The NA60 Experiment Readout Architecture

Michele Floris, Davide Marras, Gianluca Usai, André David, Peter Rosinsky, and Hiroaki Ohnishi

Abstract—The NA60 experiment was designed to identify signatures of a new state of matter, the Quark Gluon Plasma, in heavy-ion collisions at the CERN Super Proton Synchrotron. The apparatus is composed of four main detectors: a muon spectrometer (MS), a zero degree calorimeter (ZDC), a silicon vertex telescope (VT), and a silicon microstrip beam tracker (BT). The readout of the whole experiment is based on a PCI architecture. The basic unit is a general purpose PCI card, interfaced to the different subdetectors via custom mezzanine cards. This allowed us to successfully implement several completely different readout protocols (from the VME like protocol of the MS to the custom protocol of the pixel telescope). The system was fully tested with proton and ion beams, and several million events were collected in 2002 and 2003. This paper presents the readout architecture of NA60, with particular emphasis on the PCI layer common to all the subdetectors.

Index Terms—Data acquisition, heavy ion physics, PCI architecture, readout electronics.

I. INTRODUCTION

THE NA60 experiment at the CERN Super Proton Synchrotron [1] studies the dimuons produced in collisions of proton and ion beams on nuclear targets.

NA60 evolved from a series of previous experiments (the last one being NA50 [2]) from which it inherits some detectors. NA60 will improve some of the measurements made by these experiments and will complement them with some new measurements, thanks to a renewed experimental apparatus. In particular, NA50 observed an excess of dimuons in the mass region between the ϕ and the J/ψ (intermediate mass region), with respect to the expected sources (*open charm* and Drell–Yan). To clarify the origin of this excess it is necessary to separate the prompt dimuons from the muons resulting from the decays of *open charm* mesons. This requires precise tracking in the vertex region, which is achieved in NA60 thanks to state-of-the-art silicon detectors inserted in a 2.5 T dipole field, placed in the target region.

Even if the dimuon excess is seen to be due to other sources, such as the production of thermal dimuons, another signature of quark gluon plasma (QGP) production, the measurement of open charm production remains very relevant in itself and as the best normalization reference in the studies of charmonium suppression in nuclear collisions.

Moreover, the much improved mass resolution achieved thanks to the tracking in the vertex region, allows a detailed study of the low mass resonances ρ, ω, ϕ , complementing measurements done by previous experiments [4], [3].

The readout of all detectors is based on the PCI architecture: the detectors are directly interfaced with the acquisition PCs via a set of custom electronics cards. In NA50, the readout was based on VME/CAMAC. The choice to move to PCI for all the detectors is justified by the higher performances and lower cost of a PCI-based readout system. This scheme turned out to be very flexible, allowing to replace the old VME/CAMAC-like readout systems for the MS and to implement the complex custom system of the pixel detector with the same PCI card. Moreover, as it will be clear in the following sections, the readout scheme which has been implemented simplifies the interfacing with the DAQ software, as the interface is the same for all the detectors.

In the following sections, we will first describe the experimental apparatus, then the data acquisition (DAQ) system following a top-down approach: from the DAQ software to the interfacing with the front ends.

II. DETECTOR CONCEPT

The NA60 experiment measures dimuons produced in proton-nucleus or heavy-ion collisions. By measuring dimuon production, the experiment studies the production of vector mesons, the Drell–Yan process, thermal dimuon production and, for the first time in heavy-ion collision, the production of charmed mesons. Moreover, it is very important to measure the centrality of the nuclear collisions, in order to compare the more elementary-like peripheral collisions with the more violent central collisions, where new physics is expected to show up.

The identification of events where pairs of D mesons are produced is done through the measurement of the impact parameter of the muon tracks, i.e., the distance in the transverse plane between the extrapolated muon tracks and the interaction vertex, at the point along the beam axis where the collision took place. This measurement requires a very precise vertex identification, which is achieved with the new silicon detectors placed in the vertex region.

The NA60 apparatus, as can be seen in Fig. 1, is composed of the following main detectors: a muon spectrometer (MS), a zero degree calorimeter (ZDC), a vertex telescope (VT), and a beam tracker (BT), which are briefly described in the following sections.

