

2016 NA62 Status Report to the CERN SPSC

Abstract

NA62 aims to study the rare decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ at the CERN SPS. In this document we report the status of the detector construction and the general progress of the experiment since April 2015.



Contents

1	Introduction	4
2	Beam	5
3	Gigatracker (GTK)	7
3.1	Introduction	7
3.2	PCB carrier	8
3.3	Bump-bonding	9
3.4	Micro-channel cooling	9
3.4.1	Procurement of the next cooling plates from Leti	9
3.4.2	Hydraulic shock from the cooling plant to GTK_1	10
3.4.3	Reparation of GTK_2	11
3.4.4	Assembly of modules GTK_4 and GTK_5	12
3.5	TDCpix test	14
3.6	"Off-detector" DAQ status	15
3.6.1	Update on the "Off-detector" DAQ	15
3.6.2	Local PC farm performances	16
3.6.3	Upgrade local PCfarm (3 PC to 6 PC and L0 to L1)	17
3.6.4	Upgrading from 3 to 6 GTK-PC	17
3.6.5	Upgrading from L0 to L1	17
3.7	Results from 2015 data taking	17
3.7.1	Time Resolution	18
4	Large Angle Veto	20
5	Straws	24
6	RICH	26
7	LKr	30
7.1	CREAM readout	30
7.2	CREAM firmware modifications	30
7.3	Operation during the 2015 run	31
7.4	Calibration software and operation	31
7.5	Cal-L0 trigger	32
7.5.1	Installation and commissioning	34
7.5.2	Operation experience	35
8	CHOD	36
9	NewCHOD	38
10	Trigger and Data Acquisition	40
10.1	Common infrastructure	40
10.2	Common TDC-based TDAQ system	40
10.3	Other Systems	42
10.4	L0 Trigger system	43
10.5	On-line Farm and High-Level Trigger	45

10.6 Conclusions	46
11 Publications and analysis of older data	46

1 Introduction

The first physics run with an almost complete detector was carried out in 2015. Here follows a short summary on the status of each sub-detector after numerous changes and updates during the winter shutdown in preparation for the 2016 run:

- **KTAG:** The detector operates with nitrogen to design specification with an efficiency of 97.4% for detecting K^+ and less than 1 per mille pion contamination at the designed beam intensity of 750 MHz. A new gas system with pipework, that will enable operation with hydrogen if beneficial to the experiment, has been installed and the DCS and DSS systems have been updated accordingly. Verification of the tube welding has been performed using X-rays. A protocol is in preparation to ensure the safe filling and evacuation of the detector with hydrogen.
- **GTK:** During the 2015 run three detectors were gradually installed and operated. However, due to time constraints and technical difficulties, the detectors were read out only partially. The data taken are useful for performance studies and to correlate the GTK with the other detectors. The GTK is on the critical path, but the present plan is to install for the start of the run. The inlet pipes for the cooling of the three stations have been shortened in order to decrease the operation temperature of the detector.
- **CHANTI:** It has been working smoothly during the whole 2015 data taking period. The time resolution, monitored along the data taking, resulted steadily better than 1 ns. The single layer efficiency has been measured to be higher than 99% for straight tracks.
- **LAV:** A problem was discovered with some of the high voltage connectors. Discharges between the high voltage pins and the casing (at ground potential) were found during the 2015 run. All the connectors will be replaced with a new modified version.
- **RICH:** The two semi-hexagonal mirrors have been replaced with new ones with a slightly different geometry to better accommodate the passage of the beam pipe.
- **LKR:** All the high voltage connectors have now been changed.
- **IRC and SAC** were installed and used during the run. The detector read-out was connected to the CHOD electronics based on the standard TDC/TEL62 system.
- **HASC:** A Hadron Sampling Calorimeter in the very forward region has been installed and will be ready for data taking from the start of the run.
- **NewCHOD:** A new CHOD with a higher level of segmentation was built and has been installed. It will be ready once the data taking restarts.

During the 2015 run, we observed several failures of the primary power supplies for the HV crates (CAEN A4531) that are likely to have been caused by single-event effects induced by radiation. Such situations will be avoided in the future by enhanced control and monitoring measures, which are being implemented. The following actions have been taken to mitigate potential problems related to higher beam intensity conditions:

- The shielding in front of the CAEN power supplies in TCC8 has been increased.
- The sensitive equipment (PLCs) in the GTK cooling plant has been moved from TCC8 to ECN3 with the help of the EN-STI group and the R2E working group.
- Additional radiation monitoring devices have been installed for the 2016 run. These will enable us to understand better the type and intensity of the radiation in the zones of interest. In addition, the monitors are read on-line in real time. The shielding wall is completed.
- The blue wall mounting will be completed once the consolidation work on the T10 target is completed.

2 Beam

The K12 beam line was successfully commissioned during the 2015 run. It worked correctly at nominal intensity and according to specification. Also, the P42 beam line, transporting the protons that did not interact in the T4 primary target to the T10 target over a distance of almost 900 metres, operated correctly with the correct transmission. As an example, Figure 1 shows the profiles measured by filament scanners in the region of the muon veto counters, and Table 1 shows some comparisons of measured rates with simulations. The muon background rates in the main detectors are consistent with beam halo simulation results.

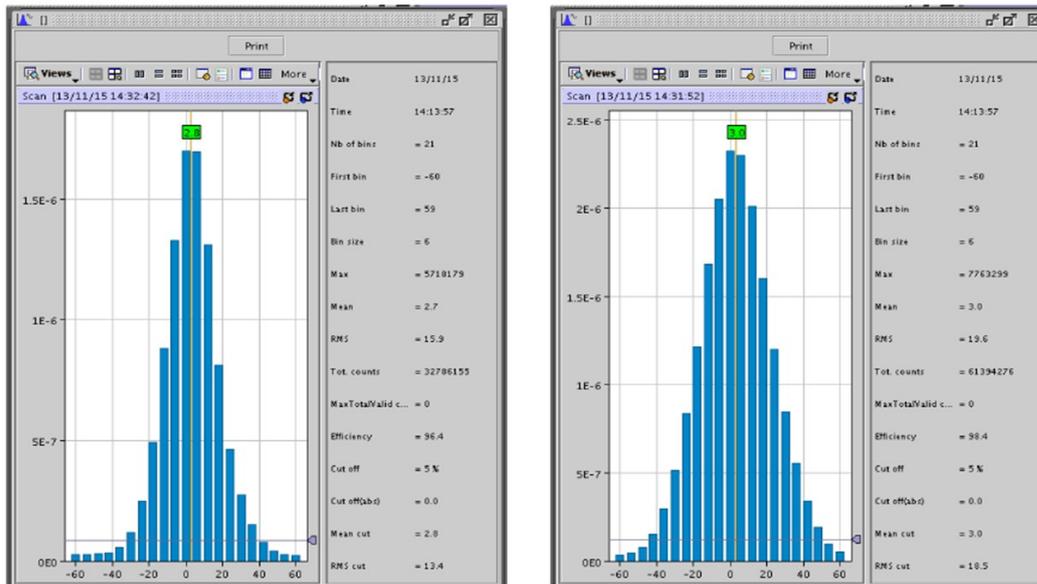


Figure 1: Beam profiles measured in the FISC filament counter downstream of the LKR calorimeter

During the year 2015, monitoring of the magnet currents was further improved. Some critical quadrupoles included in the already existing surveillance of dipole currents were also improved. Nevertheless, some radiation damage to the electronics of the experiment could be observed. On a number of occasions, the radiation damage (single event upsets

Parameter	Design value	Measured
Average K^+ momentum (GeV/c)	75	74.9
Proton flux on T10 target per 4.8s spill	$3.3 \cdot 10^{12}$	$3.3 \cdot 10^{12}$
Total flux per pulse	$2.25 \cdot 10^9$	$2.2 \cdot 10^9$
Nominal spill length (s)	4.8	4.8
SPS Duty-Cycle (flat-top/cycle times (s))	0.3	0.19-0.29
RMS divergence at KTAG horizontal (mrad)	0.07	0.07
RMS divergence at KTAG vertical (mrad)	0.07	0.07
Fraction of beam (%) p	23	22.4
Fraction of beam (%) K^+	6	6.6
Fraction of beam (%) π^+	70	70.2
Fraction of beam (%) μ^+	< 1	

Table 1: Comparison of the principal measured rates with the design values

or real damage) was concurrent with magnet trips in the P42 line, upstream of the T10 target. Therefore, a shielding wall has been installed, some 10 metres upstream of the T10 target and covering, with some margin, the section of the proton beam line tunnel. Most of this wall consists of 1.6 m of concrete, but some critical regions (around the beam and the top layer) are reinforced with iron. See Figure 2.

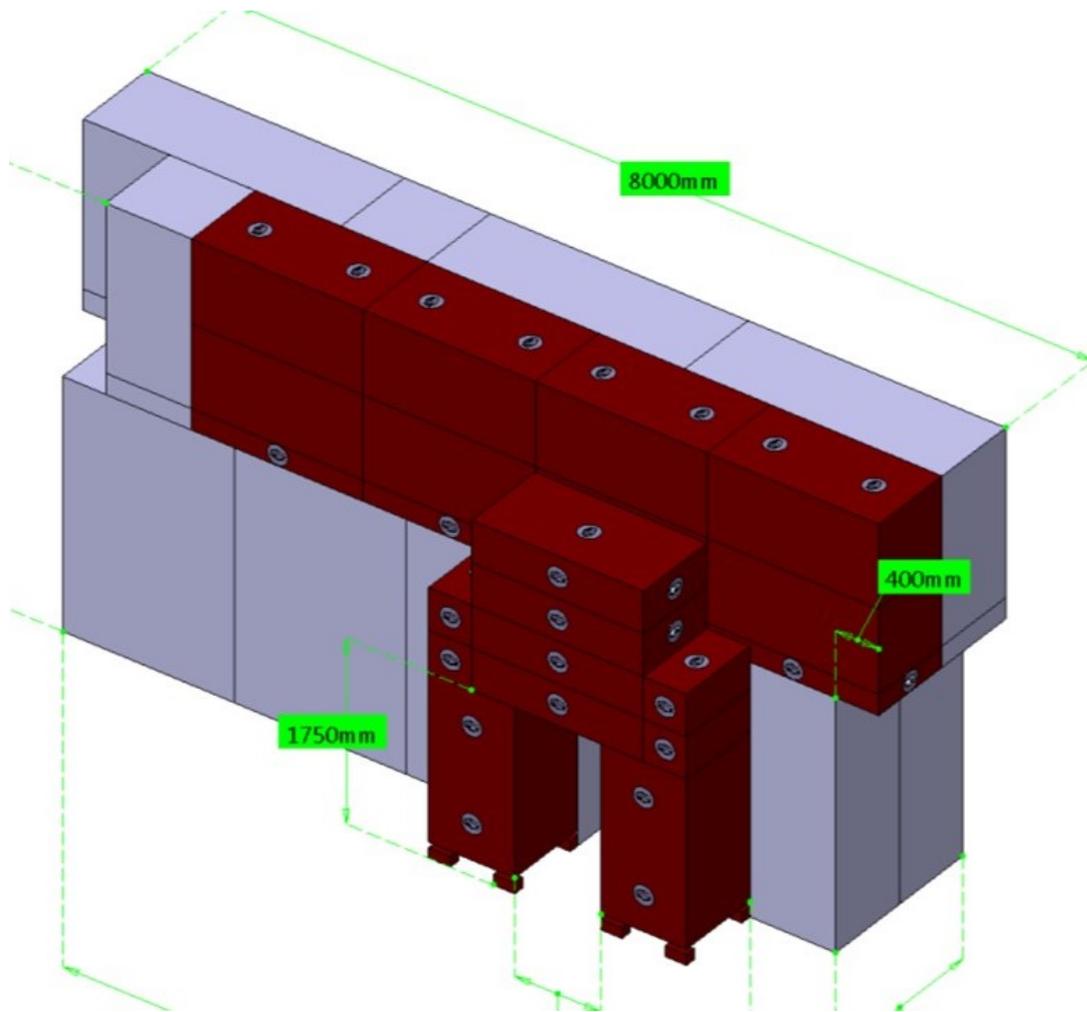


Figure 2: the new shielding wall just upstream of the T10 target

During the 2015 run, the primary targets in the TCC2 target zone, including the T4 target, developed a problem with the target head motorization. Investigations revealed an issue with a trapezoidal screw in the mechanism. A modification has been tested thoroughly and implemented on all targets, including the T4 and T10 targets. The air lock between the T10 target zone (TCC8) and the ECN3 experimental area, consisting of a double wall with over-pressure in between, has been modified to allow easier opening in case the crane must travel from one side to the other.

The infrastructure to allow the KTAG to operate with Hydrogen gas in the CEDAR has been completed. This included some maintenance work on the Hydrogen vent line and on the gas extraction system in the surface building.

3 Gigatracker (GTK)

3.1 Introduction

During the 2015 data taking, all the component parts of the GTK were installed and functional. The list of achievements follows.

- a. Three detectors were installed with thicknesses 450 (GTK1) and 100 μm (GTK2 and GTK3).
- b. The DAQ was operated in L0 trigger match mode. Thus, only the data associated to a given L0 trigger were read out and sent to the PC farm.
- c. The detector temperature was kept at about 0 $^{\circ}\text{C}$, with the cooling plant operating continuously to provide a 2 g/s flow of liquid C_6F_{14} at -25 $^{\circ}\text{C}$.
- d. The Detector Control System was operational.
- e. Data have been taken continuously and merged into the NA62 data stream. Analysis of the data is presented in section 3.7.

We observed an anomaly of the behavior of the pixels whose characteristics are briefly described.

- a. After a few months of data taking at low beam intensity, we observed that a fraction of the pixels (a few percent) were noisy, dispatching a hit even when the beam was not present.
- b. The anomaly was independent on the sensor bias voltage.
- c. The list of noisy pixels changes as a function of time: the time constant is of the order of 100 s.
- d. Even after a few months without the beam the anomaly remains present.

As described in section 3.5, we irradiated a bare chip and an assembly (chip bump-bonded to a sensor) to a much higher dose with X-rays and we did not find anything causing the noise. In the next sections, we discuss in detail the status of component parts for the next run in 2016.

3.2 PCB carrier

A new version of the GTK carrier has been studied and five boards were produced in 2015. The main features of the new boards are as follows:

- a. Decoupling capacitors have been added close to the TDCpix.
- b. The power supply distributions on the board have been duplicated to limit the voltage drop on the VDD during the activation phase of the PLLs.
- c. The new version gives us the possibility to change the working mode of the TDCpix remotely.
- d. The analog output signal is now wire bonded to a pin connector. Thus, we can verify the DACs functionality during the tests.
- e. An optical reset input is now available on the board.
- f. The wire bonding scheme has been optimized by creating a step on the GND connection inside the GTK PCB carrier.

Furthermore, all the bonding pads were also brought nearer the TDCpix by $500\ \mu\text{m}$ in order to facilitate the work of wire bonding. The production of twenty GTK carriers is foreseen this year as soon as the first prototype (GTK_Carrier_5) has been completely tested and validated.

3.3 Bump-bonding

In the past year, we had the possibility to fully exploit, in the experimental environment, the assemblies of the first production, learning valuable information on the operational and performance side. Those assemblies have been done with untested read-out chips. In the same period we also set the baseline for the second production. In May 2015, we had a meeting at IZM on the overall performance of the first batch, where all the results on laboratory tests have been discussed. Hereafter, we report the details defined for the second batch: 9 full-module and 9 single-chip assemblies, with 100 microns thin TDCpix, using the last 9 wafers (n-in-p) from the FBK sensor production. The main differences with respect to the first production are:

- a. The usage of tested TDCpix chips.
- b. The deposition of BCB passivation layers on the sensor guard rings and on the read-out chip bands facing the guard rings, to protect the TDCpix from discharges coming from the sensor HV.

