The HADES DAQ System: Trigger and Readout Board Network

Jan Michel, Michael Böhmer, Ingo Fröhlich, Grzegorz Korcyl, Ludwig Maier, Marek Palka, Joachim Stroth, Michael Traxler and Sergey Yurevich for the HADES collaboration

***Abstract*—HADES is a di-electron spectrometer operating at GSI, Germany. Currently, the HADES data and trigger system is being upgraded. The main aim is to substantially increase the event rate capabilities by a factor of up to 20 to reach 100 kHz in light systems and 20 kHz in heavy ion reactions. The data rate will be about 400 MByte/s in peak. In this context, the complete readout system has been exchanged to use optical communication.**

In this contribution we present a general-purpose real-time network protocol that has been developed to meet the new requirements. These include strong timing constraints with latencies less than 5 μs for endpoint-to-endpoint communication through up to 10 intermediate hubs in a star-like network setup. In summary, this network connects over 500 FPGAs distributed over the whole HADES detector.

Monitoring and slow control features as well as readout and trigger distribution were joined in a single network protocol. Hence, channel multiplexing with inherent arbitration by priority is implemented. Switching between channels takes less than 100 ns. For control and monitoring, a dedicated channel provides a data bus with a virtual address space spanning the whole network.

The configuration is highly flexible and thus adaptable to different experimental setups and hardware.

I. Introduction

THE study of properties of nuclear matter at high densities and temperatures is one of the challenging research topics in hadronic physics. In the laboratory this is realized by heavy ion collisions where, for a short time, such states are formed in a fireball, before the decay products emerge and are measured by detector systems.

One of the main probes to have an undistorted insight into the fireball over the entire history of the heavy ion reactions are di-leptons (e^+e^- or $\mu^+\mu^-$) as they are not hampered by final-state interactions. Now, after a successful decade of measurements of electromagnetic probes (summarized in [1], [2]) there is common agreement that only by systematic studies in various systems over a wide energy range, conclusions on fundamental questions in the field of nuclear science can be drawn: is there a phase of "free" quarks during the hot and dense phase, and what are the basic features of this new form of matter?

To measure these rare decays, the High Acceptance Di-Electron Spectrometer (HADES) has been installed at the **S**chwer**I**onen**S**ynchrotron (heavy ion synchrotron) SIS 18 at the GSI Facility for Heavy Ion Research in Darmstadt, Germany.

Work supported by EU FP6 (contract number 515876), the BMBF (grant 06FY9100I) and by the Hessian LOEWE initiative through the Helmholtz International Center for FAIR (HIC for FAIR).

Jan Michel is with the Institute for Nuclear Physics, Goethe University Frankfurt, Germany and the Helmholtz Graduate School for Hadron and Ion Research (HGS-HIRe)

Michael Böhmer and Ludwig Maier are with the Department for Physics E12, Technical University Munich, Germany

Ingo Fröhlich is with the Institute for Nuclear Physics, Goethe University Frankfurt, Germany

Grzegorz Korcyl is with the Jagiellonian University Krakow, Poland

Marek Palka is with the Institute for Nuclear Physics, Goethe University Frankfurt, Germany and the Jagiellonian University Krakow, Poland

Joachim Stroth is with the Institute for Nuclear Physics, Goethe University Frankfurt, Germany and the GSI Helmholtz Centre for Heavy Ion Research, Darmstadt, Germany

Michael Traxler and Sergey Yurevich are with the GSI Helmholtz Centre for Heavy Ion Research, Darmstadt, Germany