A. Muon Spectrometer

The MS was built for the NA10 experiment. It is composed of eight multiwire proportional chambers (MWPCs) for pre-

Manuscript received May 31, 2003; revised January 20, 2004.

M. Floris and G. Usai are with the Università degli Studi di Cagliari, 09100 Cagliari, Italy, and also with INFN, Sezione di Cagliari, 09042 Cagliari, Italy (e-mail: michele.floris@ca.infn.it).

D. Marras is with INFN, Sezione di Cagliari, 09042 Cagliari, Italy.

A. David is with IST, 1049-001 Lisbon, Portugal, and also with CERN, CH-1211 Geneva, Switzerland.

P. Rosinsky is with CERN, CH-1211 Geneva, Switzerland.

H. Ohnishi is with RIKEN, Wako-shi, Saitama 351-0198, Japan.

Digital Object Identifier 10.1109/TNS.2004.828828

Fig. 1. The NA60 apparatus. The heart of the detector is the MS made by MWPCs, trigger hodoscope stations and a toroidal magnet. The ZDC is along the beam line inside the hadron absorber. The detectors installed in the target area are shown in an expanded view in Fig. 2.

Fig. 2. Target area detail. This region contains the new NA60 silicon detectors: the BT (downstream the target) and the VT (upstream the target, inside the dipole field). Also note the segmented target which allows to have a bigger interaction length minimizing reinteractions.

cise tracking of muons and four hodoscopes of plastic scintillator slabs that provide the trigger signal for the experiment. In between the first four and the last four tracking stations sits a toroidal magnet, needed to measure the momenta of the muons. The last trigger hodoscope sits behind an iron wall, placed after all the chambers, to ensure that only muons can trigger the experiment without degrading the tracking resolution.

The MS is separated from the target region by a hadron absorber, to ensure that the abundantly produced hadrons will not pollute the trigger hodoscopes and tracking chambers.

B. Zero Degree Calorimeter

The ZDC is used to determine the centrality of the events by measuring the forward energy of non interacting beam nucleons. It exploits the Cherenkov light produced by charged particles cascading into quartz optical fibers, the sensible element of the detector. This detector has been inherited from NA50.

C. Vertex Telescope

The main challenge of the NA60 experiment is to match the muon tracks seen in the MS to the corresponding tracks measured right at the vertex, so as to eliminate the deterioration effects induced by the multiple scattering and energy loss suffered by the muons while crossing the hadron absorber. This requires a very accurate measurement of the charged particles, in the vertex region, a task accomplished by the silicon VT (Fig. 2).

While a silicon microstrip telescope is enough to do the job in the case of proton-nucleus collisions, the hundreds of charged

particles produced in heavy-ion collisions would lead to an occupancy so high that the tracking would be impossible.

This imposes the exclusive use of truly bidimensional silicon pixel detectors, given their excellent granularity. Furthermore, since NA60 looks for very rare events, identified thanks to the very selective dimuon trigger, it must work with rather high interaction rates, unlike other (“minimum bias”) heavy-ion experiments. This imposes the use of radiation tolerant pixel detectors, a technology that only recently became available and which is being used by NA60 for the first time.

In view of delays in the availability of the pixel detector, during the June 2002 proton run a silicon microstrip detector was used. Since this detector has a small material budget and covers the muon angular acceptance even at 40 cm away from the target, a dimuon mass resolution even better than with pixel could be obtained. Nevertheless, because of its low granularity, its use is restricted to proton runs. Moreover, subsequent Monte Carlo simulations have shown that the best setup for proton runs is an hybrid telescope made by both strip and pixel planes. This hybrid setup will be used for the 2004 proton run.

D. Beam Tracker

Upstream of the target, a BT [5] (Fig. 2) measures the flight path of the incoming beam particles and determines the transverse coordinates of the interaction point, at the target, with a precision of around $20 \mu\text{m}$. These microstrip silicon detectors operate at a temperature of 130 K, in order to increase their lifetime in the extreme radiation conditions of the experiment.

Fig. 3. DAQ system architecture. Data throughput is referred to ion runs, unless indicated otherwise.

III. DAQ ARCHITECTURE

At the CERN SPS particles do not arrive on the target continuously, but rather within limited time intervals called “bursts” (typically 5 s long in proton runs and 6 s long in ion runs). The period of time in between two bursts is called “interburst” (typically 10 s long).

