagged L0 triggers produced at well-defined times which induce the generation and propagation of special control patterns from each sub-detector system. If all the following trigger stages are arranged in such a way that the pair of closest events of this kind occurring before and after any selected event are always kept in the data stream (down to the offline level) this leaves the possibility to check that no anomalous TDAQ failures occurred close in time to any specific event appearing in the final sample.

6.3 The TELL1 board

The TELL1 board [9] was developed as a common system for data pre-processing, buffering and trigger processing for the LHCb experiment: it is used there by almost all sub-detectors (for a total of about 300 boards) and is going to be supported throughout its lifetime, which should overlap with that of NA62.

It is a VME-9U board containing 5 FPGAs⁷, an embedded PC with an ethernet link, DDR memory, TTC clock/trigger [6] receiver and 4 standard 1-Gbit ethernet links for outputting the data to PC-farms. Half of the board is free for installing custom daughter cards. In very crude terms the TELL1 implements a data path from 4 input sections to one output section, along which data processing and buffering is possible.

Power consumption is below 50W (depending on the daughter cards) and 21 boards can fit in a standard 9U VME LHC-type crate⁸. The board only uses the crate for power, as it has no VME capability and its slow control happens through the embedded PC (via ethernet).

The board can accomodate up to 4 mezzanine daughter cards with 200-pin connectors, each one being connected to one dedicated FPGA; two types of input mezzanine boards were developed by LHCb:

⁷Altera Stratix I 1S25 with 25K logic cells (upgradable to 80K) working at 200 MHz.

⁸Crates with a single power-only J1 connector are required, which are available from Wiener through CERN.

- a digital optical receiver card (double-size, occupying two of the four connectors) with 12 optical link receivers⁹ running at 1.6 GHz (80 MHz \times 16 bit per link);
- an ADC card (single-size) with 16 10-bit 40 MHz flash ADCs plus line receivers and pre-amplifiers.

At least another daughter card is required for NA62, providing a suitable number of TDC channels with 100 ps LSB; a prototype card of this kind was built in Mainz [10] using the 8-channel TDC-GPX by ACAM [11] (81 ps LSB, 32-fold multi-hit capability, 10 MHz continuous rate per channel, 40 MHz per chip), with 4 TDC \times 8 channels per card (128 channels per TELL1 board).

Another TDC daughter card with much higher integration has been designed in Pisa and is currently under test, based on the 32-channel CERN HPTDCs [12] (100 ps LSB, multi-hit capability, 4 MHz rate per channel), with 4 TDC \times 32 channels per card (512 channels per TELL1 board), which should lead to a significant cost advantage. A TELL1 board equipped with these cards will be 2 units wide, only allowing a maximum of 10 TELL1 boards in a crate, but this is not expected to be an issue. 8 high-density miniature connectors are used to bring differential signals to the card.

A 16-channels 1GHz 8-bit FADC mezzanine card (two units wide) was also proposed [13] for the read-out of the CEDAR, allowing a total of 32 channels per TELL1 board.

The development of a more performant output mezzanine card (*e.g.* one using a single 10 Gbit/s link) does not appear too useful, given the limitations in the internal data bandwidth of the TELL1 and the fact that distributing the output bandwidth on multiple cards can actually be valuable (*e.g.* for trigger primitive transmission).

The data I/O capabilities of the TELL1 board are as follows:

- each of the four input sections can receive data at a maximum rate of 960 MB/s;
- the board can output data at a nominal maximum rate of 4.8 Gb/s towards the mezzanine ethernet board. The actual output rate performance depends on the data packing, favouring large (>1 kB) packets.

The amount of processing possible in the TELL1 FPGAs has to be evaluated on a case by case basis: as an indication, pedestal subtraction, calibra-

⁹CERN Gigabit Optical links (GOL).

tion, zero suppression and data formatting were foreseen to be performed in LHCb (using up to 80% of the FPGA resources), with a 1.1 MHz maximum trigger rate.