Following the meeting, IZM has designed and prepared all the necessary masks for the BCB layer depositions.

During the data-taking period of NA62, the operations to complete, commission, and operate the three GTK stations in the experiment required all the available team stamina. The work on the TDCpix wafer test system, initiated before the summer, restarted in the fall, with final commissioning and full-fledged testing of three wafers completed by December. The average yield of working chips is high, around 80%. The three tested TDCpix wafers have been delivered to IZM before Christmas.

The production of the new batch is currently ongoing in IZM: TDCpix wafers processing (thinning, glass carrier bonding, bumping, and laser carrier release) started at the beginning of January 2016, while the 9 sensor wafers from FBK production, already in house, have been processed with some advance. Everything is proceeding well. The schedule foresees the delivery of the full batch by the second week of April, with 4 full-module assemblies (put on a fast-track) delivered before Easter. These 4 first assemblies will be used to procure the next set of fully operational GTK stations for the upcoming run. Finally, in this last quarter, a new order has been placed to IZM, covering the next GTK assembly production for a total of 25 full-module units.

3.4 Micro-channel cooling

3.4.1 Procurement of the next cooling plates from Leti

The cooling plates delivered by CEA-Leti so far are not fully compliant to the NA62 specifications. Nevertheless, a full set of GTK detector modules has been installed in the beam and successfully taken physics data. The cooling plates from the first batch delivered in 2014 and 2015 exhibited the following main issues:

1. Too thick in the acceptance.
2. Did not hold the nominal operating pressure.
3. The thickness in the acceptance was not uniform.

During a meeting held at CERN on September 9th, 2015, with the NA62 technical coordination, the GTK Project management, representatives of the EP-DT group, and Leti management and process engineers, a roadmap to optimize the cooling plates was discussed. The first step consists in the procurement of 14 additional cooling plates split in two series with the following properties:

- 8 cooling plates of so-called Option A (see Figure 3 (left)) for operation at low pressure like the cooling plates currently installed in the experiment. They consist of a Silicon-On-Insulator 8 inch wafer in which the micro-channels are etched bonded to a bulk silicon wafer to close the micro-channels. The buried oxide layer, represented in blue in the cross-sections of Figure 3 allows for a perfectly flat surface after thinning in the acceptance by Deep Reactive Ion Etching (DRIE) of the silicon. The thickness of the cooling plate in the acceptance will be $270 \mu\text{m}$.
- 6 cooling plates of so-called Option B (see Figure 3 (right)) for operation at an intermediate pressure or even nominal pressure according to the thickness on top of the distribution manifold. This thickness has to be agreed with the wire-bonding team at CERN because the thicker it is, the harder it is for them to bond the TDCPix chips to the read-out board. In this option, not only the microchannels are etched in an SOI but they are also closed by bonding another SOI wafer on top. This allows two perfectly flat surfaces on both sides of the micro-channels in the acceptance to be obtained. The thickness of the cooling plate in the acceptance will be $210 \mu\text{m}$.

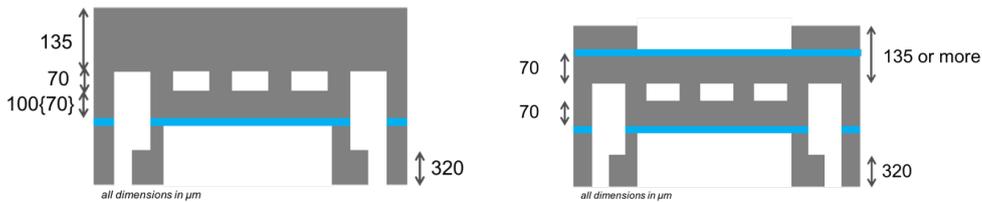


Figure 3: Cross-section of the cooling plates of (left) Option A and (right) Option B.

The cooling plates from the third batch could be meeting the specifications in terms of operating pressure and material budget in the acceptance by properly selecting the total thickness of silicon on top of the manifolds ($135 \mu\text{m}$ or more in the case of Option B) as well as the thickness of silicon between the micro-channels and the buried oxide layer on both sides

3.4.2 Hydraulic shock from the cooling plant to GTK_1

On June 18th, 2015, the cooling plant in the NA62 experimental area had been running stable at $-25 \text{ }^\circ\text{C}$ for about 20 hours. An imminent vacuum cut required to warm the

coolant circulating in the micro-channels. During the warm-up phase, the cooling plant was stopped due to a general DSS interlock. It was then restarted in normal mode. At this point the temperature had reached 0.7 °C. The cooling plant was stopped again and was restarted in warm mode. As it began to show a strange behavior, it was stopped yet another time to be restarted in normal mode. A few minutes later, it stopped again due to a DSS interlock which coincided with a loss of vacuum in the beam line. After investigation, a pressure surge during this manual operation was identified. This overpressure created a hydraulic shock which was not absorbed by the pressure relief valve (set at 4.5 bar) placed at the inlet of the cooling and led to the simultaneous breaking of the silicon in front of the inlet connectors. The problem was clearly identified to be coming from the cooling plant control system and corrective action was taken to improve the safety. Before failure, this cooling plate had successfully been commissioned in the experimental test-bench at CERN where it was tested at 6 bar following the standard procedure during the fall of 2014. It was then installed in GTK Station 3 of the NA62 beam line in November 2014 and was successfully operated for data taking until Christmas of the same year. The operation parameters were the following:

- 3 bar measured at the output of the cooling plant.
- 1.8 bar measured at the inlet of the micro-channels cooling plate.
- A pressure relief valve calibrated at 4.5 bar is installed just before the inlet of the micro-channels cooling plate.

It was then removed from the beam and installed in the laboratory of building 14 for further testing of the detector module. Following this test campaign, it was re-installed in the beam line.

3.4.3 Repairation of GTK_2

In June 2015, a bad thermal contact was identified between the micro-channels cooling plate and the tip of one of the TDCPix read-out chips of the detector module GTK_2 (see Figure 4). A dedicated tool was designed and manufactured (3D-printed at CERN) to improve the thermal contact between the chip and the cooling plate. The tool consists of a support positioned on the mechanical frame allowing for two sticks to apply a mechanical pressure symmetrically on the TDCPix tip (see Figure 5(left)) and the backside of the cooling plate (see Figure 5(right)).

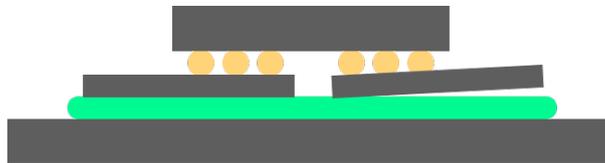


Figure 4: Schematic representation of the TDCpix detached from the cooling plate.

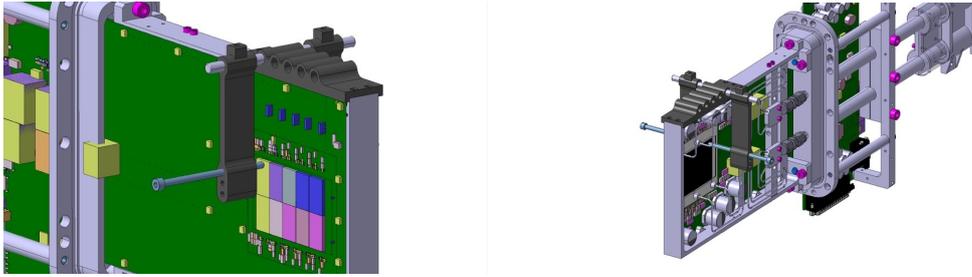


Figure 5: Dedicated tool to apply pressure on the TDCpix tip (left) and the cooling plate (right).

3.4.4 Assembly of modules GTK_4 and GTK_5

Two GTK modules were assembled in 2015: GTK_4 and GTK_5. Various steps of the assembly procedure were improved and validated with these two modules.

- Assembly of GTK_4 with a new support for the wire-bonding

A new support structure was introduced for the wire-bonding of the module GTK_4 (see Figure 6). This 3D-printed frame provided more rigidity to the cooling plate and guaranteed no displacement in the vertical axis when the bond tool attached the wires on the pads of the TDCPix. It was agreed with the wire-bonding team at CERN to implement this solution for all the later modules to be bonded.

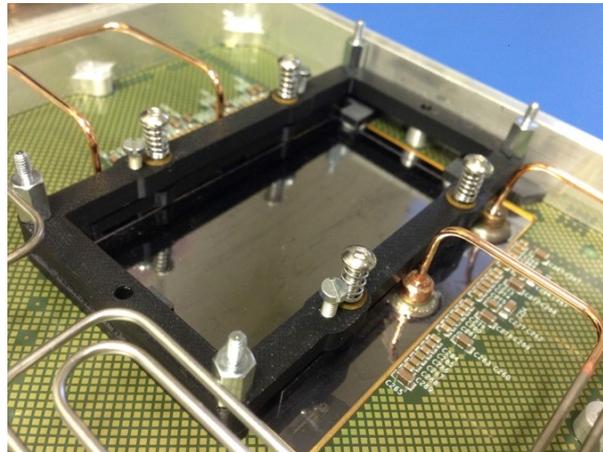


Figure 6: 3D-printed support structure positioned on the backside of the cooling plate during wire-bonding.

- Assembly of GTK_5 with a new brazing procedure and a new gluing jig.

The GTK_5 module was assembled with the first GTK_Carrier read-out board delivered by SOMACIS.

- New brazing procedure

The KOVAR connectors of the previous GTK modules were Cu-plated in order to be brazed to the cooling plates. Even though the brazed connectors of GTK_1 to GTK_4 were perfectly leak tight and robust, the brazing on other devices with Cu-plated KOVAR connectors failed to be leak tight. An extra layer of Au was plated immediately after the Cu-plating in order to avoid the oxidation of the Cu layer and to better the wettability (see Figure 7 (left)). This solution was tested and validated on a faulty cooling plate and selected for the brazing of the connectors of GTK_5 (see Figure 7 (right)).



Figure 7: Left: Au-plated KOVAR connectors with an underlying layer of Cu. Right: Brazed connectors on the cooling plate of GTK_5.

- New gluing jig

Following the issue with the thermal contact between the tip of one TDCPix and the cooling plate of GTK_2, the gluing procedure of the sensor assembly to the cooling plate was modified for GTK_5. To better control the pressure applied during the gluing, additional weights were added to the jig. In total, eleven weights were used during the curing of the 3M tape used as thermal interface between the backside of the TDCPix chips and the cooling plate (see Figure 8 (left)); one weight per each TDCPix and one weight for the sensor. The precise positioning of these weights with respect to the TDCPix chips and the sensors was obtained by guiding them in a 3D-printed insertion tool aligned to the sensor assembly and cooling plates with the help of alignment pins. After a few trials with dummy silicon sensor assemblies glued on glass dummy cooling plates, the procedure was validated and adopted for GTK_5 (see Figure 8 (right)). This method will be adopted for all the future GTK modules.

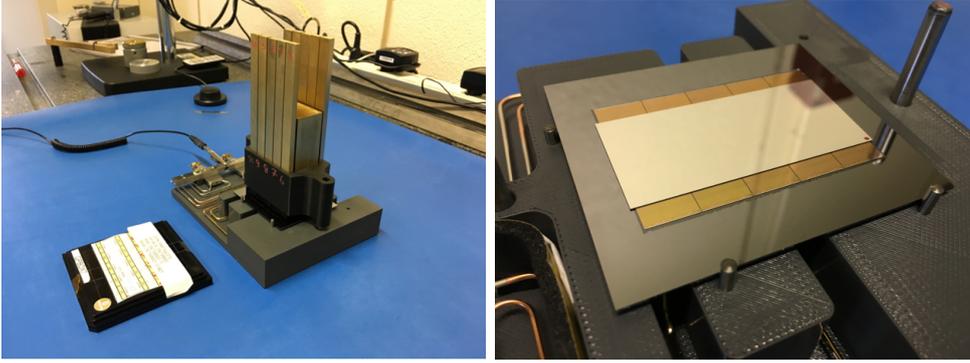


Figure 8: New gluing jig with additional weights (left) used for GTK_5 (right).

3.5 TDCpix test

- The behavior of the single TDCPix assembly has been tested with the high precision laser and found consistent with that expected from measurements performed on the prototype ASIC.
- Absolute calibration of the pixel thresholds of the assemblies installed on the GTK Carrier was performed with a 109-Cd source, using the 22 keV and 25 keV gamma peaks.
- The commissioning of the wafer prober was completed and 3 wafers were probed at the end of 2015 (180 chips in total). These wafers along with the acceptance wafer maps were sent to IZM for dicing and bump-bonding. The overall yield measured on these three wafers was 81%, consistent with that seen from measurements performed on the single chips and single chip assemblies, giving 146 known-good chips. This is largely sufficient for the second round of bump-bonding.
- Additionally, the wafer prober was used to probe one of the two remaining multi-chip assemblies from the original batch. This assembly was subsequently transferred to the integration team for the manufacturing of the GTK_Carrier_5 - the first detector using version 2 of the GTK Carrier Card. With this, we were able to verify that placement of the needles on the wire bonding pads is sufficiently accurate to leave enough space for the wire bonding and additionally, that the marks left by the probe needles would not be detrimental to the quality of the wire bonding. So far, no issues have been detected.
- Significant effort was invested in the preparation of the GTK Carrier Card production testing laboratory. This turned out to be more involved than initially anticipated as a result of recently tightened safety measures for new installations. The laboratory is situated in the basement of building 13 and will be ready for the installation of the equipment required for the production testing at the beginning of April. This room is dedicated to production testing and will not be accessible to any other experiment. A humidity controller storage cabinet has been foreseen in the EP-ESE wafer storage facility (also in the basement of building 13) for the storage of tested carrier cards awaiting use. In parallel, space was made available in the ESE/GTK laboratory (on the 3rd floor of building 14) to begin the integration

of the production testing equipment and this work is expected to be completed at about the same time as the laboratory becomes available. The installation of the cooling circuit will be undertaken by a team from EP-DT.

- At the end of the 2015 run, GTK_4 (installed in station 2) was extracted from the beam line and brought back to the ESE/GTK laboratory to enable testing of the spurious noise. A very detailed investigation to try to determine the cause of this noise is currently under way, although it has proven difficult to identify the cause of this noise.
- As part of this investigation, a single chip with no sensor and a single chip with a sensor were taken to the ESE-ME Total Ionising Dose X-Ray machine on the 5th floor of building 14. Here, the pixel array without sensor was irradiated to 6 Mrad and the end of column region to 900 krad. Although these doses significantly exceed that seen in the 2015 run, no undesirable behavior was observed.
- Subsequent to the irradiation, the time resolution of the ASIC was measured to be just under 90ps (RMS). This represents a small degradation relative to the 75ps (RMS) measured prior to the irradiation.
- The sensor on the single chip assembly was also irradiated to a dose equivalent to about 10^{13} particles/ cm² and the noise did not appear.
- The investigation, thus, remains open. However, the TDCPix chip has been qualified in terms of TID up to the level expected in a year of running.

3.6 "Off-detector" DAQ status

The off-detector readout system is based on VME cards, denominated "GTK-RO", which receive the data sent by the TDCpix ASICs, store them on temporary buffers pending the L0 trigger requests, and send packets of "trigger-matched" events to the sub-detector PC farm. The off-detector readout system comprises, also, one "GTK-TIMING_UNIT" card, in triple width VME format, providing the following set of high quality timing signals to the TDCpix ASICs:

- The DLL_CLK (320MHz) used by the TDC units built in the TDCpix ASICs.
- The DIG_CLK (320MHz) used by the data processing units inside the TDCpix ASICs.
- The COARSE_FRAME_COUNTER_RESET providing the "time origin" information to the ASICs.
- The TEST_PULSE defined with a resolution of down to 5 ps.

3.6.1 Update on the "Off-detector" DAQ

Each TDCpix ASIC is served by one GTK-RO module. The latter is actually made of two decked units: the mother board, which is a 6U VME card hosting the main functional blocks, and a daughter card featuring the interface to the TTC system and various timing functions required for the control of the connected ASIC.