NA60 has adopted a scalable philosophy for its DAQ system (Fig. 3). Each detector is divided into one or more “partitions.” During the burst, when events are produced and detected, each partition is read out independently by the electronics, whose last stage is a general purpose PCI card (see Section IV-A). This PCI card is the same for almost every detector (with the single exception of the microstrip vertex detector [6]). It has a large local memory buffer (64 MB) where data are stored during the burst. These cards sit in commodity PCs which are called local data concentrators (LDC).

The readout cycle is started by a trigger signal produced by some special logic which gets the data from the MS hodoscopes and looks for patterns compatible with the presence of a dimuon inside the acceptance of the spectrometer.

During the interburst data are read from the local memory of the various PCI cards and stored on the main memory of the LDCs. Data from the various LDCs are then collected and put together (“built”) by a single PC, called global data concentrator (GDC). The GDC sends the global event to CERN’s Central Data Recording Facility (CASTOR [7]) for permanent storage on tapes.

A. DAQ Software

The NA60 DAQ software is based on the data acquisition and test environment [8] (DATE) framework, developed by the ALICE collaboration.

This software is designed to run as a distributed system of several hosts performing the tasks of readout (LDC), event building (GDC), run control (RC), online monitoring (MON), and message handling (INFO). The operating system of DAQ processors must be UNIX-like, they must share the same filesystem and must be connected by TCP/IP network. The DAQ chain in DATE (Fig. 4) works as follows: the detector-specific program *readout* running on LDC collects the data from the front-end electronics (from the PCI card buffer, in the case of NA60) and stores them in a local buffer. The buffer is shared with the *recorder* process that off-loads it, asynchronously sending the data over the network into the GDC. Here, the *gdcServer* (one instance per LDC) stores the received data into another buffer that is off-loaded by central *eventBuilder* process. The *eventBuilder* waits until all the subevents from the LDCs arrive, builds the full event (adding a special header), performs consistency checks and writes the event into a file. In a multi-GDC environment the events are distributed among the GDCs to share distribute the load. The adopted two-stage buffering scheme naturally balances the load in the system and decouples the event building from the readout stage. The whole system is controlled from a single point: the *runControl* process on RC host opens sockets to the DAQ hosts, where certain Internet daemons are running. On LDC, the *rcServer* daemon controls the *readout* and

Fig. 4. DAQ, RC, and monitoring in DATE.

recorder processes, on GDC, the *rcServer* and *ebDaemon* daemons control the *gdcServer* processes and central *eventBuilder*. To facilitate the online monitoring, a dedicated *mpDaemon* is started upon request and sends the events from the LDC or GDC buffer to the monitoring workstation: a fraction of events is sent from LDCs to the monitoring hosts to check the functionality of the detectors and the data quality during the runs. Each DAQ host can generate information or error messages that are collected by *infoLogger* host by *infoDaemon* Internet daemons (one per connection).

The NA60 DAQ system consists of central run control (RC) host, one or more LDCs for detectors (1xBT, 1xZDC, 1xMS, and 3xVT¹), one special LDC (reading nondetector information), three GDCs and several monitoring (MON) hosts (Fig. 3). The RC host contains also the DAQ software repository and exports it over NFS to all DAQ machines. It also serves as the *infoLogger* host and main control console.

Fig. 3 also shows the data throughput expected: up to 150 MB/burst are expected in ion runs and this results in a sustained flow of about 10 MB/s into tapes. The data rate is highly dominated by the VT (Section IV-E).

IV. THE PCI SYSTEM

As aforementioned, the basic unit of the PCI system is a general purpose PCI card. Each detector is interfaced to this card via a detector specific mezzanine, which is in general different for each detector. Here, we will first describe the general purpose

PCI card, then the particular implementation for each detector in the experiment.

A. PCI Card

The NA60 readout was initially based on the PCI-FLIC [9] (FLexible Input/output Card), developed by the EP/ED-DTb group at CERN. This card contains a programmable logic (FPGA) with an embedded hardware PCI core (ORCA OR3LP26) and a large memory buffer (32 or 64 MB). The FPGA was programmed using VHDL code and a synthesizer, plus a tool from Lucent for placement and routing. The development of the VHDL code was mainly done by the NA60 team.