The memory available on the card is 3×256 Mbit DDR SDRAM chips for each of the 4 sections of the board, for a total of 384 MB per TELL1 board¹⁰, working at 120 MHz.

The TELL1 also has some user-defined LEMO inputs and an RJ-45 connector which intended to be used for data flow control. The presence of an embedded PC on the card allows to perform configuration and monitoring tasks, as well as to implement a (very reduced data rate) standalone readout system for testing, debugging and monitoring purposes.

The cost of one TELL1 board (with no input daughtercards) is about 4.6K CHF. The components which might present earlier some availability problem are the TTC interface chips and possibly the output Gbit card.

The possibility of designing an improved version of a TELL1-like general purpose board on the time scale of the experiment is also being discussed both within and outside the collaboration, and should be followed closely.

6.4 Sub-detector readout systems

It is assumed that most sub-detectors using TDCs (with the exception of the Gigatracker) will use TDC-equipped TELL1s. Such detectors should provide differential LVDS discriminated signals of duration not shorter than 10 ns to the input of the TDCs using high-quality cables. The TELL1 boards should be housed in crates placed at short distances (order 5 m) from the corresponding sub-detector to minimize signal degradation: the time performance of such a system is one of the first issues to be assessed in tests.

The use of the existing or suitably modified LHCb FADC daughter cards is also an option for NA62: as an example, with 64 channels per board, about 210 TELL1 (10 crates) would be required to handle the whole LKr calorimeter data (a number comparable to that of the present CPD system of NA48).

The sub-detector FEP PCs can be placed anywhere: for practical reasons they would best be located upstairs close to the control room, but this would probably require the use of optical ethernet (and translation from copper in case of TELL1) before the FEP, which in some cases might not be desirable.

¹⁰This amount could be doubled by replacing the memory chips.

A rough channel count results in the following figures (an additional TELL1 board for RICH trigger computation is listed separately):

Sub-detector	Crates	TELL1
CEDAR	1	1
Straw tracker	6	16
RICH	1	4+1
N. hodoscope	1	1
MAMUD	1	5
Total	10	28

6.5 CEDAR

Due to its high channel rate, the CEDAR will most likely require FADCs with relatively high sampling frequency in order to obtain the required double-pulse resolution. The clock for the above can be generated locally with small jitter and synced to the one from the TTC. With 1 GHz 8-bit FADCs the raw data rate from the CEDAR is 256 GB/s, requiring of the order of 256 MB of L0 buffer space.

6.6 Gigatracker

Due to its challenging performance, space and radiation requirements, most of the Gigatracker electronics (pre-amplifier, TDC and trigger matching logic) will be implemented in custom VLSI very close to the detector. For this reason, and due to the fact that the sub-detector will not participate in the first trigger levels, its implementation is not discussed here; it is assumed that the front-end will send its data to the TDAQ farm upon a valid L0 request.

Channel ID and leading + trailing time information might be packed in 32 bits per active pixel, leading to a net estimate of 24 bytes per event.

Being the detector with the largest hit rate, the L0 latency parameter is relevant for the gigatracker: the raw data rate estimate for each one of the three stations is 8 GB/s, corresponding to a minimum buffer space of 8 MB (actually larger in order to cope with rate fluctuations), which is negligible.

6.7 Large angle vetos (ANTIs)

Inclusion of the ANTIs in the L0 trigger would require a good online time resolution, thus either a separate TDC (possibly with a mean-timer, depending on the sub-detector layout) or online processing of FADC data. Indeed, rather than implementing two analog signal paths with a fast-shaping TDC-equipped timing branch and a slow-shaping ADC-equipped pulse-height branch, FADCs might be used to provide both informations. In this case the sampling frequency is dictated by the signal shaping time: relatively high sampling frequencies (300 MHz) are required for using the fast PM signals, with some cost and data rate disadvantages; however, depending on the counter segmentation the expected rates on individual counters might allow the use of a signal shaping (stretching) stage without hitting limitations due to random vetoing: with single-counter rates not exceeding 1 MHz a shaping period compatible with 40 MHz sampling seems feasible. Such an approach, clearly attractive for reasons of cost and simplicity, poses some more difficulties for the use of the ANTI information at the L0 trigger stage: if the LKr calorimeter information is also used in the L0 trigger (see later) the ANTIs might however exploit the same solution.