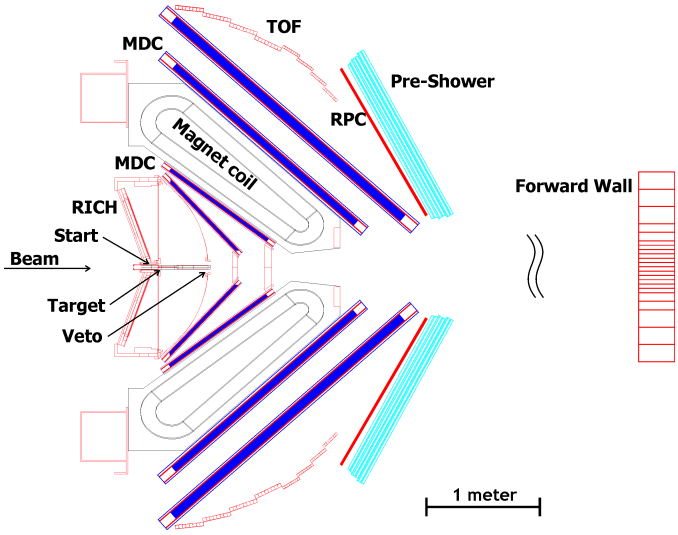

Fig. 1. A cross section of the HADES spectrometer. From inner to outer, the detectors are start and veto diamond detectors, a ring imaging Cherenkov detector (RICH), four layers of multiwire drift chambers (MDC), a superconducting magnet, a time-of-flight wall (TOF), resistive plate chambers (RPC) and a pre-shower detector. A forward hodoscope is located 7 m downstream from the target (distance not to scale). For the sake of clarity, only two of the six sectors are shown.

II. The HADES Experiment

The HADES setup [3] combines a high-acceptance magnetic spectrometer with a six-fold segmented toroidal field geometry aligned with the beam axis. A cross section of the spectrometer is shown in fig. 1.

The reconstruction of charged particles is based on their identification by specific detectors and momentum measurement. The latter is done using four planes of Multiwire Drift Chambers (MDC) toghether with a toroidal field created by a

superconducting magnet. In total, four sets of drift chambers with minimal material budget provide position measurements for charged particles in front of and behind the field region.

The spectrometer is complemented with detector systems for electron identification. A ring-imaging Cherenkov detector surrounds the target. Using a gas radiator, it provides electron identification by suppressing all hadrons with a gamma-factor below 18. Additional rejection power against fast pions comes from the signals of a pre-shower detector positioned in the polar angle region between 18 and 45 degrees in the outer shell of the detector.

The time-of-flight (TOF) detector consists of two parts: The outer TOF wall, composed of 64 scintillator stripes per sector, and an inner TOF wall, composed of 4 scintillator slabs per sector. The latter, providing only moderate granularity, has been replaced recently by a resistive plate chamber (RPC) [7]. detector system with more than 1000 single-cell, 4-gap counters.

The spectrometer is further supplemented by a forward hodoscope covering polar angles below 7 degrees, an region not covered by the other detectors. In total, the HADES spectrometer comprises about 100,000 individual detector cells.

The HADES collaboration has recently finished a series of measurements focused on a better understanding of dilepton emission from elementary [5] and light heavy-ion [4] collision systems at kinetic beam energies of 1 to 3.5 AGeV. An excess radiation of lepton pairs above contributions from hadron decays after chemical freeze-out of the nuclear fireball has been observed.

It is notable that one of the additional capabilities of HADES is that at the same time various hadronic reaction products (in particular pions and hadrons containing strangeness [6]) can be studied.

III. The Upgrade Program

HADES will continue the physics program at its current place at SIS 18 and then move to the upcoming FAIR[1] accelerator complex. There, HADES will continue its experimental program up to kinetic beam energies of 8 GeV per nucleon at the new SIS 100. Such measurements would allow to fill the gap between the low-energy "dense" nuclear matter and the "hot" region which is accessible at RHIC and LHC. At energies between 8 and 45 AGeV, the Compressed Baryonic Matter (CBM) experiment [8] will measure di-electrons and di-muons. HADES, on the other hand, will provide the link between 2-8 AGeV. The advantage of utilizing the present HADES setup is that it has all required features for di-electron measurements (as well as capabilities for the reconstruction of hadrons), and a higher geometrical acceptance than CBM for energies up to 8 AGeV. The much higher multiplicities of charged particle tracks, on the other hand, needs an upgrade of the trigger and readout system, which will be outlined in the following.