This card has been successfully used in datataking runs in 2001/2002, validating the readout architecture.

Subsequently a new card [called compact and flexible design (PCI-CFD) Fig. 5] equivalent to the FLIC from the architectural point of view but with much faster logic (based on an ALTERA APEX EP20K100E FPGA) has been developed and adopted. Thanks to this new card the microstrip vertex detector will also be included in the readout scheme described here (see Section IV-E).

The PCI-CFD has been used in all datataking periods in 2003.

The interfacing with the PCI bus is handled via a PLX 9030 PCI bridge, which allows the user to program transactions on a local TTL bus, in a very simple way. This chip can be highly customized, and the customization values are written on a EEPROM directly from the PCI bus using a Linux program. Using an external component as a PCI to local bus bridge decouples the problems of FPGA design development from the

¹In June 2001, proton run 4 LDCs were used for the microstrip VT.

Fig. 5. PCI-CFD general purpose PCI card. This card is the basic readout unit of the NA60 experiment. It contains an Altera programmable logic, a PLX PCI bridge and a large memory buffer.

PCI interfacing, simplifying considerably the card's debugging and thus leading to a design which is more reliable and easier to maintain. In the PCI-FLIC the hardware PCI interface was embedded in the FPGA.

The PCI-CFD contains a large amount of memory as well (64 MB) and the connectors to plug mezzanine cards. However this card is not fully compliant with the IEEE P1386.1 standard for PMC mezzanine cards.

It is also equipped with a silicon serial number DALLAS DS2401, which permits to uniquely identify every PCI-CFD.

Both PCI-FLIC and PCI-CFD cards are designed to work as a PCI target, which allows to get more than satisfactory performances, with less problematics with respect to a PCI master: with an accurate choice of the host computer it is possible to reach a maximum bandwidth of almost 38 MB/s [10].

In the case of the PCI-FLIC the maximum measured bandwidth was some 15 MB/s. With the PCI-CFD the bandwidth is 30 MB/s; this number is not as dependent on the host computer as it was with the PCI-FLIC, where one could get 50%–60% differences in bandwidth between different computers.

Fig. 6 shows a block diagram and data flow of the whole application. The dashed line encloses blocks which are inside the FPGA. There are four main blocks (READOUT, EVENT FORMATTING, SDRAM CONTROLLER, and PCI INTERFACE), plus a CONTROL LOGIC, which triggers the other blocks, arbitrates accesses to memory, performs the handshaking with software and handles signals from the experiment (like trigger, burst, and busy signals).

Fig. 7 shows how the handshake between software and hardware is implemented. The software is continuously polling a status register implemented in the PCI card. The lowest two bits of this register contain the SPS (BURST/INTERBURST) status, while the third one (bit 2) is used to implement the handshake itself. When the burst is over and the last event has been collected the hardware sets bit 2. As soon as the software sees this bit it starts to move the data resetting it at the end. If a new burst arrives before bit 2 is reseted, bit 3 is set (by the hardware) indicating a timeout.

The READOUT block gets the data from the detector via the mezzanine. This block is detector specific and depends on

Fig. 6. Application block diagram. The dashed line encloses blocks implemented inside the FPGA.

the mezzanine. It will be discussed in detail in the following sections.

The EVENT FORMATTING block frames the incoming data within a header and trailer on an event by event basis. The header and trailer contain all the information necessary to build the “global” event, like number of data and total words, number of burst, number of event inside the burst, time of arrival of the event, and error flags. The data format is depicted in Fig. 8.

The SDRAM CONTROLLER writes the formatted events into the memory. During the interburst it is connected to the PCI Interface to send data from the on-board memory to the main LDC memory (see Section IV-A).

B. Linux Driver for the PCI Cards

The operating system used for the DAQ hosts is Linux. The version used is CERN Linux 6.1.1, which was the CERN standard when the project started (2001). This version is based on a kernel release 2.2.x.

A driver for both cards, FLIC and CFD, has been developed for this operating system. It hides the hardware peculiarities of the PCI card (like absolute offset of its internal registers and memory), making it available as a logical entity, through an entry in the `/dev/` directory in the Linux tree. Through the driver one can simply access the card using the C call `open(“/dev/cfd0”, O_RDWR)` and using some macros that define the local offset of memory areas inside the card, instead of opening the `/dev/mem` device and accessing memory areas using absolute offsets.