Another possibility under consideration is that of using only TDCs (after suitable shaping of the PMT signals) and extracting rough pulse-height information from time of threshold; this would make the ANTI system easier to integrate in the L0 trigger and uniform with other sub-detectors.

6.8 Straw tracker

It is assumed the straw chambers front-end will have one pre-amplifier and discriminator per channel (LVDS output), with only leading-edge time being recorded; discriminated signals will be transmitted on copper to TDC-equipped TELL1s.

A cost estimate [10] for the straw tracker read-out electronics based on the TELL1 with TDC-GPX daughter-cards was around 500K EUR ($1600 \text{ TDC} \times 100 \text{ EUR} + 200 \text{ daughter-boards} \times 400 \text{ EUR} + 48 \text{ TELL1} \times 4\text{K EUR}$), or 40 EUR/channel.

With HPTDC chips (60 CHF each) the above cost estimate could be significantly reduced to about (very rough estimates) 160 (280) KEUR ($400 \text{ TDC} \times 40 \text{ EUR} + 100 (200) \text{ daughter-boards} \times 400 \text{ EUR} + 25 (50) \text{ TELL1} \times 4\text{K EUR}$), or 13 (22) EUR/channel, for 4 (2) TDC/daughter card (excluding

cables). With the HPTDC solution the system could be housed in 1 VME 9U crate per chamber.

6.9 RICH

The analog signals from the RICH PMTs will be locally connected to a fast amplifier/discriminator chip [14] (8 channels per chip) with LVDS outputs providing time-over-threshold pulses of duration 12-17 ns, suited to HPTDCs¹¹.

TDC-equippde TELL1 boards are used for read-out: digital signals from the front-end are fed to the TDCs, and the digitized times are continuously transferred from the TDC internal buffers to the TELL1 memories. Processing which has to be performed on this data include masking of individual channels, time alignment, adding coarse time information and re-formatting (some of this can occur already inside TDCs).

The number of TDC channels per board is practically determined by space considerations, as the input and output channel bandwidths arguably pose less severe constraints. With 4 (2) TDC (that is 128 (64) channels) per daughter card, the whole system would require 4 (8) TELL1 for data input, and the average data rates would be 610 (370) MB/s from each daughter card and 160 (97) MB/s out of each TELL1 (allowing in principle for a possible L0 trigger rate increase by up to a factor 3.7 (6.2) with respect to nominal), requiring 2 (1) output Gbit links, or 8 links in total, for an aggregate data rate of 640 (770) MB/s for the entire system (the system being developed has 4 TDC/card). The whole RICH system would be housed in a single VME 9U crate.

The overall event data size is estimated to be at most: [11 bit (channel ID) + 2×16 bit (relative fine time) = 6 bytes] \times 40 channels = 240 B/event + headers; with 256 B/event the overall aggregate read-out data volume for the RICH is estimated to be 256 MB/s.

For trigger purposes the RICH sub-system has to extract a single time-stamped information (trigger primitive) for each event satisfying the trigger requirements, which in the simplest implementation can be defined as a time coincidence of a minimum number of hits. This requires time matching between hits, which must be performed both within a single TELL1 board

¹¹The same preamplifier-shaper-TDC chain is used for the TOF system of the ALICE experiment [22].

among 512 (256) channels, and then between the 4 (8) individual boards.