In the first place, the large data volumes expected in experiments with heavy collision systems (Au+Au) already at SIS 18 and with higher energies at FAIR require a data bandwidth

[1]FAIR: "Facility for Antiproton and Ion Research"

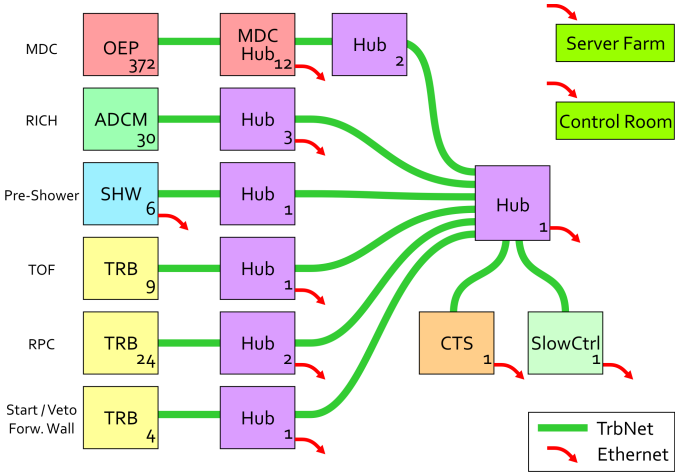

Fig. 2. A schematic view of the full network setup. The network has a tree-like structure, connecting all detector with the central control system. The numbers show the amount of boards of each type.

which cannot be achieved by the current system. And secondly, since the currently used data acquisition system was designed more than ten years ago, it was reasonable to reconsider the whole concept and make use of new technologies, like optical transmission.

The new DAQ system aims at a data rate of 100 kHz for $p + p$ collisions, and 20 kHz for heavy systems (Au+Au at 2 AGeV, Ni+Ni at 8 AGeV). To fulfill this goal, a new trigger and readout board (TRB), a multi-purpose electronic device with on-board DAQ functionality, has been developed [9]. This board has already been used during the last beamtimes for read-out of the HADES Forward Wall and Start/Veto detector. Furthermore, the TRB has been used by several other experiments, such as in prototypes for the PANDA-DIRC detector and the CBM micro vertex detector (both will be located at FAIR).

As the concept has been tested successfully during previous HADES data taking periods, its design has been extended to serve as the new standard readout module for all HADES sub-detectors by adding a high-speed connector and pluggable dedicated detector-specific add-on modules [10]. Besides the readout of the add-on modules this connector is used to connect other read-out electronics, supplying them with slow control and power.

During the upgrade, the complete data transport and parts of the digitizer hardware have been replaced by the TRB or other custom made electronics based on similar concepts.

IV. Description of the Components

The TRB, our main component, is equipped with a Virtex4 FPGA and 128 high-performance TDC [11] channels (resolution up to 30 ps). In this configuration it is used to take data from all timing relevant detectors of the HADES setup, namely TOF, RPC, the Forward Wall and the start-time and veto detectors. Moreover, the TRB features a TigerSHARC DSP processor and a 2 GBit/s optical link. For slow control purposes the board uses an EtraxFS CPU running an embedded Linux system and a 100 MBit/s Ethernet connection.

The backbone of the optical network is formed by hub boards. These are designed as high-bandwidth multi-link network devices equipped with 20 optical links, each running at a data-rates of 2 GBit/s. The routing between these links is accomplished by two large Lattice FPGA (ECP2M100). One link also serves as a bridge to Gigabit Ethernet (GbE) and provides a convenient way to send data to off-the-shelf computers.

The Pre-Shower detector employes 96 channel ADC boards equipped with three Lattice ECP2M50 FPGAs and two optical links. This board has been designed to make use of the add-on connector of the TRB.

The RICH (Ring Imaging CHerenkov) detector uses similar boards (ADCM): The digitizer part consists of a total of 60 ADC chips with 8 channels, 12 bit, 40 MSPS each. Data transport and front-end control is provided by a Lattice FPGA (ECP2M100) and a 2 GBit/s optical link. A microcontroller ensures correct voltage levels and controls the start-up sequence.