It also sets all the important environment parameters needed to achieve the best possible performances (i.e., the `mtrr` registers which allows caching and burst transfers [10]).

Finally, the driver allows to get some parameters from the PCI card through the Linux `ioctl()` system call, like the unique identifier (Section IV-A) or the number of detected cards in the system.

Fig. 7. Readout cycle and handshake between software and hardware.

Fig. 8. NA60 data format. The header and trailer are common to all detectors. The “data” words framed between header and trailer are encoded in a detector-specific format.

C. MS

MS data (both chambers and hodoscopes) are read and temporarily stored by a system of CAMAC modules developed at CERN in the mid 70 s, called receiver memory hybrid (RMH) [11]. These modules can contain only one event and must then be read before a new trigger can be accepted.

The total number of readout channels of the MS (MWPCs + hodoscopes) is about 20 000, giving an average event size of about 1 kB.

In the previous experiments data were read by a custom VME module (called MEMRMH) which contained a small memory buffer (4 MB), limiting the number of events which was possible to acquire in a single burst to about 4000. This VME module was connected to a VAX-based DAQ through a system of transputers.

Fig. 9. The MS RMH mezzanine. This mezzanine is “passive” and converts incoming signals to TTL, so that they can be used inside the PCI card. All the readout logic is implemented in the PCI card.

The new readout based on PCI does not have this limitation anymore, and the number of events which is possible to acquire is limited only by the dead time of the detector, which has also been reduced and is at present around 120 s on average (mainly due to the CAMAC system). The MS is the slowest detector in the experiment.

The mezzanine used for this detector (called RMH mezzanine, Fig. 9) is very simple. It converts the incoming signals from ECL and NIM, used by the front-end electronics to TTL. The protocol is implemented inside the READOUT block in the FPGA of the PCI card.

The RMH protocol is a double handshake protocol, similar to the VME one. Fig. 10 shows a typical RMH cycle. The PCI card is the master and asserts a “START_READ” signal which frames the whole transaction and then asks for the first data word, asserting the “ENCODE” signal. The CAMAC modules, which are the slave, respond asserting the “DFLAG” signal when valid data are on the bus. When no more data are available for reading, the slave asserts the END_OF_READ signal.

Two PCI cards with RMH mezzanine are needed to read the MS (one for MWPC and one for hodoscopes) but they share the same LDC.

Fig. 10. RMH protocol cycle.

Fig. 11. Pixel readout system layout. The pixel chips are read out by the pilot chips which sends the data to the mezzanine through GLink, via the GOL chip.

D. ZDC and BT

The ZDC and BT readout applications are similar to the MS one and use the same mezzanine (RMH). In these cases, it is necessary to read a FERA CAMAC module.

This protocol is very similar to the RMH one (double handshake), the main difference being that the PCI system is now the slave.

The average event size is less than 1 kB for both the ZDC and the BT.

Although from the point of view of the performances these two detectors could be read using the same LDC, it has been decided to use two distinct LDCs for practical reasons.

E. VT

The microstrip VT has a slightly different readout architecture [6], based on a different PCI card and it will not be discussed here.

The silicon pixel telescope, on the other hand, represents the most sophisticated application.

Fig. 11 shows the pixel detector readout system. The detector is composed of 16 pixel planes based on the ALICE1LHCB [12] readout chip (pixel chip in the following) developed in the framework of the ALICE and LHCb collaborations at CERN. This chip is a radiation tolerant 32×256 matrix of readout cells, each $425 \times 50 \mu\text{m}^2$. Its operation is highly customizable: several internal DACs can be set via a JTAG interface [13].

Each plane contains four or eight pixel chips, bringing the total number of channels in this detector up to about 720 000. The expected maximum occupancy is about 3%–4%, giving a maximum event size of about 18 kB. Assuming a maximum number of triggers of 8000 the corresponding data throughput

Fig. 12. Pixel readout mezzanine. This mezzanine contains a programmable logic, which is used to implement the readout logic, and some FIFOs which are used to temporarily store the events.

is 140 MB per burst. The VT produces more the 90% of data in each burst (Section III-A).