The first matching task is performed within the TELL1 FPGAs, generating a data rate which can be estimated as $15 \text{ MHz} \times 5 \text{ bytes} = 75 \text{ MB/s}$, which can be transferred out of the board on one Gbit link. Several possible algorithms are being considered for this purpose [17], and the most scalable one appears to be one based on a continuous synchronous time-binning and time-reordering of incoming data.

The second task (matching among different boards) can be performed in a concentrator unit receiving 4 (8) streams of trigger primitive data from each TELL1, for a (maximum) total of 300 (600) MB/s, and sending (maximum) 75 MB/s of data towards the central L0 processor. Such concentrator unit might be either a dedicated PC or another TELL1 board; in the latter case a large part of the internal structures of the TELL1 algorithms for this board would be the same as for the digitizing boards, and the communication between TELL1 boards might be done using the Gbit ethernet ports; however, even if the entire system is housed in only 5 TELL1 boards, a custom output link card should be developed to transmit the global RICH trigger information out to the central L0 processor, as the TELL1 has only 4 Gbit links.

An alternative approach, still using the TELL1 for trigger primitive aggregation, is that in which the boards are daisy-chained and trigger primitives are sent from one to the next, occupying only two Gbit links for this purpose: data bandwidth considerations indicate that in this case however each TELL1 should not simply append its own trigger data to the incoming stream from the previous, but rather actively merge it together (*e.g.* performing multiplicity sums) to keep the trigger data rate approximately constant; this is the default scheme assumed in the following.

The data-collecting TELL1 boards should respond to an incoming L0 trigger request (on the TTC link), perform trigger matching, send primitives and transfer data out of their buffer memories upon reception of a L0 trigger, to dedicated PCs (as many as required to accomodate the data rate).

6.10 LKr calorimeter

It is assumed that the LKr calorimeter cells will all be continuously flash-digitized at 40 MHz frequency as in NA48; the somewhat reduced dynamic range of photon energies in NA62 with respect to NA48 points to the possibility of reducing the data word size while maintaining the desired energy

resolution: for the purpose of estimation a minimal scheme with (non-linear) 10-bit word and 8 samples (per cell per event) is assumed (the actual viability of such a scheme has of course to be assessed).

The above assumptions correspond to a continuous raw data rate of 675 GB/s out of the front-end digitizers into temporary buffers. The L0 latency will affect in the most severe way the LKr detector, requiring a minimum buffer space of 675 MB ¹².

With the NA48 signal shaping about 8 samples (at 40 MHz) are required for accurate time and pulse-height reconstruction, which corresponds to a read-out data rate to PCs (after L0) of 130 GB/s if no zero suppression is performed. Zero suppression is the most challenging issue in the read-out of a high resolution detector with a significant amount of channels, since if it cannot be done on a single-channel basis (*i.e.* thresholding) it requires data exchange with neighbouring channels (which must cross individual module boundaries), and requires additional buffering. As discussed above, the decision of not performing any zero suppression (at least not in hardware) on veto detectors such as the LKr calorimeter was taken early on as a consequence of the basic design principles of the experiment.

In order to transfer the required amount of data into the TDAQ farm high-bandwidth links (order of 10 Gb/s) are required. These in turn have to be matched to the available memory bandwidth into the PCs, suggesting the use of PCI-Express [19] as the system bus; this is now starting to become the standard on PCs: it is a scalable full-duplex bus specifying a bandwidth of (1× up to 16×) 200 MB/s in each direction; using 4× buses (PCs with 3 such slots are available now), or 800 MB/s (only one direction used) per link, (64 channels per link), a number of links of order 200 would be required for the whole LKr calorimeter, and a number of PCs of order 100 (the use of 10 Gb/s links on 8× PCI-Express might somewhat reduce the overall number of links).

It is clearly worth considering whether the above figures can be somehow reduced. Even if zero suppression is not performed in the hardware, the LKr system could profit in a very significant way from on-the-fly lossless data compression: a factor 10 reduction in data volume would basically eliminate the issue completely.

