**THE EUROPEAN
PHYSICAL JOURNAL A**

AGATA phase 2 advancements in front-end electronics

J. Collado[1], S. Capra[2,3], A. Pullia[2,3], N. Karkour[4], Ch. Houarner[5], V. Gonzalez[1], G. Wittwer[5], A. Boujrad[5], M. Kogimtzis[6], J. Lawson[6], A. Goasduff[7], O. Stezowski[8], Ch. Bonnin[9], L. Charles[9], V. Alaphilippe[4], N. Dosme[4], C. Esnault[4], L. Gibelin[4], X. Lafay[4], E. Legay[4], D. Linget[4], M. Cahoreau[4], D. Sidler[10], G. Vinther-Jørgensen[11], J.V. Civera[12], G. Duchêne[9], E. Clément[5], I. Lazarus[6], A. Gadea[12,a]

1 Departamento de Ingeniería Electrónica, Universitat de Valencia, Burjassot, 46100 Valencia, Spain
2 INFN Sezione di Milano, 20133 Milan, Italy
3 Dipartimento di Fisica, Universitá di Milano, 20133 Milan, Italy
4 Université Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France
5 Grand Accélérateur National d'Ions Lourds, CEA/DRF-CNRS/IN2P3, 14076 Caen cedex 5, France
6 STFC, Daresbury Laboratory, Daresbury WA4 4AD, UK
7 INFN Laboratori Nazionali di Legnaro, 35020 Legnaro, Italy
8 Université de Lyon 1, CNRS/IN2P3, UMR5822, IP2I, 69622 Villeurbanne Cedex, France
9 Université de Strasbourg, CNRS, IPHC UMR 7178, F-67000 Strasbourg, France
10 Department of Computer Science, ETH Zürich, 8092 Zürich, Switzerland
11 Technical University of Denmark, 2800 Kgs. Lyngby, Denmark
12 Instituto de Fśica Corpuscular, CSIC-University of Valencia, 46980 Valencia, Spain

Abstract The AGATA collaboration has a long-standing leadership in the development of front-end electronics for high resolution γ-ray spectroscopy using large volume high purity germanium detectors. For two decades, the AGATA collaboration has been developing state-of-the-art digital electronics processing with high resolution sampling ADC, high-speed signal transfer and fast readout to a high throughput computing (HTC) farm for on-line pulse shape analysis. The collaboration is presently addressing the next challenge of equipping a 4π array with more than 6000 channels in high resolution mode, generating approximately 10 MHz of total trigger requests, coupled to a large variety of complementary instruments. A next generation of front-end electronics, presently under design, is based on industrial products (System on Module FPGA's), has higher integration and lower power consumption. In this contribution, the conceptual design of the new electronics is presented. The results of the very first tests of the pre-production electronics are presented as well as future perspectives.

2 electronics aims to replace both the phase 0 electronics described in Ref. [1], designed in 2005 and decommissioned in 2022 for the second LNL campaign, and the current phase 1 electronics designed from 2010 and used since 2015, developed in collaboration with the GALILEO INFN project [2,3]. The phase 1 development was motivated by the costs of the phase 0 system extrapolated beyond the initial phase of AGATA, the component obsolescence and the onset of the RoHS (lead-free) initiative.

We aim to have a unique version of the front-end electronics for the full phase 2 of the AGATA project. The new system will solve the long-standing issues with the optical cabling for large number of detectors and will reduce as much as possible the customised parts, using industrial products while keeping the performance and the requirements listed in the AGATA specifications [1]. The present contribution reviews the design as well as the advanced concepts introduced in this new electronics.

1 Introduction

The AGATA collaboration is currently upgrading its front-end electronics for the phase 2 of the project. This phase

2 Conceptual design

The AGATA array is operating large volume, highly segmented HPGe detectors with ultimately 180 crystals equipped with digital electronics for excellent energy resolution, good time resolution, good position resolution for each interaction

a e-mail: gadea@ific.uv.es (corresponding author)

Springer

and high counting rate capabilities. As described in [1], the fundamental ingredient to achieve this performance is 100 MHz 14bits ADC sampling of the 36 segments and the two gains (7 MeV and 20 MeV range) core analogue signals, that are pre-processed and then the information is readout, together with 1 μs of each signal (including the baseline and the rise of the pre-amplifier signals), and sent to the HTC farm. The HTC farm performs the on-line Pulse Shape Analysis (PSA) of all signals to extract the interaction points of the γ ray (in the crystal reference frame) with a comparison algorithm using a reference position-signal data base [4,5]. The energies as well as a digital constant fraction timing, with respect to the time stamps, are calculated by the pre-processing board and recorded together with the 1 μs traces.

At the front-end electronics level, the specifications as defined in the project definition [1] remain unchanged for phase 2. It is expected to run at a maximum of 50 kHz per crystal level 0 trigger rate, i.e. the Digital Constant Fraction Discriminator (dCFD) or Leading Edge discriminator (LE), using a 100 MHz clock distribution for time stamping, and a user defined level 1 electronic hardware trigger, to reduce by a factor 10 the single detector acquisition rate, based on multiplicity and coincidence conditions. Finally, it must make use of high readout bandwidths to the HTC farm for PSA, event building between AGATA crystals, event merging with complementary instruments, and tracking algorithms before disk storage. In phase 0 and phase 1, the clock synchronisation for the complementary detectors was achieved by the AGAVA VME board and the trigger decision was performed using the global trigger system (GTS) protocol to transfer the time information to a trigger processor [6]. The AGATA specifications for phase 2 include a significant effort in integration, reliability, long term availability of components and more flexibility in the data flow distribution. The previous phases of AGATA made use of dedicated FADC connected by optical fibres to pre-processing boards (ATCA [1] and GGP [2,3]) readout by PCIe [7] to a unique workstation per crystal.