The MDC system uses a slightly different read-out concept [12] due to very strict space constraints. The signals from more than 27,000 sense wires are digitized by TDC boards which are controlled by small FPGAs (Lattice ECP2M20). A total of 372 TDC and FPGA boards (Optical End Point - OEP) are mounted directly on the frames of the MDC detector.

Due to space constraints the OEP functionality had to be integrated on a 4 cm by 5 cm PCB, including a board-to-board connector as well as connections for the time reference signal, the optical fiber and power.

Despite the functionality to configure and read out the TDC chips it features a variety of functions: Voltage regulators provide the power to the read-out electronics, an ADC monitors all voltages. Two seperate flash ROMs store two designs that can be loaded into the FPGA - one can be changed using the slow control system, one provides a "golden image" which ensures that the FPGA can be loaded again after a failure of the second flash ROM.

A 250 MBit/s fiber optical transmitter (FOT) provides the link to the MDC data concentrator board (MDC-Hub). Here, data from two chambers are collected. This board is equipped with 32×250 MBit/s FOTs and two 2 GBit/s transceivers, one of the latter used for Gigabit Ethernet. The large amount of data collected is handled by five FPGAs (Lattice ECP2M100).

An overview of the complete system can be found in fig. 2. All front-ends are controlled and monitored by two central devices: One is the slow-control system providing the interface between operators and the detector for configuration and monitoring purposes. The other is the central trigger system (CTS) controlling both trigger distribution and data read-out.

V. The DAQ Network

A. Architecture

As mentioned above, the main goal of the HADES upgrade is to increase the event rate capability up to 100 kHz and a total data rate up to 400 MByte/s in peak (250 MByte/s sustained). Data will be transported using an optical network which provides the necessary high bandwidth and also reduces

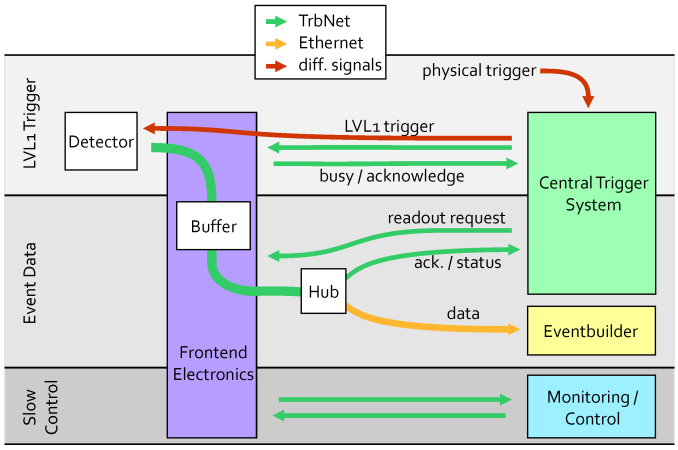

Fig. 3. Several types of data and information are transported in parallel using one common network setup

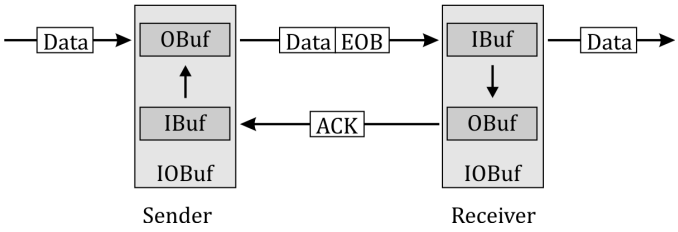

Fig. 4. Data transfer is controlled by a handshake mechanism. The sender attaches to each block of data (up to approx. 1 kByte) a special word (EOB) which forces the receiver to send back an acknowledge (ACK) to show that it is able to receive the next block of data.

the electromagnetic noise introduced in the detector during data read-out.