Pixel chips are clocked at 10 MHz. During event readout 32 bit pixel rows (256 per chip) are extracted in subsequent clock edges, so that the readout time is 25 s per chip. Pixel chips sitting on the same plane are read out sequentially. Reading out four chip planes therefore takes around 100 s while reading out eight chip planes takes around 200 μs . The dead time is reduced by a multi-event buffer present inside the pixel chip, which allows to accept up to three new events during the readout stage. The presence of the multi-event buffer and the partitioning of the detector lower the effective dead time below the one of the MS (120 μs , Section IV-C).

The pixel chip is interfaced over a data/control GTL bus to a second radiation tolerant chip (the PILOT chip [14]). This chip on one hand converts the GTL levels to CMOS and distributes all readout control signals and JTAG commands coming from the counting room to the pixel chip. On the other it transmits the data downstream to a radiation tolerant serializer chip (the GOL [15]) which finally sends the data to the PCI system using the industry GLink protocol [16].

The pixel readout mezzanine, shown in Fig. 12, contains a FPGA which implements the communication with the PILOT chip through a serial LVDS line (JTAG commands, clock distribution and PILOT configuration commands) and gets the data via two HP 1034 GLink receivers from two different planes.

Fig. 13. The graphical user interface used to configure the pixel detector.

At this level data are zero suppressed, encoded and stored temporarily in a FIFO, which is subsequently accessed by the PCI card.

Commands are sent to the mezzanine, the PILOT chip and the pixel chip by a custom software which writes specific registers on the PCI card. The FPGA application of the PCI card in the case of the VT contains one extra block with respect to those depicted in Fig. 6, which acts almost independently and is used to send commands to the mezzanine through a serial link. Commands to be sent are written in a specific PCI register, then another register is accessed to trigger the transmission of the command. The mezzanine interprets the received command and acts accordingly (either it sets some internal parameters, either it forward a JTAG command to the pixel chip/PILOT chip system). Some commands produce an output which is read by the PCI card through the same FIFOs used for data and written in another PCI register. The software can access output resulting from certain commands through this PCI register (a typical example is the reading of a status register inside the mezzanine).

A complete software package has been developed for the Linux operating system which allows to easily configure the detector. It consists of two control programs, which are used respectively to configure internal parameters of the mezzanine (for example, it needs to know how many pixel chips are on the plane) and to send JTAG commands to PILOT chip and pixel chip. The JTAG controller is software-based: the hardware simply forward the bit sequence it gets from the software. This program is interfaced to a mysql database which contains the optimal settings for every pixel chip. Those two control programs are also used by DATE to configure the detector at run startup.

A graphical user interface (GUI, Fig. 13) has been developed for these programs, which allows to easily configure the detector.

Finally, a program performs a threshold scan of the detector, to find dead or noisy pixels and tests the performances of the detector.

V. DAQ PERFORMANCES

The new readout system has been used in several different conditions over the past two years.

In October 2001 (with the PCI-FLIC) the PCI electronics for the MS was used for the first time, in a proton test run. During this run, we reached a trigger rate as high as 9000 triggers in a 5 s burst. This run was meant to commission the new MS readout electronics only.

In the June 2002 proton run, all detectors (except the microstrip vertex detector) were working with the PCI-based electronics (with the PCI-FLIC), though only one pixel plane was available. Tracking in the vertex region was assured by the silicon microstrip VT. 800 000 dimuon events were collected.

In the low-energy Pb run of October 2002, using the PCI-FLIC card, three pixel planes were read out with the new system, allowing to test a small pixel telescope. The MS was not used during this run, while the other detectors (BT and ZDC) were included in the acquisition and read out with the PCI system. Around 30 000 000 events were collected.

In August and September 2003, in two test beams with pions and protons at the CERN SPS, eight planes of the pixel detector were tested using the PCI-CFD on beam for the first time.

In the October 2003 indium run, all the detectors were present, including a complete silicon pixel telescope made of 16 planes. During this run the pixel telescope, the BT and the ZDC were read out using the PCI-CFD card, while the MS was read out using the PCI-FLIC card. The average event size of the pixel detector was about 10 kB, and the trigger rate was about 6000/7000 triggers/burst. For testing purposes, short runs were taken with a trigger rate of more than 9000 trigger per burst. Even if the DAQ can stand this trigger rate it was decided to run with 6000 triggers/burst to minimize the pileup rate.