Some investigations were carried out [8] using actual LKr data from non-zero-suppressed events collected in NA48/2. Using general purpose lossless

¹²The NA48 CPD system used (10 years ago) 160 MB for such purpose.

compression algorithms (not at all optimized for the LKr data) the achievable data compression factor (regardless of encoding/decoding time) was measured and found not to depend significantly on the number of clusters (0 to 2 in the analyzed sample) actually present in the data. Out of the three compression algorithms tested in more detail a reduction of the data to 0.41 to 0.43 was obtained for two of them (bzip2 and 7-zip) and 0.56 for the third (gzip), with no significant change as a function of the “aggressiveness” of the required compression. Tests with on a total of 16 general purpose algorithms showed no better (but sometimes significantly worse) compression ratios.

The time required for compression (and – to a smaller extent – decompression as well) is of course an important issue; as an indication the average compression time per event was measured on a 2.8 GHz Intel Xeon processor to be around 8 ms (for gzip and bzip2) or 2 ms (for 7-zip), while the decompression time (on the same machine) was around 2 ms (for gzip and 7-zip) or 5 ms (for bzip2); the timing of dedicated algorithms running in hardware (at lower clocking frequencies) requires further study.

Given that a factor 2 can be obtained with a generic (blind) compression algorithm, it seems reasonable to expect that significantly better results could be achieved with an appropriately tailored approach (starting *e.g.* with simple differential encoding) in which the learning phase of the above adaptive algorithms is hardcoded. More work is required in this direction.

Another possibility is of course that of performing a “smart” (*i.e.* “safe”) zero suppression of LKr data, not an easy solution also considering the fact that such data will be necessarily split among different boards and therefore either inter-board communication is implemented or a partial (less performant) zero suppression system is used, based on some halo expansion algorithm working on a limited part of the calorimeter data (not suppressing data corresponding to the border regions of a board).

This task is performed somewhat more easily in software, *e.g.* within the (dedicated) LKr FEP PC of the TDAQ farm, but clearly in such case it would not alleviate the read-out rate issue, and only help in reducing the data load on the TDAQ farm itself. For this purpose the use of graphic cards as computing platforms (GPGPU) was considered [8], as they usually have higher processing power and memory bandwidth than general purpose processors, are cheaper and improve faster in time, and moreover they are designed to perform transformations which are quite omomorphic to the zero-suppressing “halo” algorithm on a bidimensional structure such as the LKr calorimeter.

A different approach has been proposed by the Roma 2 group [15], in which a fully interconnected hardware system buffers LKr calorimeter data and produces trigger primitives exploiting full interconnections between boards, but data is only read on custom high-speed links to PCs after a L1 trigger. Of course, since L1 is defined to be performed on individual sub-detector data, a sufficiently flexible and performing hardware system capable of communication with the TDAQ farm might actually participate in the formation of this trigger level, without modifying the proposed scheme and allowing a delayed data readout.

In any case it should be remarked that the evaluation of L0 trigger primitives from a FADC detector requires some special effort because precise time information has to be extracted first from a set of data samples, before it can be combined with that from other channels.

It is clear that more studies are required to converge on a full design for the LKr calorimeter read-out and trigger handling (and for FADC sub-detectors in general), for which other integrated solutions might be promising [21].

6.11 Neutral hodoscope

If used, the neutral hodoscope PMTs can be read out with the same system used for the RICH, by a single TDC-equipped TELL1 daughter-card. If the data is used in forming the L0 trigger such daughter-board can be housed on a TELL1 board also handling (trigger) data from other sub-detectors.

6.12 Muon veto

Both the fast muon veto plane and the whole muon veto system [18] is read by TDC-equipped TELL1s; the boards handling the data from the fast plane are interconnected among them and to the central L0 processor as described for the RICH.

6.13 Small-angle detectors

The detectors with small number of signals such as the IRC and SAC can be expected to use a limited number of channels of the same solution eventually adopted for the LKr calorimeter.