Due to space constraints, most boards have only one single bi-directional optical fibre connected. This requires that all individual accesses including trigger distribution, data transfer, control and monitoring have to be incorporated in one common network protocol. Their demands are, however, contradictory: the distribution of triggers usually requires almost no bandwidth but a very low latency of less than 5 μs for a broadcast to all systems. This constraint comes from the fact that for each event the trigger information has to be transported to all front-ends and a busy-release has to be sent back to the central trigger system. Data transfer needs a high bandwidth to transfer big amounts of data in a streaming-like mode, while control and monitoring data relies neither on latency nor on bandwidth.

The different data streams on the network are shown in fig. 3, details can be found in the following subsections.

In order to merge the different data types, a network protocol (Trigger and Readout Board Network - TrbNet) based on three virtual channels has been developed (up to 16 channels are foreseen for future extensions). To guarantee the latency requirement, all data transported on the network is divided into packets of 80 Bit each. Taking into account the standard link data rate of 2 GBit/s this gives the possibility to switch between two channels within less than 100 ns.

All channels carry an inherent priority to give triggers and

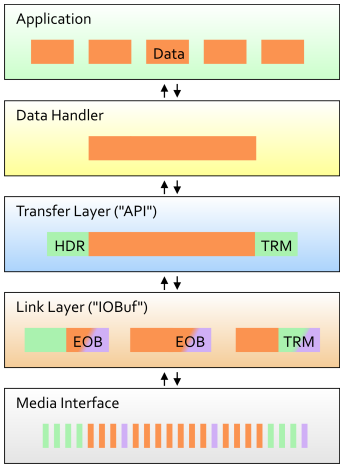

Fig. 5. The network interface is structured in five logical layers from the application layer to the media layer.

data precedence over slow control accesses. Thus, even in case of an active data transfer, a trigger information packet can be transported with almost no additional latency using the same link. In that case, the data transfer is stalled for about 100 ns and can be continued afterwards without data loss.

Each 80-bit packet starts with a 16 bit header containing the packet type and the channel number. The packet type can either be a data packet, transporting a payload of 64 bit or can be one of several network control types.

The data transfer protocol also includes a handshake mechanism on each point-to-point link. At the end of a data block (typically up to 1 kByte of data) a marker word (End of Block - EOB) is added and the data transfer is stopped on this channel until the receiver acknowledges the receipt of the data block as shown in fig. 4. This handshake prevents buffer overflows. In addition, the EOB encodes a checksum over the full data block to determine if data has been corrupted during transport.

B. Network Layers

The structure of a network endpoint can be described in a model similar to the Open Systems Interconnect (OSI) network layer structure as shown in fig. 5.

The application provides data to be sent over the network, typically organized in 32 bit words. This data is taken by the channel specific data handler, packed into TrbNet packets and equipped with necessary information, e.g. event information on the data channel.

The transfer layer adds a header (HDR) containing network addresses of both sender and receiver, the length of the message and information about the type of data transferred. The end of the transfer is signalled by a termination word (TRM) which encodes basic status information about the sender and also contains a checksum.

The link layer sub-divides data into blocks and adds the necessary handshake information to prevent buffer overflows. Additionally, checksums are calculated and checked when receiving data.

The lowest level is the physical media layer, establishing and controlling the physical connection between two FPGAs. On this level, further error corrections takes place. Each 80 bit network packet is secured by a 8 bit CRC checksum. In case a mismatch or any other transmission error on the 8b/10b encoded level is detected, a automatic request for retransmission is sent to the opposite side of the link. Hence, corrupted data is discarded and transmission is started again from the indicated position.

The media interface is the only place where data from different channels is merged into one data stream. A multiplexer between the link layer and the media layer combines and splits the data streams to keep the interdependence of all channels as small as possible.

C. Trigger and Readout Process

Sending a trigger to the detector front-ends is a three-step process. Based on analog signals from all timing relevant detectors, a trigger decision is made by the CTS. This decision is then distributed to all front-ends. The distribution is done using a dedicated differential signal acting as a low-jitter (less than 20 ps) timing reference. It arrives within 500 ns after the event took place, stops the analog-to-digital conversion and starts the read-out.