The "Off-detector" readout system reached its full implementation in time for the start of the 2015 run with the installation of the 30 GTK-RO daughter cards delivered at the end of 2014 and the installation and commissioning of the GTK-TIMING_UNIT modules.

The firmware for the FPGAs on the GTK-RO motherboard has been upgraded, in preparation for the 2015 run, to move from the "calibration" mode (stand-alone mode used for the commissioning of the detectors) of data acquisition to the "trigger-matched" mode foreseen for the integration of the GTK DAQ system with the rest of the data acquisition systems.

Also, the "calibration" mode firmware has been upgraded to be compatible with the higher luminosity foreseen for the 2015 run with respect to the 2014 run and to increase the flexibility in the control of the serial link between the GTK-RO and TDCpix dedicated to the ASICs' configuration and monitoring.

The commissioning of the "trigger-matched" firmware, i.e. testing in the field what had been developed using ideal models of the ASICs' behavior and the experience acquired with real detectors in 2014, started in July and took longer than expected. Experienced problems had been first, with the reliability of the DDR2 DRAM memory controller instantiated in the FPGA firmware. An upgrade to a more recent version of the controller IP module solved these issues. Subsequently, the debugging of the "trigger-matched" firmware involved the search of the conditions that induced meta-stabilities in the different state machines controlling the GTK-RO operation and resulting in stalls of the flow of data toward the NA62 PC FARM.

The first trigger matched data were collected at the end of August. By the time of the high intensity part of the 2015 run, all installed GTK assemblies were read out in "trigger-matched" mode.

Concerning the hardware development, the activities being carried out in preparation for the 2016 run are toward improving the interface to the run control systems in order to make the initialization of on-detector and off-detector readout systems more automatic and repeatable.

3.6.2 Local PC farm performances

In the 2015 Run, the GTK-PC farm was composed of 3 sub-detector PCs. Each PC is meant to receive event fragments from the 10 GTK-RO boards of a station, check data consistency (for debug and monitoring purposes), assemble the fragments, and send the complete events to the NA62-PC farm. Performances have been studied in different input/output rate conditions.

No input limitation has been observed in run conditions, namely event fragments from the 10 GTK-RO boards are sustained by a GTK-PC at 100% of nominal intensity with $\sim 250\text{k}$ event/burst and at 50% nominal intensity with $\sim 900\text{k}$ event/burst.

On the other hand, limitations have been observed in the output capabilities at high trigger rates (w.r.t. average 2015 run conditions), i.e. 50% of the nominal intensity with $\sim 900\text{k}$ event/burst. This is explained by a software bottleneck in the final assembler of complete events. In order to overcome this output limitation, the GTK-PC will be read-out at L1 trigger level, lowering the number of events to assemble and send to PC farm.

3.6.3 Upgrade local PCfarm (3 PC to 6 PC and L0 to L1)

At the moment, each GTK-RO board is sending events to the GTK-PC eight at a time, for a total amount of data of ~ 1150 B/packet. Due to the use of ethernet, the max theoretical throughput sending the data to the farm with this packet size is ~ 960 Mbps, equivalent to ~ 0.83 Mevts/s. To reach the 1 Mevts/s we expect, we need a second 10 Gb link, both in input (from the GTK-RO to the GTK-PC) and in output (from the GTK-PC to the PCfarm).

3.6.4 Upgrading from 3 to 6 GTK-PC

This upgrade requires minimal changes on the GTK-PC side. The most relevant is to provide 3 more ports on the central router. There is a plan to add one 8-port module to the switch: 7 ports would be usable for detectors, and 3 would be assigned to the 3 new GTK-PC. The PC purchasing order was approved on the 2nd of February with 6 weeks foreseen for the delivery.

3.6.5 Upgrading from L0 to L1

The GTK-PC software has been modified to receive Multi Requests Packets (MRP) issued by the PC farm and containing multiple L1 trigger requests. The implementation of the GTK-PC software, inherited from the 2015 version, allows us to fetch (pre-indexed) event fragments using the L1 trigger information. This allows us to assemble and send only (L1) requested events overcoming the output bottleneck. This solution has been implemented and will be tested before the start of the 2016 run.

3.7 Results from 2015 data taking

As stated above, during the 2015 run the three stations were installed and equipped with detectors, and all the infra-structure such as vacuum, mechanics, cooling, and data acquisition were in place and operational. GTK data, synchronized to the NA62 level-0 trigger, were collected during the run and enabled the detector to be partially commissioned.

In 2015, the detectors were operated at a temperature of 0°C by circulating C_6F_{14} at 2 g/s, which allowed a safe pressure drop of around 3.2 bar across the cooling plates. Pixel-signal thresholds were adjusted and pixel-to-pixel thresholds equalized with a test-bench procedure performed before installation. The thresholds were set to 0.7 fC, corresponding to 30% of the most probable charge of a minimum ionizing particle (i.e. 2.4 fC). During the data taking, the bias voltages were scanned in the range between 200 and 300 V. All the GTK-RO cards were installed and optical-fiber communication to TDCpix at 320 MHz and 3.2 GHz was proven to be fully functional and reliable.

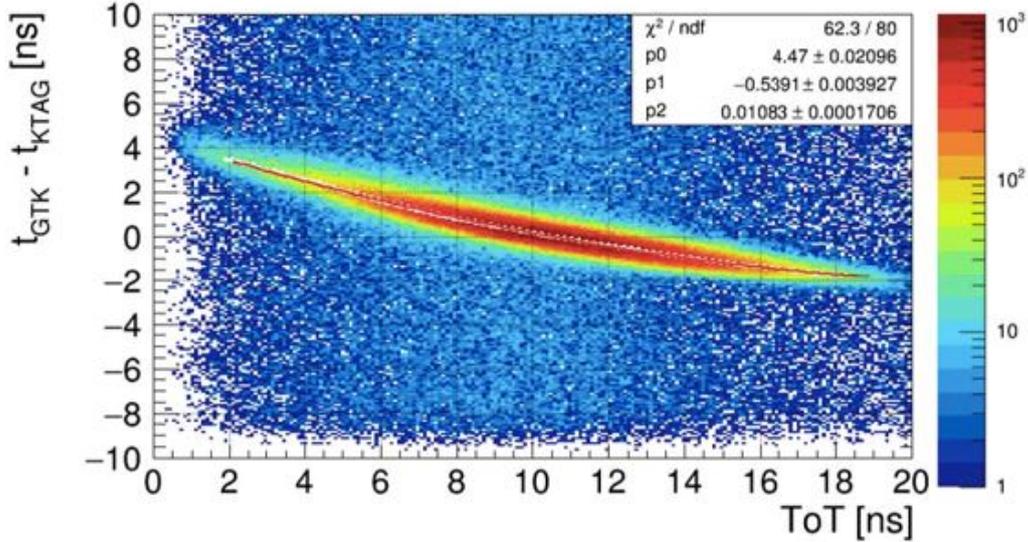


Figure 9: Time-walk as a function of time-over-threshold.

The operation of the GTK system in 2015 confirmed that the mechanics, the cooling, and the electronics were properly working, although, further calibration and data analysis are needed to exploit the full potential of this detector.

3.7.1 Time Resolution

The time of a hit in a given station is reconstructed by applying two corrections:

1. A time offset correction, which depends primarily on the station and the chip and, in second order, on the column and row of the pixel. This correction is straightforward and is derived by comparing the GTK hit time to any other detector providing precise timing, such as the CHOD or KTAG.
2. A time-walk correction to account for Landau fluctuations of the signal.

Figure 9 shows the time-walk as a function of the time-over-threshold (ToT). The corrected hit time is obtained by correcting the raw time with the fitted variation of its ToT.

The measured time resolution for the three GTK stations is shown in Figure 10 and includes a small contribution (80 ps) from the KTAG. These measurements were obtained using a bias voltage of 300 V for GTK1 and 2, and 216 V for GTK3. The time resolutions are approximately 50 ps larger than the results obtained with a prototype detector, when operated at similar bias voltages. We ascribe the worsening of the time resolution to the contribution of the systematic errors of the time offsets.

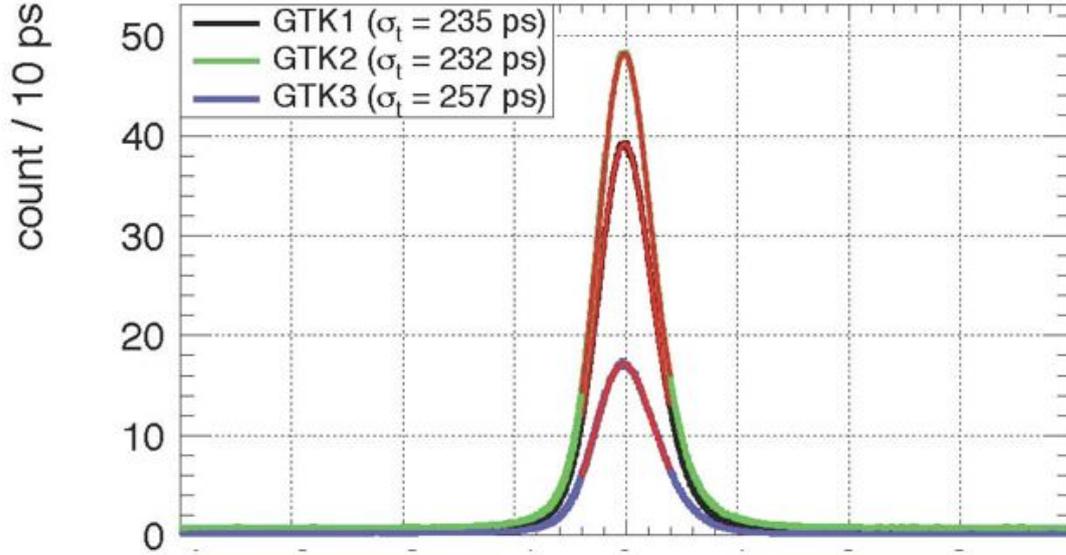


Figure 10: Time resolution for the GTK stations, when operated at a bias voltage of: GTK1 and 2 at 300 V, GTK3 at 216 V.

Momentum and directions of kaons were reconstructed and associated to the tracks measured in the Straw Spectrometer. A preliminary impact of the GTK on the physics analysis is demonstrated by extracting the squared missing mass $|P_K - P_\pi|^2$, for a selected sample of $K^+ \rightarrow \pi^+\pi^0$ decays (see Figure 11). The stations were aligned in the transverse plane using a sample of $K^+ \rightarrow \pi^+\pi^+\pi^-$ decays. The K^+ was reconstructed from the three pions using the Straw Spectrometer and its propagated position at each GTK station was compared to the GTK hit position. The error on the missing mass, $1.5 \times 10^{-3} \text{ GeV}^2/c^4$, was indeed in agreement with the expected values as from the GTK specification. This method gave transverse alignment values consistent with the $K^+ \rightarrow \pi^+\pi^0$ analysis.

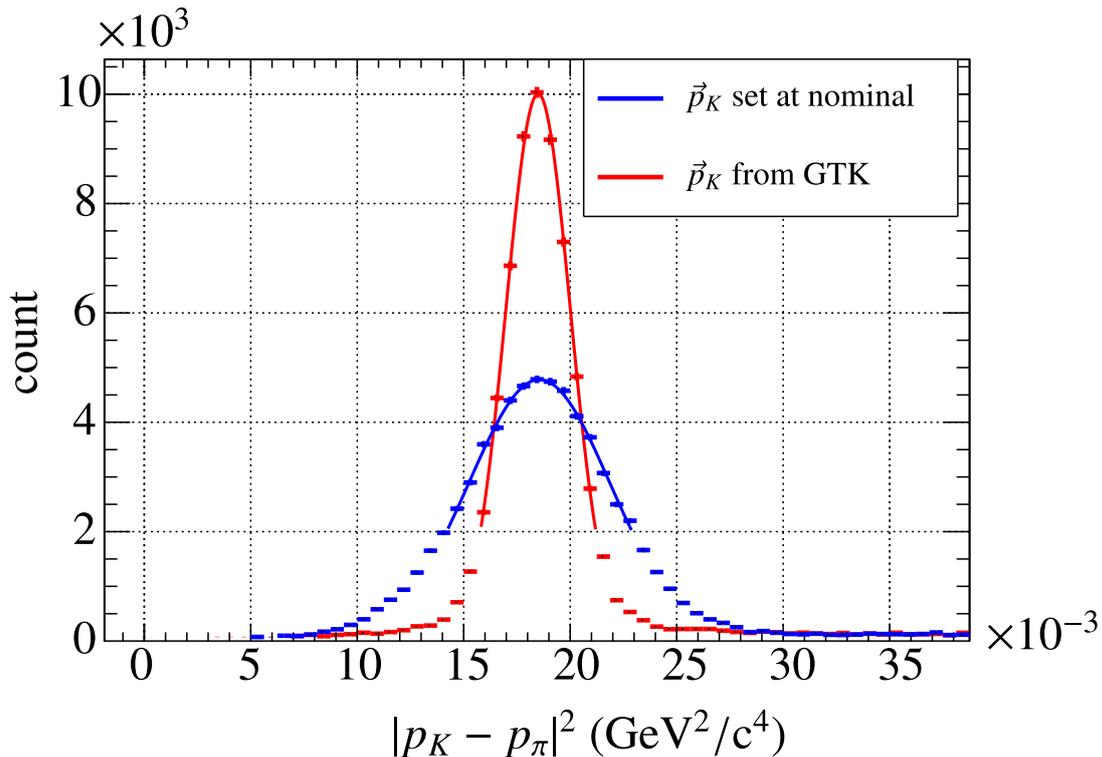


Figure 11: The squared missing mass distributions for $K^+ \rightarrow \pi^+\pi^0$ decays computed without (blue line) and with (green line) the GTK track reconstruction.

4 Large Angle Veto

The principal achievements of the photon-veto working group in 2015 were in the commissioning of the Large-Angle Veto(LAV) system. Particular progress on the LAV system was made in the following areas:

- Development of the level-zero trigger firmware.
- Implementation and optimization of the reconstruction code.
- Analysis of 2014 data and measurement of system performance.

The LAV system consists of 12 detector stations arranged at intervals of 6 to 10 m along the vacuum tank along its entire length. The first 11 stations are incorporated into the tank itself and are operated in vacuum; the 12th station is placed immediately downstream of the RICH and is operated in air. The diameter of the stations increases with distance from the target, as does the number of blocks in each, from 160 to 256, for a total of 2496 blocks. Each station consists of four or five rings of blocks, with the blocks staggered in azimuth in successive rings. The total depth of a five-layer station is 27 radiation lengths. This structure guarantees high efficiency, hermeticity, and uniformity of response.

The first LAV station, A1, was constructed in 2009 and served as a prototype. By August 2014, all twelve stations were completed, delivered to CERN and installed on the beam line. During the construction of the LAV detectors, more than 2500 lead-glass

blocks from the OPAL electromagnetic barrel calorimeter were processed (structurally reinforced, cleaned, fitted with new HV dividers, tested, and characterized).

The particles traversing the LAV detectors mainly consist of photons from kaon decays, as well as muons and pions in the beam halo. For each incoming particle, the veto detectors are expected to provide a time measurement with 1-ns resolution and an energy measurement of moderate precision (of order 10%). To maintain the detection efficiency as high as possible for muons and low energy photons, the system must be operated with thresholds of a few millivolts, i.e., well below the signal amplitude for minimum-ionizing particles (MIP). With an intrinsic time resolution of < 1 ns for the lead-glass blocks and a rise time of 5 ns for the Hamamatsu R2238 PMTs, the requirements on the precision of the time measurement are not difficult to satisfy. On the other hand, the amount of energy deposited in the LAV stations for photons from π^0 decays spans a very wide range, from about 100 MeV up to 30 GeV. Using the measured average photoelectron yield of 0.3 p.e./MeV and a nominal gain of 1×10^6 for the R2238 PMT, one expects a signal charge of 4.5 pC for a MIP, corresponding to a signal amplitude of 20 mV on a 50Ω load. At the upper end of the photon energy range, signals from 20 GeV showers can reach an amplitude of 10V. The readout chain for the LAV stations consists of two different types of boards, a dedicated front-end board (LAV-FEE) developed for the LAV detector, and a common digital readout board called TEL62, used by many of the NA62 detectors.