VI. CONCLUSION

NA60 is one of the first high-energy physics experiments to use an architecture fully based on the PCI bus for its DAQ system. This bus offers attractive features, like high performances at a relatively low price and very easy interfacing with the acquisition software. Moreover, using PCI hardware cores and bridges eases the PCI card development.

The readout system described here has been successfully used during several datataking runs over the past two years.

ACKNOWLEDGMENT

The authors would like to thank the EP/ED—DTb group at CERN and, in particular, H. Muller for the technical support and the many ideas in the early stages of the project.

They would also like to thank the ALICE pixel team—in particular, F. Formenti, G. Stefanini, K. Wyllie, A. Kluge, M. Burns,

P. Riedler, and M. Campbell—for their constant support in the development of the NA60 pixel detector. A special acknowledgment goes, of course, to the entire NA60 Collaboration.

REFERENCES

- [1] *Study of Prompt Dimuon and Charm Production with Proton and Heavy Ion Beams at the CERN SPS*, May 15, 2000. NA60 Proposal.
- [2] *Study of Muon Pairs and Vector Mesons Produced in High Energy Pb-Pb Interactions*, Nov. 1991. NA50 Proposal.
- [3] M. C. Abreu *et al.*, "Low mass dimuon production in proton and ion induced interactions at the SPS," *Eur. Phys. J.*, vol. C14, pp. 443–455, 2000. NA38/NA50 Collaborations.
- [4] G. Agakishiev *et al.*, "Enhanced production of low mass electron pairs in 200-GEV/U S-Au collisions at the CERN SPS," *Phys. Rev. Lett.*, vol. 75, pp. 1272–1275, 1995. CERES Collaboration.
- [5] P. Rosinsky, "The cryogenic Beam Tracker of NA60 for heavy ion proton beams," *Nucl. Instrum. Meth. A*, vol. 511, no. 1–2, pp. 200–204, Sept. 2003.
- [6] A. David *et al.*, "The readout system of the NA60 silicon microstrip tracker," *Nucl. Instrum. Meth. A*, pp. 118–126, Aug. 2003.
- [7] J. P. Baud *et al.*, CASTOR status and evolution. presented at 2003 Conf. Computing in High Energy and Nuclear Physics. [Online] <http://www.slac.stanford.edu/econf/C0303241/>
- [8] M. Arregui *et al.*, "The ALICE DAQ: Current status and future challenges," *Comput. Phys. Commun.*, vol. 140, pp. 117–29, 2001.
- [9] H. Muller, F. Bal, A. David, D. Dominguez, and J. Toledo. A flexible PCI card concept for data acquisition: The PCI-FLIC. presented at 12th IEEE-NPSS Real Time Conf.. [Online] <http://ific.uv.es/rt2001/>
- [10] A. David. (2001, June) Performance Issues in PCI Reading. [Online]. Available: <http://adavid.home.cern.ch/adavid>
- [11] J. B. Lindsay, C. Millerin, J. C. Tarle, H. Verweij, and H. Wendler, "A fast and flexible data acquisition system for multiwire proportional chambers and other detectors," *Nucl. Instrum. Methods*, vol. 156, p. 329, 1978.
- [12] K. Wyllie *et al.*, A pixel readout chip for tracking at ALICE and particle identification at LHCb. presented at 5th Workshop Electronics for LHC Experiments. [Online] http://lebwshop.home.cern.ch/lebwshop/LEB99_Book/LEB99_Proceedings.html
- [13] *IEEE standard test access port and boundary-scan architecture*, IEEE STD. 1149.1a, 1993.
- [14] A. Kluge. The ALICE on detector pixel PILOT system—OPS. presented at PIXEL 2002 Workshop. [Online] <http://www.slac.stanford.edu/econf/C020909/>
- [15] P. Moreira *et al.*, "A 1.25 Gbit/s serializer for LHC data and trigger optical links," *Proc. 5th Workshop Electronics for LHC Experiments*, pp. 194–198, Sept. 20–24, 1999.
- [16] Agilent HDMP-1032, HDMP-1034, Transmitter/Receiver Chipset [Online]. Available: <http://www.semiconductor.agilent.com>