This timing signal is followed by a digital trigger packet sent over the optical network containing information about the trigger type, the trigger number and further control information for the readout process. This packet arrives at all network nodes, including all front-ends, within less than 5 μs. This is required, since the handling of data in the front-ends (e.g. zero supression) depends on the trigger type. Now all front-ends can store the acquired data in internal buffers. Afterwards a “busy release” packet is sent back to the CTS to confirm the successful trigger and to report the ability to accept the next trigger.

All front-ends are able to store about 300 average sized events before read-out has to take place. This gives the data acquisition system the opportunity to handle large variations in beam intensity on time scales shorter than 1 ms.

The read-out process for each event is again triggered by the CTS by sending a request to all front-ends. The latter forward their data to the next network hub. Here, data from several modules are combined to build one data stream as shown in fig. 6.

After combining up to 12 distinct data streams into one data block, it is forwarded either to the next network hub or fed into the so-called sub-event builder. This logic block inside the network hub packs the detector dependend data stream into a common data structure. This sub-event is then transported to the event builders (EB) in a server farm, using standard Gigabit Ethernet. Each event builder collects about 40 sub-events from the hubs, and merges them into one complete event which is sent further to the data storage system.

D. Slow Control

One crucial feature of the whole DAQ network is the ability to control and monitor every single board individually. Therefore, a versatile slow control functionality has been included. Each FPGA in the network can be identified using an

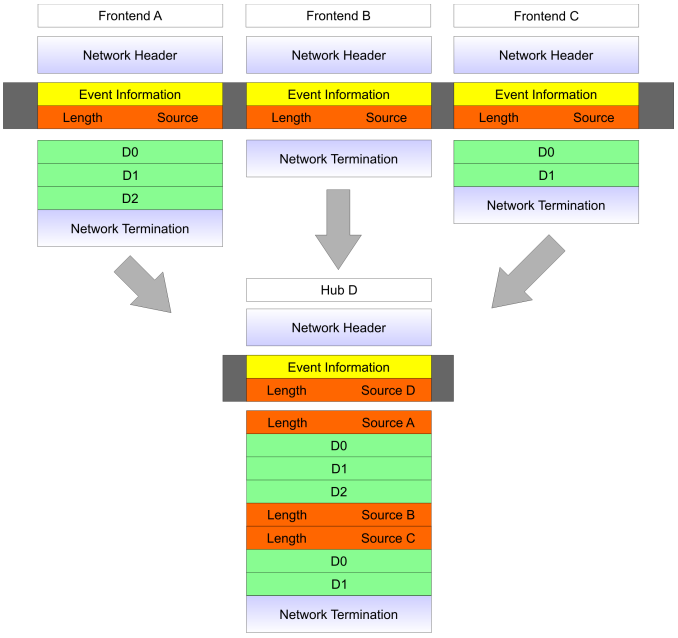

Fig. 6. Data Transport and Combining. Data from three front-ends, each equipped with event-, length- and source information are sent to a hub. There, all data streams are combined into one. Afterwards the data is forwarded either to the next hub layer or sent to the event builders (data servers) via Gigabit Ethernet using a UDP protocol.

on-board ID-chip. Based on this ID a DHCP-like mechanism is used to assign a 16 bit network address to each FPGA. This address is used to address this chip as well as to mark the source of event data. Additionally, several broadcast addresses are defined to access all boards or a subgroup of boards simultaneously.

All FPGAs provide a 16 bit internal address bus that allows to connect any kind of register, memory or higher level logic. On this data port all control, status and debugging registers are connected. To further simplify monitoring the system, in all FPGAs a standard set of registers is implemented that contain all relevant status information of the sub-system.

Most boards provide the option to upgrade the FPGA design through the optical network. The design is stored in a flash chip which can be programmed remotely. Afterwards, a reboot of the FPGA can be triggered. Some boards also offer a dual boot feature to prevent the risk of a memory corruption when upgrading the ROM.