The LAV-FEE board, a custom 9U board designed by the LNF Servizio di Elettronica, splits the analog signal from the PMT into two copies and converts each into a logical LVDS signal, using two comparators with independently adjustable thresholds. The duration of the LVDS pulse is equal to the time during which the analog signal is above the programmed threshold. The basic idea is to exploit the time-over-threshold (ToT) technique to measure the signal charge over a broad interval. The LVDS signals are then sent to the TEL62 readout board, in which a custom-designed TDC mezzanine digitizes each signal into leading and trailing times. The FPGAs on board the TEL62 are used to correct raw hit times for slewing and to produce a level-zero (L0) trigger primitive, which is sent to the L0 trigger processor using a dedicated Gigabit Ethernet interface. The system is designed to sustain hit rates of up to 100 kHz per channel and to be able to transmit a data volume to the L1 PC farm of up to 2.4 Gbit/s for each station.

During the 2014-2015 runs, there was significant progress on the implementation of the L0 trigger for the LAV system. The NA62 L0 trigger receives inputs, or "primitives", from each detector system. These are generated by the FPGAs on the TEL62 boards, which examine buffered data during acquisition, and then forwarded to the central L0 trigger processor, which decides whether to initiate readout. The LAV12 L0 firmware reconstructs physical hits by searching for matching leading edges on the low- and high-threshold channels of each block within a programmable coincidence window, nominally 6 ns; if a coincidence is found, a slewing correction is applied to obtain the hit time. The hits are sorted and hits within a programmable window up to 25 ns in length are merged. The primitive, if generated, contains the timestamp of the acquisition window to read out; the time resolution is 100 ps.

Primitives are continuously generated with a rate of up to 10 MHz, and during the 2015 run were used as a zero-bias monitor of the activity on LAV12. As illustrated in Figure 12, the spill structure is evident in the LAV12 primitive rate. During the 2014 run, the L0 primitives provided the first evidence for the 75-Hz microstructure in the spill seen in the zoom in the bottom panel of the figure.

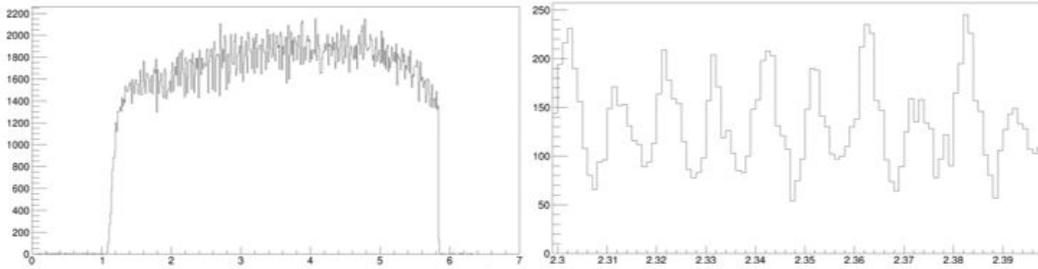


Figure 12: Continuous monitoring of LAV12 hits from L0 trigger primitives, in bins of 0.5 s (left), and zoom in bins of 0.02 s (right), showing evidence of 75-Hz spill structure.

Much progress has been made on the LAV reconstruction code and on the data analysis. The performance of the LAV system has been studied both with data collected under standard conditions, i.e., with a trigger based on the identification of K^+ decay products, and during dedicated muon runs, with the beam dump closed and the muon-sweeping magnets turned off, so that the experiment is traversed only by high-energy muons that penetrate the upstream shielding.

Muon runs were used to establish the threshold settings and to study the efficiency for the reconstruction of hits left by MIPs. Penetrating muon "tracks" in the LAV system were identified by the correlation between hits on blocks at the same azimuthal angle in different stations, as illustrated in Figure 13; for certain configurations, it is additionally possible to require hits on the blocks immediately upstream and downstream of a given block in order to better determine its efficiency. We find that the MIP detection efficiency saturates for values of the low threshold below 6 mV. This led us to adopt a value of 5 mV for the low-threshold working point.

Runs in standard conditions are used to measure the time offsets for each channel with respect to the signals from the detectors that provide the experiment's event-time reference (CHOD and KTAG). Hit reconstruction is then performed and slewing corrections are applied. A hit may be reconstructed from up to four time measurements, corresponding to the leading- and trailing-edge times on each of the high and low thresholds. The algorithm used to correct for slewing depends on how many and which of the edges are used to reconstruct the hit. For example, if both leading edges are present, the slewing correction is based on the difference between the high- and low-threshold crossing times; if only the low threshold is crossed, the slewing correction is based on a fit to the measured distribution of leading-edge time vs. time-over-threshold. After the application of slewing corrections, time resolutions at the level of 1 ns or better are obtained for all LAV stations, as shown in Figure 14. We note that these results are obtained with samples including all hit edge configurations: not only complete hits built from leading/trailing edge pairs for both thresholds, but also hits crossing only the low threshold.

We have attempted to study the detection efficiency for photons using a clean sample of $K \rightarrow \pi\pi^0$ decays in which the π track is reconstructed, one of the two photons from the π^0 is detected as an LKr calorimeter cluster, and the second photon is expected to be found in one of the LAV detectors on the basis of kinematic closure. At present, the resolution on the extrapolated direction of the photon is not sufficient to allow the efficiency to be determined for individual LAV stations; this may be possible in the future by means of a complete kinematic fit making use of all available information on the K^+

trajectory from the Gigatracker. As a first attempt, we have focused on estimating the global efficiency for the entire LAV system. We thus consider events to be successfully matched if they contain at least one LAV block fired within 5 ns of the $K \rightarrow \pi\pi^0$ event time from the reference detector. MC studies demonstrate that the photon detection inefficiency as determined by this method is dominated by geometrical inefficiencies and upstream photon conversions; the intrinsic inefficiency arising from the LAV detectors itself is less than one-sixth of the observed inefficiency. Relying on the MC estimate for the contribution from the former effects, we find the intrinsic inefficiency to be a few 10^{-3} , with about 5% of the detected photons observed as a signal on an isolated block crossing only the low threshold. These results are preliminary; as noted above, further refinements to the method will be implemented. We note, however, that obtaining an accurate tag for the determination of the single-photon detection efficiency at the level of 10^{-3} or lower is difficult; determination of the experiment's π^0 overall rejection power is both easier and more relevant for the measurement of $BR(K^+ \rightarrow \pi^+\nu\bar{\nu})$.

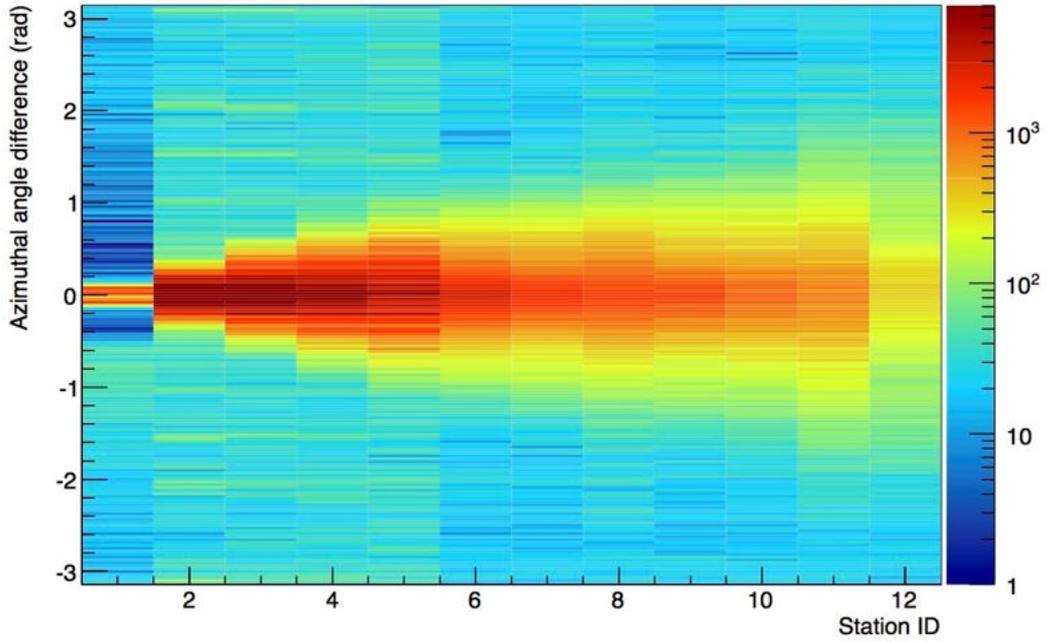


Figure 13: Distribution of difference in azimuth, $\Delta\phi$, between clusters on different stations vs. the number of stations between the most upstream and downstream clusters.

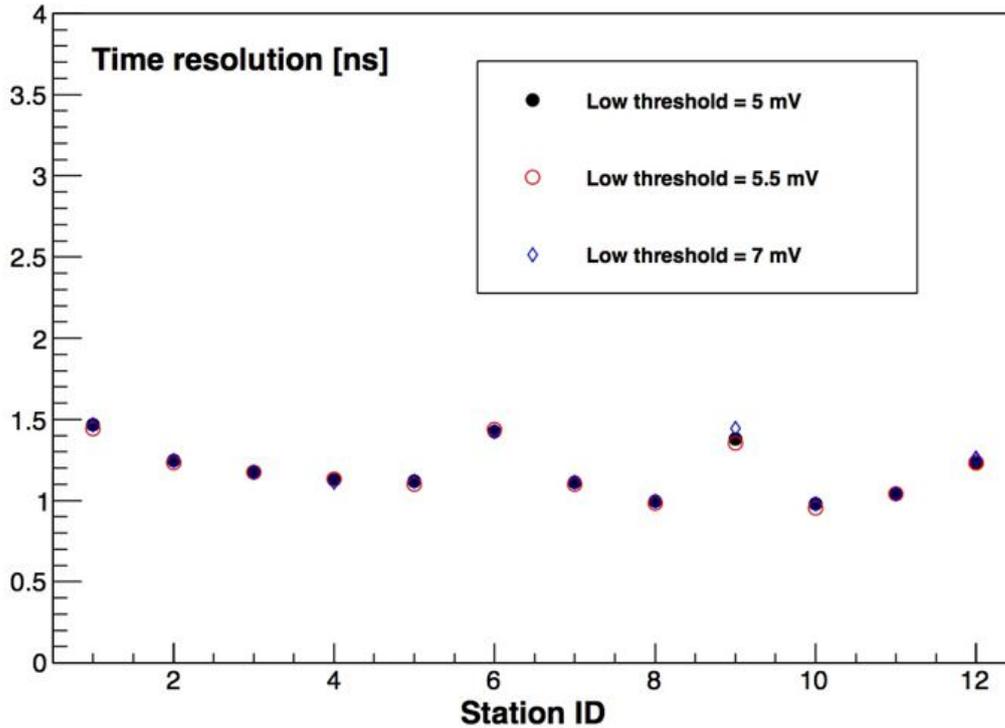


Figure 14: Difference between LAV hit time and event time from KTAG (ns) for different thresholds, for the twelve stations.

5 Straws

The four straw chambers were taking data during the run in 2015. All straws except one were operational and they participated in the data taking on a regular basis. One straw was sealed and disconnected after a small leak was found during the 2014 run (see Figure 15). Straw readout and DAQ was tested, commissioned, and successfully used during the data taking campaign. In 2015, the straw detector was capable of taking data efficiently at 50% of the nominal intensity. The firmware and software has, in 2015, undergone a number of improvements in real data taking conditions. The beam intensity variations reached up to four times nominal intensity and handling the unexpected increase in data flow related to beam intensity variations requires further upgrade of SRB firmware and DAQ software. The aim is to have the basic high-rate data firmware and DAQ ready for mid-April and further updates on on-line VME monitoring in the course of May-June 2016. The advance in functionality and firmware improvements should allow the straw tracker to play an important role in trigger level 1. The analysis of the 2015 data is ongoing and the track time resolution is shown in Figure 16.

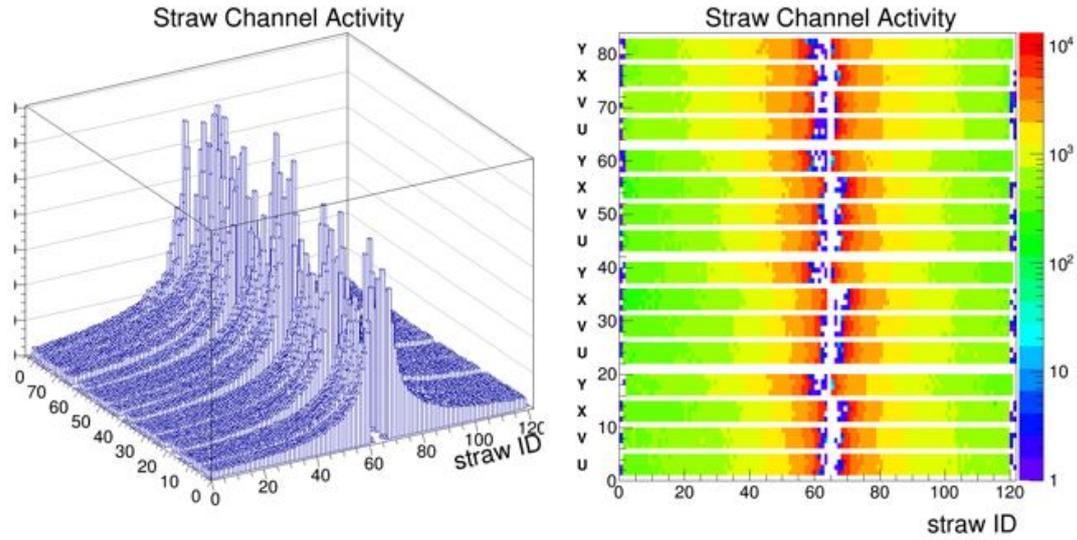


Figure 15: Online plot showing the hit activity straw by straw in all four chambers. One straw (blue point in the last layer in chamber 4) was disconnected and sealed in 2014 due to a small leak. The histogram shows triggered data as it is sent to the PC farm.

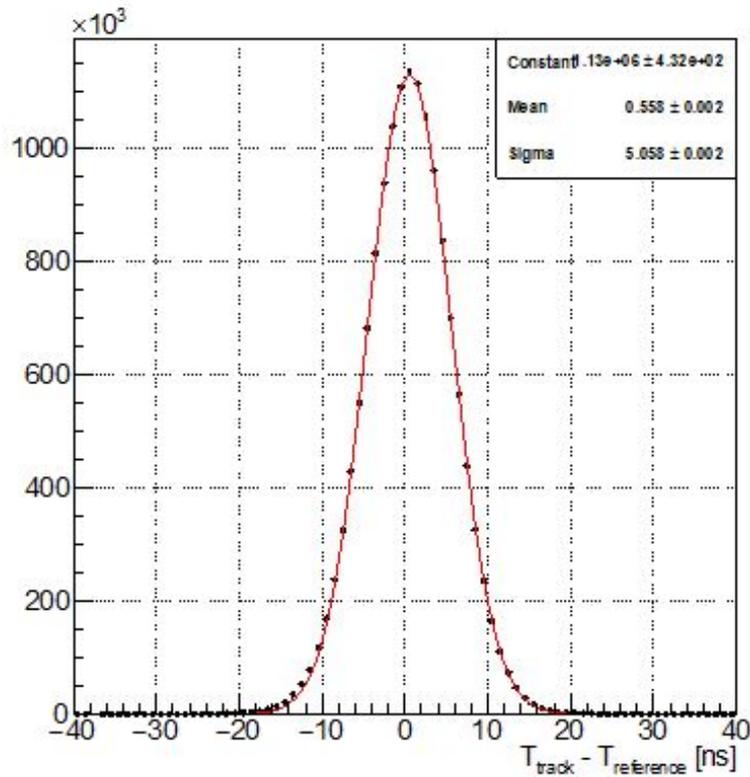


Figure 16: Time resolution for single-track event. The t_{track} is obtained from tracks having at least 27 straw hits per track and at least four hits in every chamber.