6.14 Central L0 processor

The central L0 processor should perform the final time matching of the L0 trigger primitives and produce a list of times in which the overall L0 trigger conditions are satisfied. Such list should be transformed into a synchronous time-ordered set of triggers, complete with locally generated timestamps, event numbers and trigger type words. Only the latter is broadcast, together with the single time-synchronous L0 valid signal, by using a suitable TTC transmitter. Timeout and error control should also be performed by this device, which also offers the possibility of centrally choking the L0 trigger rate in case of data congestion (such a situation is however considered an error condition).

The L0 central processor is also itself a data generator, as information on the received primitives and trigger decision (plus event number, timestamp and trigger type) should be produced and inserted into the data flow as an additional sub-detector.

The actual implementation of the central L0 processor is not defined yet: apart from a custom module some possibilities include a PC or a specially equipped TELL1 board.

6.15 Data format

Each subdetector will transmit the following data block information upon a read-out (L0) request; the data transmission format between a sub-detector L0 buffer and its FEP PCs is not specified, to allow sub-detector specific optimization, but it should at least include the local event number and timestamp, as well as the received trigger type word.

When entering the TDAQ farm infrastructure the data from each sub-detector should have a common format as follows (32-bit alignment is assumed at this stage):

32 31 30 2 1 0			
Detector ID		Event number					
Data block byte count							
Trigger type		Reserved		Reserved		Timestamp (high)	
Timestamp (low)							
Detector event data							
Reserved (trailer)							

The structure above comprises a 20 byte header/trailer (the data block byte count includes all headers and trailers).

6.16 TDAQ farm

The TDAQ farm itself will be formed by a cluster of PCs, interconnected via high-speed switches; the size and topology of the cluster is not specified here. High-bandwidth links will be required to transfer data from the sub-detector readout cards into the TDAQ farm (in the FEPs), and to move data within the farm itself as well: the two sets of links need not necessarily share the same implementation, although this would simplify the system. For the sub-detector to PC link a high-speed quasi-unidirectional point to point link is required for which a suitable driver device or FPGA core is readily available to be implemented on the read-out cards, allowing high customization capabilities; for the intra-farm link a commercial solution for which high-speed interconnect switches are available is necessary (*e.g.* 10 Gbit Ethernet, Myrinet, Infiniband); besides cost, speed and availability, important factors for the technology choice are also customizability and reliability of the link and drivers.

The Myrinet solution was investigated somewhat in more details: preliminary tests indicate that the bandwidth can be obtained with relative ease, and customizability can be high, although the CPU usage is significant; however, integrating a Myrinet driver in a custom front-end requires the implementation of a XAUI (10 Gbit ethernet) interface in an FPGA, which can be a relatively expensive and not too easy task: the advantage of using a commercial link appears to be lost when it is used in a non-standard environment.

While most sub-detectors using TELL1 might use 1 Gbit ethernet for the connection to the TDAQ farm, this is not an option for the LKr system which requires much higher bandwidth; a possible solution based on custom high-speed links is being developed in Roma [16]. In any case the use of PCI-Express [19] as a PC system bus is mandatory to achieve the required performance. Typical expected performance figures for the near future are 20 GB/s memory bandwidth and 2 GB/s input + 2GB/s output data rate from a single PC, while still leaving significant computing resources to the CPU [20].

Each PC will hold in its memory a series of buffers storing event fragments (before L2 trigger) or entire events (after L2 trigger): each such buffer

(uniquely identified within a burst by event number and sub-detector ID) will have to be explicitly freed (if L1 or L2 is not passed) or marked as valid for L1 and L2.

Among the TDAQ farm PCs some will perform the task of forming high level trigger decisions and broadcasting them to allow releasing of event (or sub-event) buffers in memories. This task might even be performed dynamically by different PCs (particularly for L2).

6.17 Slow control

Slow control of hardware devices is expected to be handled via standard commercial links (*e.g.* ethernet).