Additionally, SPI and I2C interfaces allow to access configuration ports of the front-end electronics such as ADCs.

The computer interface for slow-control purposes is formed by a software library which gives the possibility to add software modules without coping with details of the network protocol. The library features both local access from the embedded Linux system on the TRB as well as remote, Ethernet based access from any other computer.

VI. Event Building And Data Storage

Data collected in the detector is sent from the FPGA network to the server farm (event builders). For this purpose, the TrbNet hubs are equipped with Gigabit Ethernet functionality.

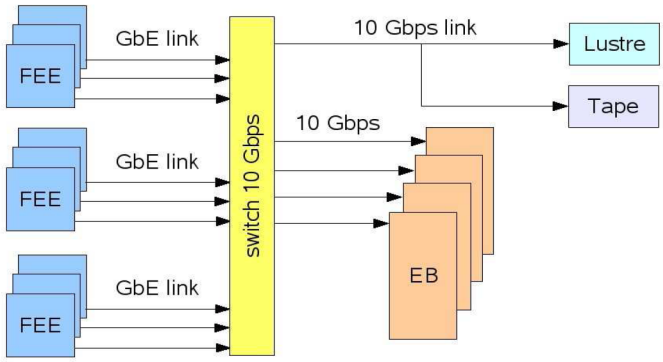

Fig. 7. The event building system. A 10 Gbps ethernet switch connects the detector, the servers and the storage system

To achieve the high data rates of above 250 MByte per second, a 10 GBit/s Ethernet infrastructure has been installed as shown in fig. 7.

Such data volumnes can not be handled by a single server. Therefore, data is dynamically distributed to a farm of servers, reducing both data rate and work load per server. Each server is equipped with 12 CPU cores, 44 TB hard disks, 64 GB memory and two 10 GBit/s ethernet links.

After collecting data from all parts of the detector, it is forwarded to the data storage. Here, several schemes are possible: data can be stored locally on a 160 TB hard disk array, or can be sent to a distributed file system (Lustre) housed in the GSI computing center, or can be written to tape for long-time storage.

VII. Performance

As of October 2010, the upgraded HADES read-out system has been fully installed and tested in two comissioning beamtimes. Performance measurements have shown that the data acquisition network meets all requirements:

Most optical links in the data acquisition network run at 2 GBit/s link speed. The maximal available bandwidth is 150 MByte/s per link due to protocol overhead and 8b/10b encoding. The Gigabit Ethernet interfaces can send up to 50 MByte/s each to the event builders (with pipelining the Ethernet frame generator 100 MByte/s would be possible).

The complete data acquisition system has a theoretical total throughput on GbE links of 1.3 GByte/s. Nevertheless, the total ethernet bandwidth is limited by the number of installed network switches and software capabilities on the server side to about 700 MByte/s.

The latency to transport trigger information from the central trigger system to the front-ends has been measured to be in the range between 1.9 μs in case of the RICH detector and 4.7 μs for the MDC system. This descrepancy is explained by an additional level of network hubs as well as the slower optical links used to connect the MDC front-ends.

The rate capability has also been greatly improved. The RICH detector is now able to take data at a continuous rate of 103 kHz with small events. In calibration mode, the data rate is limited by the ethernet bandwidth available for RICH

of 150 MByte/s. Nonetheless, the detector can run this mode at an event rate of 1.25 kHz which is a speed-up by a factor of 25 compared to the old data acquisition system.

VIII. Conclusion

We reported about the new HADES DAQ system which is currently being commissioned. The measured network performance meets or exceeds the requirements.

The approach to connect all sub-systems with one common network protocol, one common electronics platform and using common building blocks wherever possible has been proven to be very successful. Although all sub-systems have been developed in distinct test setups, integration into the complete system could be accomplished in a very short time.

The Trbnet network protocol will also be used in the prototype of the micro vertex detector (MVD) for CBM and the PANDA-DIRC detector. The modular and configurable structure of all TrbNet components allows to easily adapt to the requirements of other experimental setups.