6 RICH

The RICH detector is needed to suppress μ^+ contamination in the π^\pm sample by a factor of at least 100 between 15 and 35 GeV/c momentum, to measure the pion crossing time with a resolution of about 100 ps and to produce the L0 trigger for a charged track. The detector consists of a 17 m long tank (vessel), filled with Neon gas at atmospheric pressure. Cherenkov light is reflected by a mosaic of 20 spherical mirrors with 17 m focal length, placed at the downstream end, and collected by 1952 photomultipliers (PMTs) placed at the upstream end.



Figure 17: Left: the RICH vessel installed in the NA62 cavern. Right: the beam pipe seen from inside the RICH vessel; the entrance window and the location of the PMTs disks can clearly be seen.

The RICH detector installation was completed in the summer of 2014 and the detector was used for the first time during the pilot run at the end of 2014. The RICH was then operated and run-in during the NA62 Commissioning Run in 2015.

In order to illustrate the performance of the RICH detector, samples of charged pions, muons, and electrons were selected using calorimetric and spectrometer information. In Figure 18 (left), the number of hits per Cherenkov ring as a function of particle momentum (measured by the spectrometer) is shown for electrons, muons, and charged pions. In Figure 18 (right), the Cherenkov ring radius as a function of momentum is shown without any selection on the type of particle: electrons, muons, charged pions, and scattered charged kaons can be clearly seen.

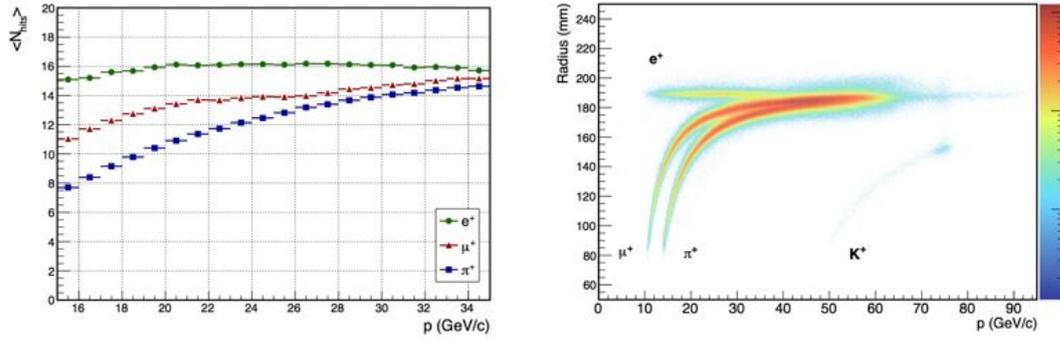


Figure 18: Left: number of hits per Cherenkov ring as a function of particle momentum; samples of electrons, muons and charged pions are selected using spectrometer and calorimetric information. The number of hits in the electron case is not exactly constant due to the limited acceptance of the photo-multipliers. Right: Cherenkov ring radius as a function of particle momentum; electrons, muons and charged pions can be clearly seen; charged kaons from the scattered beam can also be seen. Particles with momentum higher than 75 GeV/c are due to background muons from the experiment target.

The squared particle mass can be reconstructed using the velocity estimated by the Cherenkov angle measured by the RICH and the momentum measured by the spectrometer. In Figure 19, the squared reconstructed particle mass is shown (left), selecting the fiducial momentum region between 15 and 35 GeV/c for samples of electrons, muons, and charged pions identified using spectrometer and calorimetric information. Cutting on the reconstructed mass, charged pions can be selected and muons can be rejected: the right picture indicates the efficiency for pion selection as a function of the muon suppression for different choices of the cut on the reconstructed mass; for a 86% pion efficiency a 1.3% muon survival probability is quoted. It must be noted that in 2014 and 2015 the RICH mirrors alignment was not optimal and the need for better pion-muon separation was the main reason for the detector maintenance carried out in the 2015-2016 winter shutdown.

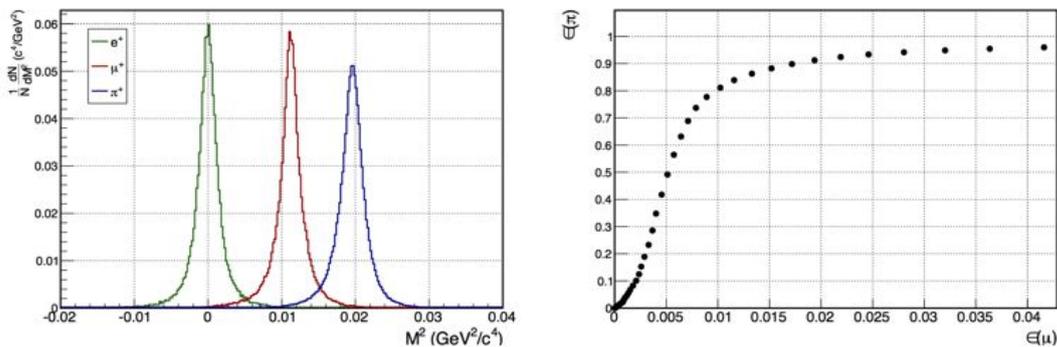


Figure 19: Left: squared particle mass reconstructed using the velocity estimated by the Cherenkov angle measured by the RICH and the momentum measured by the spectrometer for electrons, muons and charged pions selected using spectrometer and calorimetric information, with particle momentum between 15 and 35 GeV/c. Right: charged pion efficiency versus muon efficiency for different choices of the reconstructed mass cut, with particle momentum between 15 and 35 GeV/c.

The mirror mosaic is composed of 18 spherical mirrors of hexagonal shape (350 mm side) and 2 mirrors of semi-hexagonal shape located close to the beam pipe.

Each hexagonal mirror has a cylindrical hole, 12 mm wide, in its rear (i.e. not reflecting) surface, close to the geometric center. An aluminum dowel, inserted into the hole and connected to a support panel, sustains the mirror. Two thin aluminum ribbons, attached at the mirror rear surface at about 250 mm from the hole, keep the mirror in equilibrium and allow its orientation. The semi-hexagonal mirrors have two holes and a single ribbon. The mirror alignment system is based on piezo-motor actuators connected to the mirrors by thin aluminum ribbons: all the piezo-motors (two for each mirror), with their encoders, were mounted on the support panel in the summer of 2014, but some of them were not working properly after installation, preventing a fine alignment of the system. During the 2014 and 2015 run, some mirrors were re-aligned with a remote control using information from the data collected, but this was not possible for all the mirrors. Furthermore, the two semi-hexagonal mirrors turned out to be too close to the beam pipe and were almost out of acceptance. All these reasons pointed to a major intervention of maintenance after the end of the 2015 run.

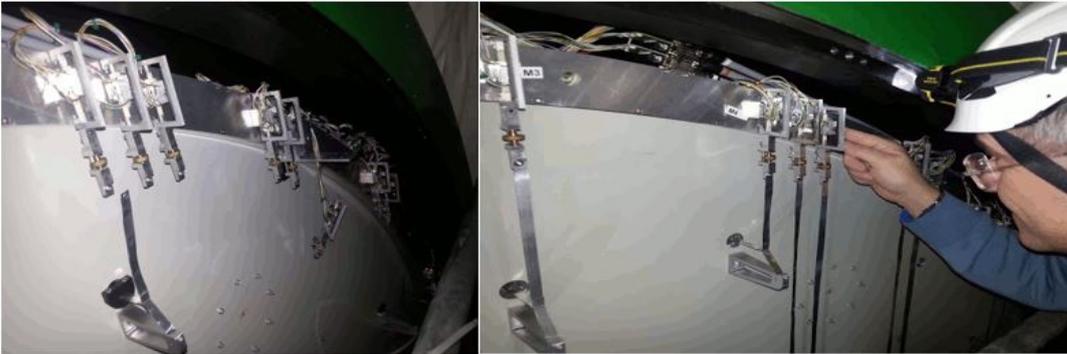


Figure 20: Left: downstream view of of the mirror support panel; the piezo-motor actuators are visible on the top, the aluminum ribbons connecting the mirrors have been disconnected after safety clamps have been installed to keep the mirror in position (visible at the bottom of the figure). Right: some piezo-motors have been replaced with new ones, connected to the aluminum ribbons and tested to properly work.

First, the RICH vessel was emptied of Neon in November 2015 and prepared for access. In December, the downstream endcap was opened to allow access to the rear part of the mirror support panel where the piezo-motor actuators are located. All 38 motors were checked and tested. About a dozen of the motors were replaced with new ones; the entire system was tested again and proved to work properly. In January 2016, the two semi-hexagonal mirrors were replaced with new ones with a slightly different geometry to better accommodate the passage of the beam pipe. In the following weeks, the entire mirror system was aligned with a laser system, confirming the possibility to move each mirror in the requested way. The final alignment will be done using data collected by the experiment as soon it restarts in spring 2016.



Figure 21: Left: Front view of the mirror system after the two central semi-hexagonal mirrors have been removed. Right: Installation of one of the two new semi-hexagonal mirrors.



Figure 22: Left: Re-installation of the beam pipe through the mirror system. Right: Detail of the beam and the central mirrors.

7 LKr

7.1 CREAM readout

The CREAM readout was completed and operational for the 2014 run. For 2015 a series of improvements in the configuration procedures have been prepared. The initialization software was carefully tuned and completely integrated with the run control. It was also found that a power cycle of a CREAM crate, or the firmware reload, modify the setting of the baseline level (defined with a proper setting of dedicated DACs). A specific baseline restore procedure was set up, iterating the setting of the DACs, the reading of the pedestals with the computation of a new DAC offset to bring the pedestal to the right value, until the pedestals of all the channels are within a fixed minimal value from the desired baseline. The procedure takes about twenty minutes and in normal conditions should be performed only when a major problem happens.

A series of monitoring plots was added to the experiment Online Monitoring system, including x-y hit map, distributions of the pedestals and of the sample of the maximum of the pulse.

As an additional network monitor tool, all the CREAM switches are configured to have a port which mirrors one of the CREAM outputs. All these mirrored ports are connected to a single switch to which a PC is connected. From the PC one could monitor (using network tools like Wireshark) the output of anyone of the mirrored CREAM ports to spot problems like sending data to a wrong IP address.

7.2 CREAM firmware modifications

The CREAM firmware produced by CAEN and thoroughly tested during the commissioning has been upgraded with a series of features found important after the debugging phase:

- Improvement of the zero suppression threshold with a cut on the difference between the minimum and maximum sample, like it was done for the 2012 run with the old CPD readout.
- Addition of a programmable detector ID in order to use the CREAMs for MUV1 and MUV2 readouts also.
- Increase of the depths of the L1 input FIFOs to avoid situations where the data was not sent to the correct PC.
- Implementation of selectable ways of forming the trigger sums, in order to find the best one for the optimal synchronization of the CREAMs with the TELDES boards in the LKr L0 system.

Later, during the run, it was realized that the possibility of having event packing in output from the CREAM should have been very useful to optimize the network traffic. This possibility has been studied in detail; a specification paper has been sent to CAEN to check the feasibility and to get an economical offer. This implementation was eventually done at the beginning of 2016.

7.3 Operation during the 2015 run

The intensive operation of the readout during the 2015 run has shown a few weak points that have been cured for future runs.

Several times, with a frequency proportional to the beam intensity, CREAM modules became stuck, likely because of radiation induced SEU. The module could be recovered reloading the FPGA firmware and the data acquisition could be restarted, initializing the faulty module. Another effect of radiation induced SEUs was the regular hangup of one of the network switches used to collect the data from the CREAM Ethernet outputs to be sent to the PC farm. In this case, an entire crate was not sending data and an intervention into the area was required to cycle the power of the faulty switch. In order to improve the handling of such events, a plan to install Ethernet controlled power distribution units (PDU) to power the switches has been prepared and it will be implemented for the 2016 run.

A couple of times, an unannounced cut of the cooling water allowed the temperature in the CREAM racks to rise, which eventually forced the cut of their power distribution. These events were not announced by the DCS, so a plan to improve and consolidate the monitor of the rack cooling has been prepared: readings of the cooling water temperatures and of the status of the pumps will be added to the DCS with a proper set of alarms. In addition, the alarm thresholds for the crate temperature will be tuned; additional existing monitor equipment provided with the Ethernet controlled PDUs will be integrated in the DCS.

7.4 Calibration software and operation

Before and during the 2015 run, calibration data were regularly taken to verify the stability of pedestals and to compute pedestals, slopes and the constants of the digital filter, used to compute energy and time of the calorimeter cells.

The stability of the pedestals is shown in Figure 23. At various times, power cut problems or iterated firmware reloads (due to SEUs) destroyed the baseline equalization giving rise to the large errors. The pedestal, on average, moved by a maximum of 3 counts.

A typical linearity plot is shown in Figure 24, with a very good behavior, translating in a non-linearity of better than 1 per mille.

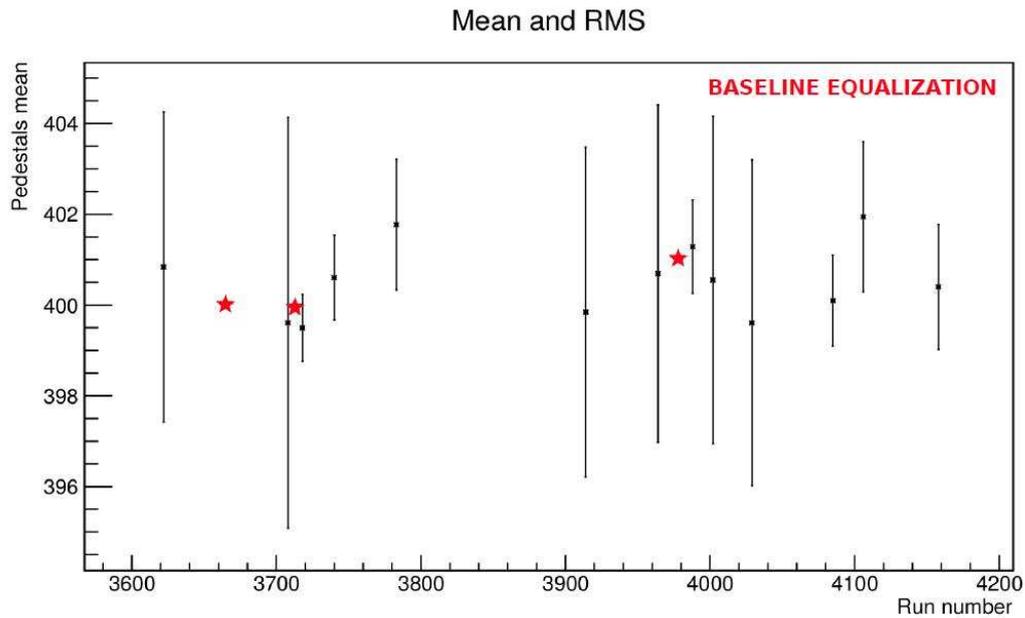


Figure 23: Average value of the pedestals, as a function of run number

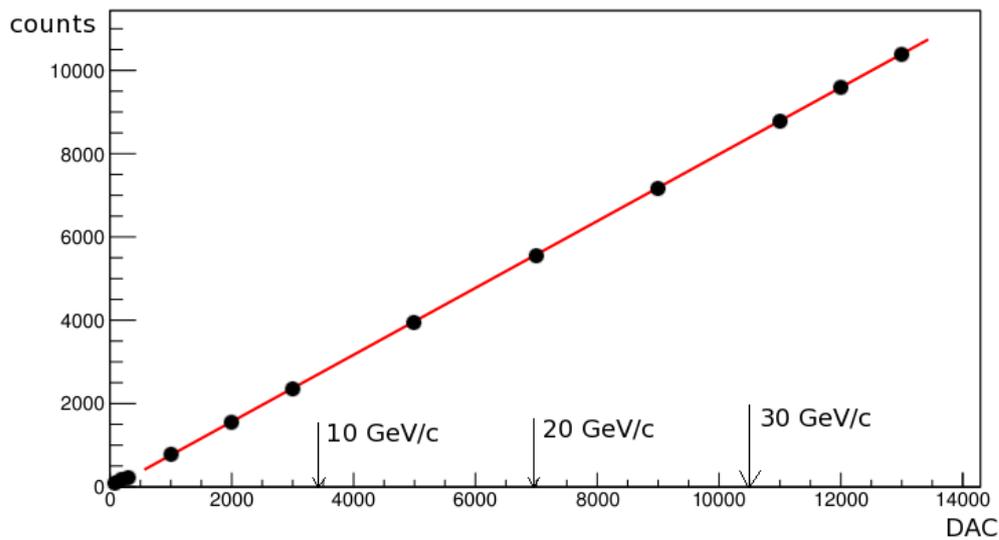


Figure 24: Linearity plot as a function of calibration voltage

For the following runs, there is a plan to integrate the calibration procedure in the standard runs, constantly acquiring pedestal and calibration data together with physics data, to have a tool for better monitoring of the parameter variation.