7 Conclusions

Some key issues and general specifications for a TDAQ system suitable for the NA62 were discussed above; this work should stimulate corrections, suggestions and the progress of design in order to converge towards a realistic scheme.

Apart from the check and correction of the sub-detector specifications, the finalization of the detector layout and a following more refined simulation, and technology choices, some of the most important open issues which remain to be addressed are listed below as a stimulus for further discussion.

- Which online time resolution can be achieved for the fast muon veto plane?
- Which is the expected hit distribution for the RICH? Which individual multiplicity threshold can be implemented on subsets of channels?
- Can the neutral hodoscope be used as a replacement for the LKr calorimeter at the L0 trigger stage? Which is the expected working threshold? Which is the expected online time resolution?
- How will the LKr calorimeter be read-out? How can a fast hardware online computation of the trigger primitives be implemented? Which online time resolution can be achieved? Can the LKr be read at L0 at all? Which bandwidth reduction can be expected from lossless data

compression? At which cost in terms of resources and compression/de-compression time? Which trigger algorithms can be implemented? Which is the rate dependence on the energy thresholds?

- How will the large-angle vetos information be digitized? Which online time resolution can be achieved?
- Which is the efficiency of the proposed L0 condition on the LKr calorimeter for the signal?
- Which time-matching algorithms can be efficiently implemented for the L0 trigger?
- Which algorithms should be foreseen in the TDAQ farm to implement L1 and L2 for achieving the required data bandwidth reduction? How should these be split?
- Which data processing model can be assumed for the TDAQ farm? Will data be processed synchronously during bursts or not?

References

- [1] P-326 collaboration – “Proposal to Measure the Rare Decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ at the CERN SPS” – CERN-SPSC-2005-013, June 2005.
- [2] N. Doble – <http://cern.ch/doble> and presentations to P-326 meetings.
- [3] Flyo MonteCarlo program (version 29.1.2007): <http://www.pi.infn.it/na48/flyo/flyoindex.html>.
- [4] N. Doble – Revised muon halo rate computation for HIKA+ beam, March 2007 (private communication).
- [5] M. Lenti – A RICH for the P326 set up, note P326-06.03, October 2006.
- [6] The Trigger, Timing and Control system for the LHC: <http://ttc.web.cern.ch/TTC/intro.html>.
- [7] M. Lenti – presentation at P-326 meeting, February 2006.

- [8] M. Sozzi – presentation at P-326 meeting, September 2006.
- [9] G. Haefeli *et al.* – “TELL1: specification for a common read out board for LHCb” – LHCb 2003-007 note, October 2003.
- [10] M. Hita-Hochgesand – presentation at P-326 meeting, February 2006.
- [11] ACAM: <http://www.acam-usa.com>
- [12] HPTDC, J. Christiansen et al., Digital Microelectronics Group, CERN.
- [13] B. Hallgren – “CEDAR readout using the LHCb TELL1 and 1 GHz ADC mezzanine board” – NA48 Note 05-02, May 2002.
- [14] F. Anghinolfi *et al.* – “NINO: an ultra-fast and low power front-end amplifier/discriminator ASIC designed for the multigap resistive plate chamber” – Nucl. Instr. Meth. Phys. Res. A **533** (2004) 183.
- [15] A. Salamon – presentation at P-326 TDAQ working group meeting, September 2007.
- [16] A. Salamon – presentation at P-326 meeting, December 2006.
- [17] E. Imbergamo – presentation at P-326 TDAQ working group meeting, September 2007.
- [18] V. Kurshetsov – presentation at P-326 meeting, September 2007.
- [19] PCI-Express: <http://www.pcisig.com>.
- [20] A. Hirstius – presentation at P-326 TDAQ working group meeting, February 2006.
- [21] L. Musa – presentation at P-326 TDAQ working group meeting, February 2005.
- [22] ALICE TOF system: <http://alice.web.cern.ch/Alice/Projects/TOF/>.