References

[1] I. Tserruya, "Electromagnetic Probes,", WIS-IT-2009-1, Mar. 2009, arXiv:0903.0415 [nucl-ex].

[2] R. Rapp, J. Wambach and H. van Hees, "The Chiral Restoration Transition of QCD and Low Mass Dileptons,", Feb. 2009, arXiv:0901.3289 [hep-ph].

[3] G. Agakishiev *et al.* , "The High-Acceptance Dielectron Spectrometer HADES," *Eur. Phys. J. A*, vol. 41, pp. 243, Aug. 2009 [arXiv:0902.3478 [nucl-ex]].

[4] G. Agakishiev *et al.*, "Study of dielectron production in C+C collisions at 1 AGeV," *Phys. Lett. B*, vol. 663, pp. 43, May 2008, [arXiv:0711.4281]; G. Agakichiev *et al.*, "Dielectron production in C-12 + C-12 collisions at 2-AGeV with HADES," *Phys. Rev. Lett.* vol. 98, 052302, Feb. 2007, [arXiv:nucl-ex/0608031].

[5] G. Agakishiev *et al.*, "Origin of the low-mass electron pair excess in light nucleus-nucleus collisions," *Phys. Lett. B*, vol. 690, pp. 118-122, Jun. 2010, arXiv:0910.5875 [nucl-ex].

[6] G. Agakishiev *et al.*, "Phi decay: a relevant source for K- production at SIS energies?,", *Phys. Rev. C*, vol. 80, pp. 025209, Aug. 2009, arXiv:0902.3487 [nucl-ex]; G. Agakishiev *et al.*, "Deep sub-threshold Ξ^- production in Ar+KCl reactions at 1.76A GeV," *Phys. Rev. Lett.*, vol. 103, 132301, Sep. 2009, [arXiv:0907.3582 [nucl-ex]]; G. Agakishiev *et al.*, "Lambda-p femtoscopy in collisions of Ar+KCl at 1.76 AGeV,", *Phys. Rev. C*, vol. 82, pp. 021901, Aug. 2010, arXiv:1004.2328 [nucl-ex]; G. Agakishiev *et al.*, "In-Medium Effects on K^0 Mesons in Relativistic Heavy-Ion Collisions," arXiv:1004.3881 [nucl-ex].

[7] D. Belver *et al.*, "The Hades Rpc Inner Tof Wall," *Nucl. Instrum. Meth. A* vol. 602, pp. 687, May 2009.

[8] V. Friese, "The CBM experiment at GSI/FAIR," *Nucl. Phys. A*, vol. 774, pp. 377, Aug. 2006; P. Senger *et al.*, "Cbm At Fair,", in *Proc. Crit. Point of Onset and Deconfinement 2006,* PoS, CPOD2006, 018 2006; V. Friese, "The CBM experiment at FAIR," in *Proc. Crit. Point of Onset and Deconfinement 2007,* PoS, CPOD07, 056, 2007.

[9] I. Fröhlich *et al.* (HADES), "A General Purpose Trigger and Readout Board for HADES and FAIR-Experiments," *IEEE Trans. Nucl. Sci.*, vol. 55, pp. 59, Feb. 2008.

[10] M. Palka *et al.*, "The new data acquisition system for the HADES experiment," in *Proc. IEEE Nucl. Sci. Symposium (NSS), Dresden, Germany, 18-25 Oct 2008.*, pp. 1398 - 1404, Oct. 2008.

[11] J. Christiansen, HPTDC, Digital Microelec. Group, CERN, http://tdc.web.cern.ch/tdc/hptdc/docs/hptdc_manual_ver2.2.pdf

[12] A. Tarantola *et al.*, "The Upgrade of the Multiwire Drift Chamber Readout of the HADES Experiment at GSI", in *Proc. IEEE Nucl. Sci. Symposium (NSS), Dresden, Germany, 18-25 Oct 2008.*, pp. 2146 - 2149, Oct. 2008.