7.5 Cal-L0 trigger

The liquid krypton electromagnetic calorimeter trigger processor, which was originally designed to work on the LKr data, was upgraded in the last year to process data from all the experiment's electromagnetic and hadronic calorimeters: LKr, IRC, SAC, MUV1 and MUV2 read out through the CREAM system.

In this new configuration, the Calorimeter L0 (Cal-L0) trigger processor is used to select events with a π^+ in the final state ($E_{hadronic}^{pion} \gg E_{hadronic}^{muon}$) and to veto events ($K^+ \rightarrow \pi^+\pi^0$) which have significant electromagnetic energy in the LKr, IRC or SAC detectors.

The Cal-L0 system continuously receives low granularity trigger sums corresponding to 16 calorimeter cells provided by the CREAM system (16 bits @ 40 MHz) and identifies clusters from the electromagnetic (LKr, IRC and SAC) and hadronic (MUV1 and MUV2) calorimeters. The Cal-L0 prepares time-ordered lists of reconstructed clusters together with the arrival times, positions, and energy measurements of each cluster. Electromagnetic and hadronic cluster search is executed in parallel for each of the five detectors to provide trigger decisions based on complex energy and cluster multiplicity combinations.

As such, the system also provides a coarse-grained readout of the calorimeters that can be used in L1/L2 software trigger levels. Two different kinds of data can be read out from the Cal-L0 processor upon reception of a Level 0 trigger: raw data and reconstructed cluster data.

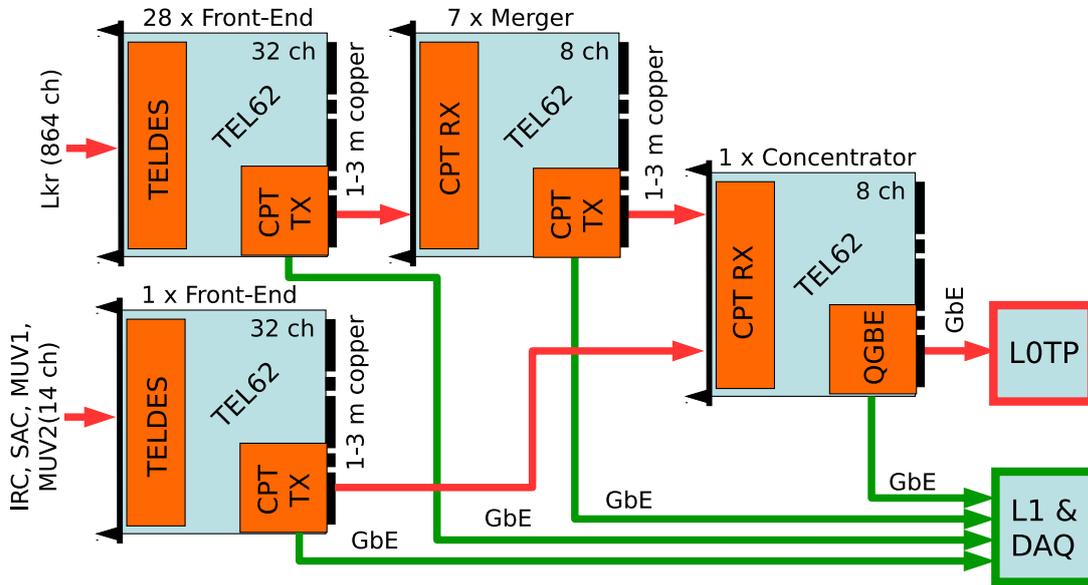


Figure 25: Schematic diagram of the Cal-L0 trigger. The system is composed of 37 TEL62 boards with dedicated daughter boards. Three different layers of boards are used: Front-End, Merger (only for LKr calorimeter) and Concentrator. The processor is installed in three 9U racks in the electronics barrack above the LKr calorimeter.

The Cal-L0 trigger processor is a parallel system, composed of Front-End, Merger, and Concentrator boards, all based on the TEL62 cards complemented with custom mezzanine boards. The basic functions of the three stages are:

- The Front-End boards receive trigger sums from the CREAM and perform peak searches in space and compute time, position and energy for each detected peak.
- The Merger boards (only used for the LKr) receive trigger data from the Front-End boards and merge peaks from different Front-End boards into single clusters.
- The concentrator boards receive reconstructed clusters from the five NA62 calorimeters, perform cluster counting, calculate sums for electromagnetic and hadronic

energy, and generate trigger primitives for the L0 Trigger Processor.

All together, the system is composed of 37 TEL62 boards, 185 mezzanine cards and 221 high-performance FPGAs. The custom made mezzanines facilitate the data transfer, i.e.,

- The TELDES receives digital data (16 bit @ 40 MHz over 15 meters long ethernet cables) from the Calorimeter REAdout Modules (CREAM).
- The Cal-L0-TX and Cal-L0-RX mezzanines provide fast transfer for trigger (48 bit @ 80 MHz over 2 meters long individually shielded twisted pairs cables) and readout data (over standard Gbit Ethernet).

The most relevant parameters of the installed system are reported in Table 2

Table 2: The Calorimeter L0 Trigger receives digitized data from five NA62 calorimeters: LKr, IRC and SAC electromagnetic calorimeter and MUV1 and MUV2 hadronic calorimeters.

Detector	Physical channels	Trigger channels
LKr	13,248	864
IRC	4	1
SAC	4	1
MUV1	176	6 (vertical) + 6 (horizontal)
MUV2	88	3 (vertical) + 3 (horizontal)
Total	13,520	884

7.5.1 Installation and commissioning

All the Cal-L0 hardware was installed before the beginning of 2015 run.

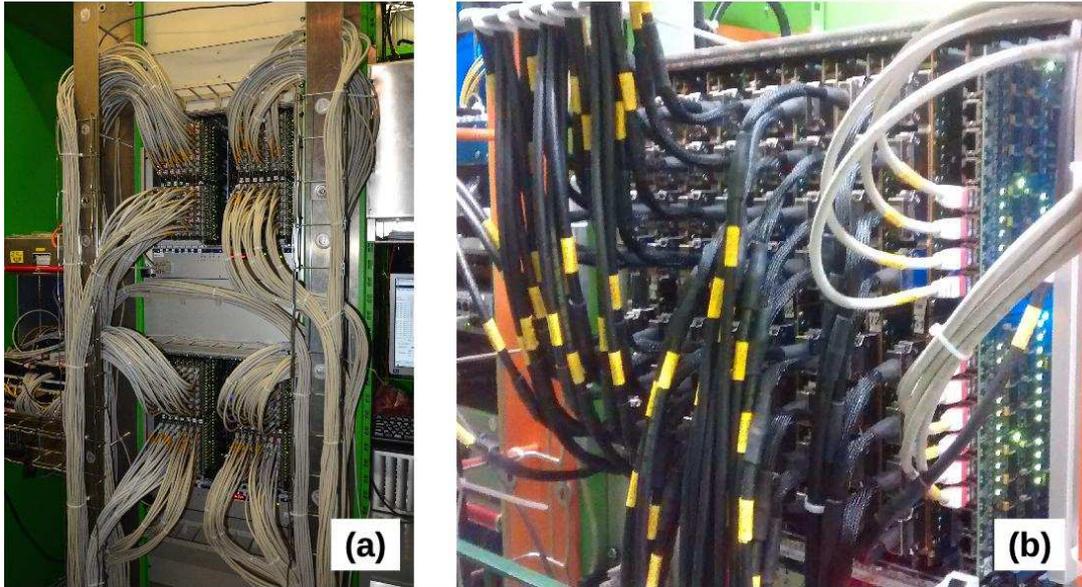


Figure 26: (a) Front-End crates receiving trigger sum links from the CREAM system. (b) Cabling between Front-End and Merger crates.

Firmware and software were continuously improved during the run, allowing a trigger on energy sums in the calorimeter at different thresholds and maintain a stable Gigabit Ethernet communication with the L0 Trigger Processor at 100% beam intensity.

After the end of the run, the Run Control software was completed and an automatic procedure was developed to detect and eventually correct Single Events Upset in the system.

Trigger clustering algorithm and readout firmware are currently being implemented.

7.5.2 Operation experience

Valuable experience with the operation of the Cal-L0 Trigger were obtained during the first high intensity physics run in 2015. The trigger processor functioned as expected with up to 100% beam intensity, i.e. no data transmission errors, no internal buffer overflows, or other hardware blocking conditions were found during the run.

No hardware replacement was needed during the run.

Single Events Upsets in the system were found to occur at the level of 2 SEUs per hour at 100% beam intensity, a rate which can be easily managed with the implemented SEU automatic recovery procedure.

During the 100% intensity runs the Cal-L0 trigger was used to select events with $E_{hadronic}^{pion} \gg E_{hadronic}^{\muon}$ and to veto events with significant electromagnetic energy in the LKr, IRC or SAC detectors.

Some preliminary measurements were performed and in particular a good agreement between the off-line reconstructed energy and the trigger energy threshold was found both for the Lkr (see Figure 27.a) and the MUV1 (see Figure 27.b) calorimeters.

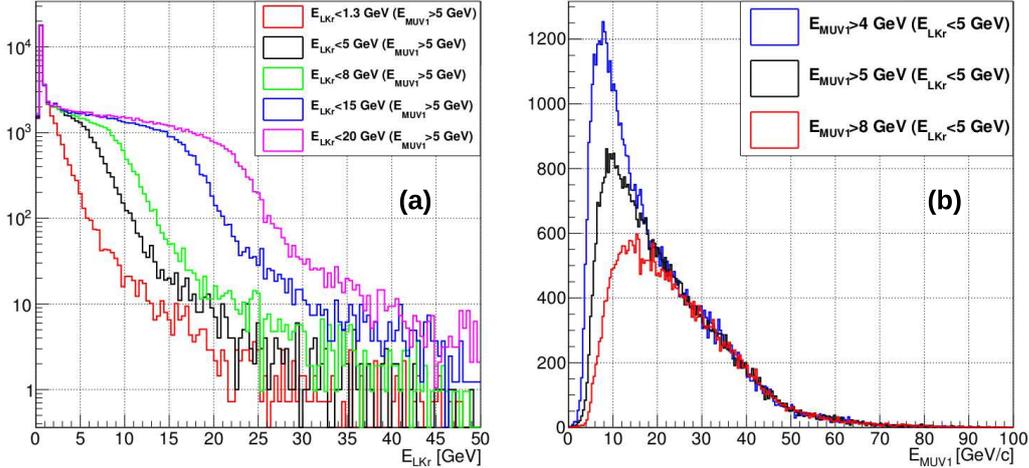


Figure 27: (a) Offline reconstructed energy in the LKr calorimeter for different LKr energy trigger thresholds. (b) Offline reconstructed energy in the MUV1 calorimeter for different MUV1 energy trigger thresholds.

8 CHOD

The NA62 experimental apparatus includes a fast scintillator system, called the charged particle hodoscopes, yielding signals when crossed by charged particles. The hodoscopes provide information on track impact point positions and precise measurements of the track crossing time. They produce fast signals for the L0 trigger, and allow background suppression when used in anticoincidence. The system consists of the NA48-CHOD detector from the former kaon experiment NA48 and a new CHOD detector optimized for the high beam intensity conditions of NA62.

The NA48-CHOD exploits a high granularity design based on the coincidence of signals in a plane of vertical and a plane of horizontal scintillator slabs. Refurbished in 2012, it has been operated successfully since the 2014 run. It consists of two planes of 64 vertical and 64 horizontal BC408 plastic scintillator counters, each one read by a Photonis XP2262B photomultiplier (PMT) through a fishtail plexiglass lightguide. The 128 counters are assembled into 4 quadrants per plane, installed upstream of the LKr calorimeter. A custom control system has been developed for the remote control of the HV system, based on SY403 mainframes from CAEN, available since NA48. Four modules of the front-end electronics developed for the Large Angle Veto (LAV FEE), consisting of Time-over-Threshold (ToT) discriminators with double threshold setting availability have been exploited to process CHOD signals. The nominal track rate crossing the NA48-CHOD acceptance at full beam intensity is 13 MHz, due to beam kaon and pion decays and, to a lesser extent, to beam halo muons. Single tracks amount to about 95% of this rate. Further non-negligible contributions to the total hit rate are given by interactions in the material of upstream detectors and back-scattering from the LKr calorimeter. All these effects lead, at nominal beam intensity, to a total hit rate of about 35 MHz in each detector plane, estimated with 2015 data. The slabs close to the beam pipe must stand single hit rates up to 2 MHz, much higher than the outer ones. To cope with this high data rate, the DAQ system exploits one TEL62 board fully equipped with four

TDCB. Half of the TDCB input channels are left unused to keep input rates into buffer memories within suitable working limits. Thanks to a feature of the TEL62 board, a further reduction of input data can be obtained by masking one of the two sets of the differential signals discriminated at two different thresholds. For the same purpose, the analogue pulses in each quadrant can be rearranged to fit the rate limits for groups of channels. During the 2015 run, the NA48-CHOD has been providing L0 trigger signals (and L0 trigger reference time) to select events with charged particles in the final state. A loose trigger selection requires at least two counters hit within an adjustable time window. A more refined track trigger selection is provided by the coincidence between the signals of at least one vertical and one horizontal counter of adjoining quadrants within an adjustable time window. This coincidence allows track times to be corrected for the hit impact point position on the detector slab. The firmware for the trigger logic and for the impact point time correction is implemented in the TEL62 FPGAs. The use of ToT discriminators allows hit time corrections due to small signal slewing. Given the height of the NA48-CHOD pulses compared with the discrimination thresholds, slewing effects do not remarkably affect the track time precision. A preliminary distribution of the time difference between vertical and horizontal slabs is shown in Figure 28(left), based on data collected in 2015 at 1% of the nominal beam intensity. The hit times have been corrected for the impact point position on the slab and for slewing effects. The sigma of the Gaussian fit is 400 ps. Using the same sample of data collected in 2015, a preliminary distribution of the time difference between the NA48-CHOD track candidates and the KTAG beam kaon candidates is also shown in Figure 28 (right). The sigma of the Gaussian fit superimposed to the data is 215 ps, with a contribution lower than 100 ps from the KTAG detector. Similar preliminary sigma values have been measured at higher beam intensity (10% and 40% of the nominal one). These results are consistent with the performance measured by the NA48 experiment.

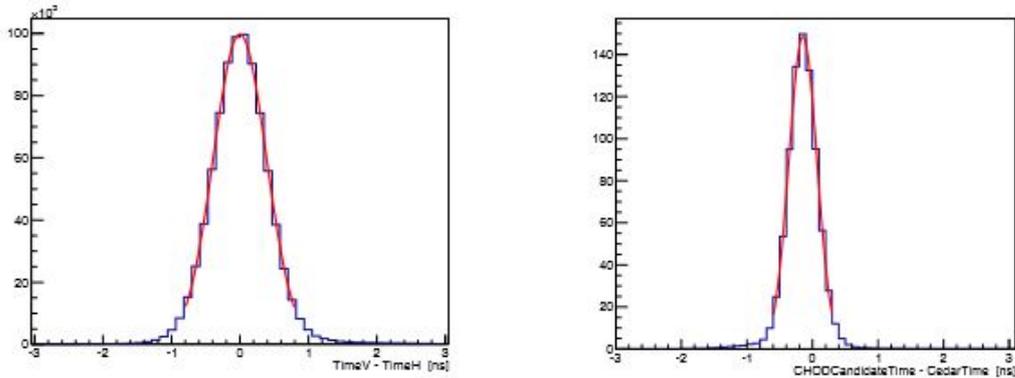


Figure 28: Left: distribution of the time difference between one horizontal and one vertical scintillator slab hit by a track (2015 data, 1% nominal beam intensity, preliminary results). A Gaussian fit with 400 ps sigma is superimposed to the data. Right: distribution of the time difference between NA48-CHOD track candidates and KTAG kaon signals (2015 data, 1% nominal beam intensity, preliminary results). A Gaussian fit with 215 ps sigma is superimposed to the data.

9 NewCHOD

The lateral acceptance region for charged particles downstream of the RICH and upstream of the LKr calorimeter is defined by the LAV12 detector with an inner radius of 1070 mm, and the IRC detector with an outside radius of 140 mm. The CHOD is a two-dimensional array of scintillator tiles installed upstream of the LAV12 with the main function of providing a basic element for the L0 trigger when at least one charged particle crosses the circular crown with the dimensions defined above. The subdivision of the acceptance surface into two-dimensional tiles leads to an optimized distribution of hit rates, and different groups of tiles can be selected to contribute to specific trigger requirements.

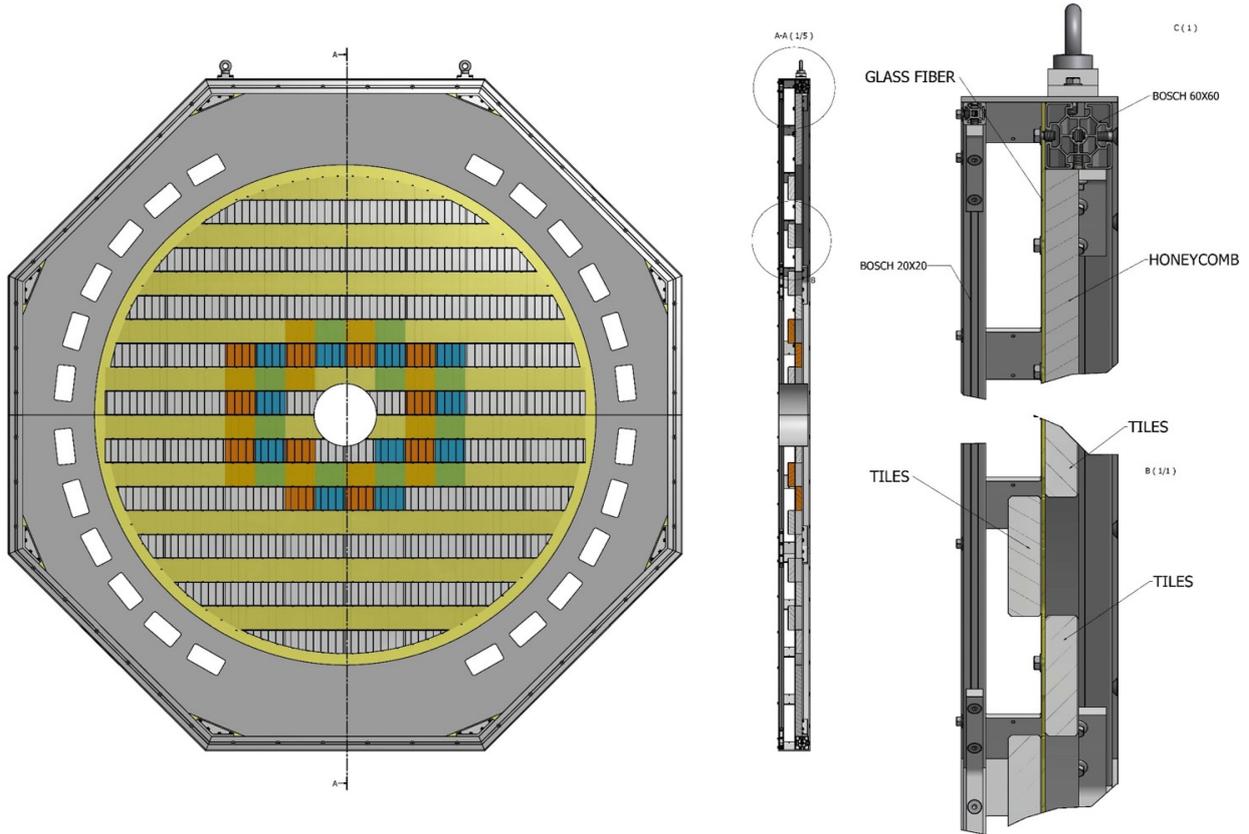


Figure 29: The new CHOD detector is mounted on the front face of LAV12 (left). It consists of one detection layer with 152 scintillating tiles. The tiles are mounted front and back of the thin G10 support panel (bottom right). At the periphery the structure is stiffened with honeycomb and aluminium construction profiles (top right).

In each quadrant, a 30 mm thick plastic scintillator is divided into 38 tiles. Except for three tiles near the external circular edge, their heights are 108 mm and their centres are spaced vertically by 107 mm, resulting in a 1 mm overlap. This is possible by placing alternatively one row of tiles on the upstream and the next on the downstream side of a 3 mm thick central support foil (G10 with 35 micron Cu lining on both sides), suitably perforated for the passage of 4.4 mm wide, 0.25 mm thick, 316L steel panduits (two per tile) to secure the tiles firmly in their positions (see Figure 29). Laterally most tiles are either 134 mm or 268 mm wide. The scintillation light, collected and transmitted by 1 mm diameter KurarayTM Y11 S wavelength shifting fibers, is detected by arrays of $3 \times 3 \text{ mm}^2$ SensLTM SiPMs on mother-boards located on the periphery of the CHOD.

Fiber lengths range from 135 cm to 200 cm, depending on the relative positions between a set of tiles and their motherboard.

The overall lateral dimensions of the CHOD are those of an octagonal box with inscribed circle diameter of 3100 mm. The structure is divided into two horizontal halves at the height of the beam. The thickness of the box, including the 0.5 mm thick foils which seal it longitudinally, is 140 mm. The structure is built with Aluminum BOSCH profiles with additional lateral 140 mm wide, 5 mm thick Al plates, with suitable openings (except on the top and bottom sides) for the mother boards (20 in total) located on the vertical and 45 deg inclined sides of the octagon.

On the tiles, 1.5 mm deep and 1.1 mm wide transverse grooves are milled every 16 mm. The fibers are glued in the grooves with EPO-TEKTM 301 optical epoxy. The odd and even numbered fibers of one tile are connected to different SiPMs. In total there are 152 tiles with 304 SiPMs arrays and the corresponding pre-amplifies. Pre-amplifiers are plugged on the mother boards in direct connection with the SiPMs. Their LEMOTM outputs can be chosen to correspond to either the individual SiPM arrays detecting the light from the two interleaved sets of fibers of each tile, or to their linear OR. The signals are shaped using constant fraction discriminators to improve the trigger time resolution, and read out by a 512-channel TEL62 board equipped with four TDCB mezzanines. The SiPMs can be accessed directly without opening the front or back of the octagonal box. Before mounting, each tile has been tested with cosmic rays, and it has been checked that the efficiency of the coincidence between the pre-amplifier outputs of the two sets of fibres exceeds 99%

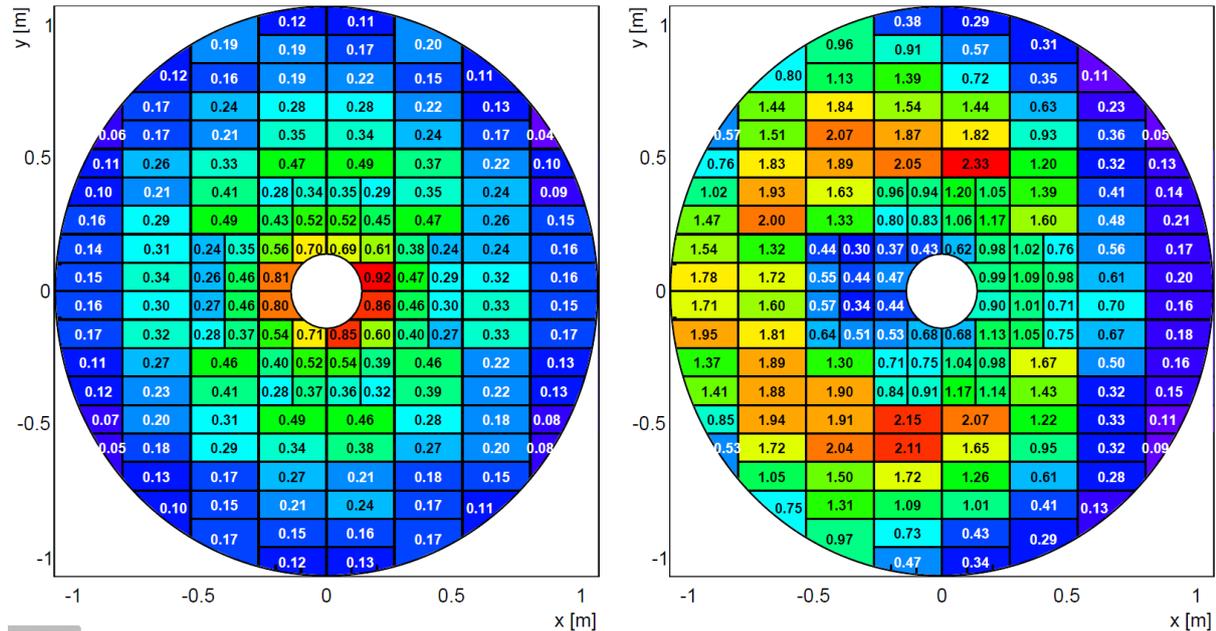


Figure 30: Left: the expected rates in CHOD tiles at nominal beam intensity, calculated with MC simulations (in MHz). Right: acceptances for $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ decays satisfying the signal selection conditions $105 \text{ m} < z_{\text{vtx}} < 165 \text{ m}$ and $15 \text{ GeV}/c < p_\pi < 35 \text{ GeV}/c$ in each tile, calculated with MC simulations (in percent).

The expected track crossing rate in the CHOD at nominal beam intensity evaluated with Monte Carlo simulations is about 13 MHz. However the corresponding total hit rate over the counters is about 45 MHz, due to the high hit multiplicity events produced by interactions of photons and beam pions upstream of the detector. CHOD tile layout

details and the expected signal rates in tiles at nominal beam intensity (from beam kaon decays, beam pions and halo muons), obtained with Monte Carlo simulations, are shown in Figure 30 (left). The lateral distribution of pions from $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ events satisfying the standard signal momentum and decay vertex position selection conditions, obtained with Monte Carlo simulations, is shown in Figure 30 (right).

10 Trigger and Data Acquisition

As the focus shifted from detector commissioning to operation, the preparation for the 2015 run saw an increasing participation in Trigger and Data Acquisition activities, although progress continued to be paced by manpower shortage issues. While the ambitious advance plans to have the complete system ready before the start of the run could not be fulfilled, the run itself saw a continuous improvement in the capabilities of the TDAQ system, with all pieces eventually being integrated together for the first time during the final weeks.

While waiting for the commissioning of the missing detectors and parts of the TDAQ, a large fraction of the time was devoted to low-intensity running, useful to assess detector performances in conditions virtually unaffected by pile-up effects. Some running time was devoted to sampling beam conditions closer to full intensity, which was important for uncovering and addressing problems of different kinds, which only become evident under stressful rate and radiation conditions. This allowed the gathering of much more information than previously available on the possible limitations of each element of the TDAQ, although a complete coherent picture will still require some time to be achieved. Overall, hints from the last day of the 2014 run were confirmed, with the TDAQ performance being limited by significant time non-uniformities of instantaneous rates.

10.1 Common infrastructure

The common TDAQ infrastructure ran smoothly throughout the run, although back-pressure signals on L0 triggers were still not tested. The limitations experienced in 2014 related to the common scheme for sub-system runtime initialization and configuration were overcome by a faster, more flexible and less error-prone system, which is now deployed everywhere. The Run Control system and the On-line Monitoring also had continuous improvements both before and throughout the run, reaching a stable status.

10.2 Common TDC-based TDAQ system

All TEL62 and TDCB boards were produced, with roughly 10% not passing initial acceptance tests. Due to the lack of foreseen manpower for complete testing and fixing, the effort focused on securing the availability of all hardware required for running, with a limited number of spares available while running; boards proved to be rather robust, with very few failures experienced during the run, most of those being related to crate issues, incomplete acceptance testing, or some weakness in the cable connections to the TDCB as a consequence of the high-density of inputs.

The major task, undertaken after experiencing hints of rate handling problems at the very end of the 2014 commissioning run, concerned improving the performance of the TEL62 firmware. While average hit rates were found to be well within design specifications, rate saturations do occur even below nominal beam intensity, with instantaneous

rates sometimes exceeding design capabilities due to both catastrophic events in which a very large fraction of detector channels fire, and very significant beam intensity fluctuations along the spill, at different time scales; while short-time fluctuations are normally absorbed by the built-in de-randomizing capabilities of the system, localized high rates beyond design capabilities lasting for time periods much longer than the time scales of the DAQ processing, e.g. related to 50 Hz beam intensity fluctuations, cannot be handled in a lossless way without over-designing the system for sustained rate capabilities several times the nominal. The presence of anomalous rates initially resulted in failures of single sub-systems which, somewhat distressingly but quite appropriately according to the experiment design, requiring a tight coupling of the TDAQ elements to guarantee vetoing efficiency, hindered overall data taking. A large effort was put into a "soft" handling of such failures to allow data taking to proceed in a controlled way during intensity peaks.

Before the start of the 2015 run, a large effort was put in increasing the capabilities of the system, in order to allow it to fully handle the maximum hit rate which can be provided by the hardware, and allow controlled behavior when exceeding specifications. The laboratory test system was much improved to allow injecting high hit rates into a single-board system with variable space and time distributions. The missing 10% of performance was achieved with a major rewrite of the core part of the TEL62 firmware, which can now read all the data the TDC chips are capable of producing before failure. A 50x increase of the size of de-randomizing buffers and the introduction of rate limiters to reject data in a controlled way at an early stage during high instantaneous rate peaks helped to make the system more smoothly running. At some stage, lab testing cycles show diminishing returns, as issues related to real experimental conditions on the large TDAQ system cannot be reproduced on a smaller scale test bench.

With the core DAQ part under control, the focus shifted to the integration in the TEL62 firmware of the primitive generators for the L0 trigger provided by the different sub-detector groups, which led to further instabilities due to the negative interference of parts of the code, due to imperfect control of the time constraints within the design, which could result in data corruption: overcoming these issues was the most time-consuming task which took place in repeated cycles during the first part of the run, also facilitated by speed-ups by many orders of magnitude in automated test capabilities.

Finally, when beam intensities around 50% of nominal were tested, some new issues appeared, such as failures related to bursts of large-size events which limited the L0-triggered output rate capability, related to L0 trigger rate, detector occupancy, and chosen readout window sizes. Some of these issues were quickly solved while some are still under investigation.

High-intensity running also confirmed the presence of radiation-related firmware effects: a hardware monitoring system was quickly set up and showed that most FPGAs indeed report some memory corruption within one hour of full-intensity running; while such an occurrence does not necessarily indicate a malfunctioning, it often does, and memory corruption was indeed shown to be present when some misbehavior appeared. Since only integrated dose information, rather than particle type flow, is available at selected points in ECN3, one cannot directly compare radiation fluence conditions during the run with those in which the TEL62 boards were tested at facilities in the past, but the order of magnitude of the corruption rate is consistent with expectations. Other systems, such as CREAM and network switches, also occasionally showed some radiation-related effects.

Such failures require a (quick) firmware reload and a (lengthier) re-initialization phase;

both monitoring and recovery capabilities were integrated in the Run Control system, and it is expected that additional shielding of electronics in most irradiated areas can reduce the number of re-initializations required. Additional data integrity controls are being implemented in order to allow a better control of possible data corruption, both on-line and in the collected data.

Statistical off-line data checks allowed us to measure data anomalies related to the documented intrinsic non-linearities of the HPTDC chips. These contribute to the intrinsic time resolution of the electronics by about 10%. On the hardware side, several faulty boards were recovered while a fraction remains for which the attempted repairs were not sufficient and require more thorough investigation to be fixed; pressure on this task, competing for manpower, was somewhat relieved by the additional production of 10 extra TDC boards, some of which were allocated to new sub-detectors (HASC) or to reducing the load on old ones (CHANTI), while others are spares.

While any machine-related improvement to the time structure of the beam would be extremely beneficial to the performance of the TDAQ system, in order to reduce instantaneous rate effects some measures were taken, such as removal of pulse trailing-time digitization where it was not useful, and further spreading of hot detector channels over a larger number of TDCs (e.g. in the CHANTI detector). The exploitation of the output links towards the PC farm is non-optimal, and discussions took place on moving towards a complete fair sharing of the network between data and trigger primitives, but such a major modification was considered to be unrealistic to be attempted in the short time available between the 2015 and 2016 runs.

10.3 Other Systems

As usual, since details on these systems are also available elsewhere, we will summarize some general information on other DAQ systems, mostly focusing on their integration within the general TDAQ infrastructure.

The calorimeter(s) readout system, based on the CREAM modules, ran smoothly within the global TDAQ system, and also accommodated MUV1 and MUV2 signals, as well as SAC and IRC. All these systems are read after the L1 trigger. Some rate limitations were discovered during higher intensity running and have been addressed after the end of the run with a new firmware revision.

The integration of the readout system for the drift chambers within the TDAQ system was one of the highlights of the late part of the 2015 run. The required L0 trigger matching firmware was completed, eventually allowing the reading of the chambers' data together with the rest of the event, resulting in a much improved understanding of the detector behavior and a faster turnaround in fixing remaining issues. Initial issues with data corruption, changing time alignments, and limitations in packing events together, which forced the TDAQ system to run at reduced rates, were eventually overcome. Also, the large observed readout inefficiency issues at high intensities have been addressed and fixed since the end of the run. Together, with other bug fixes, the system is now in a state which should experience no limitations at nominal intensity and beyond.

The Gigatracker readout system was also integrated within the TDAQ system in the late part of the 2015 run, although the standalone system remained active as well after the L0 trigger matching was deployed. The system experienced some instabilities and some intensity-dependent hangups which prevented it from running at higher intensities. The amount of data was found to be larger than expected, and hit limiters were introduced

in the preparation for the 2016 run, the GTK readout infrastructure was expanded, but it was also decided to move the readout of this detector to after the L1 trigger, thus relaxing the requirements on its performance. Some issues in the collected data might point to readout issues and are currently under investigation.

As mentioned above, the SAC and IRC sub-detectors (few channels with high hit rates) were read out through the CREAM flash-digitizers, while still also using the TDC-TEL62 readout; this double readout system turned out to be very useful in allowing a quick time alignment of the calorimetric L0 trigger to the other primitive generators, and will be maintained in 2016. No significant progress could be made on the dedicated high-rate digitizing system foreseen for SAC and IRC.

10.4 L0 Trigger system

Commissioning of the digital L0 trigger was the priority for TDAQ in 2015. Such a system was used exclusively throughout the run, starting with single elements and gradually growing to include more of them; it proved to be versatile, easy to control, and synchronize, although it also showed some limitations in the implementation, mostly related to manpower issues. Firmware development of the NA48 CHOD and RICH L0 primitive generator stopped after the 2014 run and could resume shortly before the 2015 run on the CHOD only, for which FPGA resource utilization was reduced by using a more coarse-grained impact-point correction. A thorough review of possible L0 triggers was done before the run, but their implementation was limited by the actual performance of the primitive-generating firmware modules. The elements of the L0 trigger are:

- LAV: time-clustered hit multiplicity from LAV station 12;
- MUV3: time-clustered hit multiplicity from entire detector;
- NA48 CHOD: time-clustered, impact-point and slewing corrected multiplicity from entire detector;
- RICH: time-clustered hit multiplicity from entire detector by exploiting a single additional TEL62 board reading the analogue OR of discriminator outputs;
- Calorimeters: energy, cluster counting, time and position based on 4x4 LKr calorimeter cells data, plus MUV1 and MUV2 energy.

While the hardware daughter-cards for communication among different TEL62 boards are now fully available and tested, it was not possible to devote some effort to their integration in the firmware, therefore, RICH and LAV L0 primitives are generated on single boards, and multiple-board primitive generation is postponed to next year, but this has a relatively limited impact on the L0 trigger performance.

The LAV, MUV3, and NA48 CHOD worked satisfactorily during the run at low intensities, but the latter two were known not to be able to fully withstand the high-intensity rates, reasons being understood but not eliminated during the run. Due to its halted development, the RICH and CHOD were forced to share the same firmware, definitely a non-optimal solution.

The calorimetric L0 trigger hardware was fully installed, and energy-based primitives were available for triggering in the last weeks of the 2015 run, while cluster-counting information is expected to be ready for testing some weeks after the start of the 2016 run.

Data were collected in 2015 with different mixtures of the above primitive generators, and an assessment of trigger efficiencies from data was started. A fully independent readout system for L0 trigger primitives from LAV, CHOD, RICH, and MUV3 ran continuously and turned out to be very useful for trigger consistency checking and debugging; this system will be retained for future runs, although, due to data storage limitations, data will be saved for only a sample of bursts.

Both of the two alternative L0 trigger processor systems were improved and tested before the run, although development of the PC-based system later stopped due to the overlap of the development team with GTK readout activities. The latter system was briefly tested during the run by recording produced triggers, but the FPGA-based system was used throughout the run for triggering. Several new issues were uncovered at increased intensities, which temporarily prevented data taking due to e.g. lack of strict time ordering of primitives and remaining close-in-time illegal triggers, all of which have been fixed since.

Intrinsic limitations were the maximum number of 7 participating detectors (common to both systems), the limited number of 8 trigger masks, and most importantly, the requirement for a single "reference" detector to be present in each produced trigger, hindering the collection of samples for efficiency measurements.

One common issue is the maximum input rate of primitives, which is strictly limited by the capacity of Gbit Ethernet links to 13 MHz: this is the expected primitive rate for tracks at nominal intensity, thus insufficient when intensity fluctuations are taken into account.

After the run, the FPGA-based system was further improved, and the capability of providing triggering without the "reference" detector was introduced, while also simplifying the way time coincidences are performed, to allow better control.

While the input rate limitation might not be an immediate issue when running significantly below nominal intensity, it was discussed at length, and an easy solution might have been found by reducing, in a clever way, the size of primitive data - an easy solution which avoids major changes in the hardware or network and is expected to be tested during the 2016 run.

Because of the remaining uncertainties related to future developments of the crucial L0TP systems, an effort is being made to gain most information by having both of them simultaneously receiving input throughout the entire 2016 run, although only one of them will be actively used for triggering.

The experimental GPU-based trigger was partially deployed in 2015 and ran throughout the run in a parasitic way, collecting data from half of the RICH detector and producing primitives, which were recorded but not used for actual triggering. While it could not stand the rate because of limited input capacity and the missing hardware implementation of the data merging between two boards (which had to be performed much more slowly in a PC), it looks promising and helpful both to reduce the rate of the standard L0 trigger and to complement it for more physics channels with fast higher-quality primitives. A higher performance custom network card is expected to be used in 2016 to overcome the above mentioned limitations.

After the run, development of a completely new L0 primitive generating firmware suitable to both NA48 CHOD and RICH was started; writing has been completed and the system is now being tested, and is expected to be deployed for the 2016 run.

As a byproduct of the GPU trigger development, the firmware preparing the RICH data for GPU processing was much improved and debugged, and is now going to be

adapted to work for the standard RICH L0 primitive generation, as a fallback solution.

The new NA62 CHOD detector is also expected to be contributing to the L0 trigger, and firmware for it is being prepared by modifying, in a rather limited way, the one used for the MUV3.

The calorimetric trigger is expected to be significantly improved during the 2016 run, both with the availability of the mentioned cluster counting information, and with the implementation of the foreseen readout of trigger primitive data to the PC farm, which is important for higher level triggering: the hardware boards required for the latter are now ready for production, and both the above features are expected to be ready some weeks after the start of the run. Moreover, a full integration with the Run Control is expected to allow a simpler and flexible operation of this trigger component.

10.5 On-line Farm and High-Level Trigger

The farm experienced no significant hardware failures, but the data assembly and collection efficiency ranged around 80-90%: full understanding of the misbehaviors was hindered by the limited availability of debugging tools and the lack of dedicated expert developers to replace the original software designer.

Several issues were found, such as difficulties in consistently recording End-Of-Burst events which contain precious monitoring information, which failed for reasons related to their large size or incorrect implementation by some sub-detectors. Crashes of individual farm nodes were as frequent as one per burst at high intensities, but most of the reasons for these instabilities were since understood and fixed.

Understanding of the system improved significantly as the run progressed, but the complex interaction of the different TDAQ elements prevented a complete assessment on the capabilities of the current on-line farm to withstand nominal-intensity running. It is estimated that such an assessment will be possible in late summer, after collecting further experience with the current improved system.

Some basic L1/L2 infrastructure was finally available, together with general required features such as downscaling, bypassing, and just-in-time decoding.

Simple L1 algorithms were implemented, based on multiplicities in KTAG, RICH, CHOD, and LAV, with no event rejection at the beginning of the run. No L2 algorithm was implemented, and the overall rejection factor achieved was around 5, far smaller than required for full intensity running, and somewhat influenced by the presence of a beam-track component passing the L0 trigger. Algorithms ran within acceptable time budgets, but the complete integration of DAQ and HLT processes resulted in a rather crash-prone system and is going to be addressed in the future.

A major development after the 2015 run was the implementation of a L1 algorithm working on the straw spectrometer's data (not used at L0): after development based on 2015 data, it was integrated within the L1 environment to be used during the 2016 run. Further studies on 2015 data showed that a set of proposed L1+L2 algorithms involving the straws' and calorimeters' information could achieve a rejection factor of 25 with an acceptable signal efficiency. Since calorimeter data is only read after L1, the implementation of such algorithms requires the data from the calorimetric L0 trigger to be available in the event, which will not be possible until a later time during the 2016 run.

10.6 Conclusions

Apart from some smaller parts, the deployment of the entire TDAQ system was finally completed during 2015, with all its elements being eventually present and integrated together in the last period of the run.

The digital L0 trigger was implemented and exclusively used to collect data throughout the run, and the energy-based calorimetric trigger was integrated in it, thus completing the set of participating elements. Straws and GTK readout systems were fully integrated within the TDAQ, with all detectors being now read in trigger matching mode. High instantaneous rate effects limited the performance of some systems and measures were taken to limit their effect.

Some simple high-level triggers were implemented, although the achieved rejection factors were still far below those required to allow exploiting the full intensity beam with no event downscaling.

The performance of several parts of the system suffered in general from fluctuating instantaneous hit rates higher than expectations, which could be absorbed only partially when lasting for extended time periods.

A large amount of data was collected in different running conditions, setting the teams in a better situation for understanding and addressing the known issues. The TDAQ experience gained during the 2015 run was collected in a detailed note to help assessing the performance and planning for priority actions. The improvements prepared in view of the 2016 run are going to be tested in a coherent TDAQ test run two weeks before the beam start.

With the chosen strategy of running at the highest achievable rate consistent with tape writing bandwidth and good data quality and efficiency, given the present rate reduction achievable by the trigger levels and introducing no downscaling factors, the collaboration set as a realistic goal for the 2016 run to reach stable and well-controlled running conditions at a beam intensity half the nominal one, to collect a good set of physics-grade data as soon as the performance of all systems is declared to be satisfactory. This should eliminate, for the time being, issues in the sharing of the output bandwidth between data and L0 trigger primitives, so that major changes related to a more efficient usage of the on-line network are postponed and will be considered for next year, in case an ongoing network simulation proves that the present configuration cannot handle nominal beam conditions in the presence of intensity spikes and that the advantages are large enough to warrant the effort.

Further improvement in automated data quality checking by each sub-detector team is required in order to reduce the turnaround time for identifying and fixing remaining errors, and a thorough assessment of the L0 trigger efficiencies must be completed.

11 Publications and analysis of older data

Since the last NA62 SPSC review in April 2015, the collaboration has completed the publication of the NA48/2, "Search for the dark photon in π^0 decays" [1], and approved two physics analyses based on older data sets recorded by NA48/2 and by NA62, respectively:

- a search for Majorana Neutrinos in $K^\pm \rightarrow \pi^\pm \mu^\pm \mu^\pm$ decays, based on 2003-2004 Data;

- a measurement of the π^0 transition form factor (from a large sample of $K^\pm \rightarrow \pi^\pm \pi_D^0$ with π^0 Dalitz decay), based on 2007 Data.

Both preliminary results have been presented at Winter Conferences and draft papers are being prepared for review within the corresponding Collaborations and for timely publication. More analyses of both NA62 and NA48/2 data are still on-going: the high quality data recorded in former years are still valuable for exploring poorly known very rare kaon decays or conducting high precision measurements. For Master and PhD students working on hardware-related subjects, they are also an ideal practice as training for data analysis. The above topics include:

- a search for long-lived sterile heavy neutrinos ν_H the 50 - 350 MeV/ c^2 mass range produced in the $K^\pm \rightarrow \mu^\pm \nu_H$ decay (2007 Data);
- a precision study of $K^\pm \rightarrow \pi^0 l^\pm \nu$ form factors in both K_{e3} and $K_{\mu3}$ modes (2004 Data);
- a new measurement of the $K^\pm \rightarrow \pi^\pm \pi^0 e^+ e^-$ decay (2003-2004 Data);
- the first observation of the $K^\pm \rightarrow \mu^\pm \nu e^+ e^-$ decay (2003-2004 Data)
- a search for the rare and still unobserved $K^\pm \rightarrow \pi^0 \pi^0 \mu^\pm \nu$ (2003-2004 Data)

The newest generation of students is now working on the NA62 data recorded last year and will work on the data collected in the coming 2016-2018 data taking runs. The collaboration is actively contributing to major International Conferences and topical Workshops with NA62 Detector contributions and recently published or preliminary physics results from NA62 and NA48/2 data analyses. In the past year (April 2015 to April 2016), the collaboration speakers presented 34 talks and 1 poster to Physics Conferences and gave 14 talks and 14 posters to Instrumentation Conferences. More contributions are already foreseen in future 2016 Conferences.

References

- [1] J. Batley et al., "Search for the dark photon in π^0 decays", Physics Letters B 746 (2015) 178