The motivation to design a new electronics, beyond the existing phase 0 and phase 1 ones, have multiple aspects that are summarized in the following points:

– In the earlier versions of the AGATA electronics the samples transfer between digitiser and pre-processing electronics was done with several (7 for phase 0 and 4 for Phase 1) 60–100 m long 12-fold optical fiber cables. This was introducing fragility in a moving set-up as AGATA.
– Obsolete critical components in the previous versions of the pre-processing electronics will not be present.
– The introduction of commercial mezzanines as main building blocks, improves production yield, development time and maintenance.

– Avoiding customized point-to point data transfer after pre-processing, avoids integration of customized parts in commercial servers and allows flexibility in the architecture of the Data Flow and processing server farms improving the rate capabilities and quality of the data processing to achieve a better resolving power.
– Improvements in the data transmission rates will allow, in the future, the implementation of complex software triggers in addition to the hardware GTS and SMART Trigger Processor to increase the acquisition rate and selectivity to reach the most rare events.
– Pre-processing FPGA resources, presently used at < 50% level (except for the memory blocks, used at ≈ 80% level), allowing improvements in the pre-processing algorithms (see Sect. 10).
– The new Digi-Opt12 sampling board include improvements in the signal-to-noise quality and remove obsolete fundamental components.
– Extended monitoring for the input samples and for the pre-processing signals.

3 AGATA phase 2 electronics architecture

The AGATA pre-amplifiers designed for Core [8] signal and segments [9] have a differential signal output that is distributed to the digitiser by MDR-26 (camera link Mini Delta Ribbon) cables [10] produced by HTK-HONDA. Presently, all signals corresponding to a capsule are connected to a full digitiser with 7 MDR-26 cables, managing the full crystal signal with base bandwidth of ≈ 30 MHz. These elements remain in the phase 2 infrastructure. For phase 2, the basic concept of the first stage remains with the separation of the analogue to digital conversion and the pre-processing in two dedicated hardware units, even if they are located together. The analogue part will remain unchanged, possibly with upgrades on the warm pre-amplifier part to enhance the reliability.

A schematic drawing of the AGATA Phase 2 electronics is shown in Fig. 1. It is very relevant to notice that the architecture of this electronics, excluding the sampling ADC boards and the aggregation hardware, relies on commercial System-on-Module (SOM) products. The signals, for the core and segment contacts of the AGATA encapsulated detectors [11], are sampled using four Digi-Opt12 ADC boards. While the Digi-Opt12 board was designed for the previous AGATA phase, it is now being upgraded in several aspects. For the readout of the Digi-Opt12 the optical link, used in the previous phase, is being replaced by high bandwidth copper connections, that drive the signals to the new pre-processing board. The detector signals are not filtered at the level of the sampling, thus, there is a continuous flow of data from each one of the segment and core digitiser channels to the

pre-processing electronics, in correspondence with the 100 Msps sampling frequency and the 2 bytes sample size.

The PACE-CAP pre-processing board, composed of a motherboard and a SOM module with the pre-processing FPGA, is located near the sampling boards inside the same enclosure. The part of the PACE-CAP motherboard receiving the sampled data includes the hardware of the time multiplexing, a new concept described in later sections, that optimises the use of the transceivers but requires a particular de-serialisation (SERDES) [12] in the pre-processing SOM. The energy deposited in the segment and core contacts is determined already in the newly designed pre-processing board called PACE-CAP, since it is advisable to use as large integration windows (in the Moving Window Deconvolution (MWD) [13] algorithm used by AGATA), as the counting rate allows, to maximise the resolution, to account for possible ballistic deficits in the charge collection. Generally windows from 10 (High resolution mode) down to 2.5 μs (high counting rate mode) are being used, to be compared with the 1 μs length traces that are transfer to the PSA stage. The time stamp distribution will remain initially the GTS [6] and will be upgraded, by mid phase 2, to a new system, called SMART (SFP connectivity and MicroTCA for Advanced Remote Trigger), based also on System-on-Module (SOM) modules and the μ-TCA standard supported in the long term by the community (by CNRS). This will reduce risk, R&D time and costs.

Another new fundamental aspect is the change in readout dispatching. The Phase 1 and Phase 0 electronics used a point-to-point data link based on the PCIe standard [7] between the pre-processing board and the workstation, in phase 2, one or more pre-processing boards can distribute the readout to one workstation via a fast Ethernet standard protocol like UDP. Removing the hardware link between the pre-processing board and the data processing using HTC opens many opportunities in load balancing, specific tasks or events dispatching, the use of different technologies for the PSA with mixed CPU and GPU processing with either conventional de-convolution process or artificial intelligence approaches. Ultimately, such an approach coupled to the high bandwidth readout from PACE-CAP via STARE (the newly developed Ethernet read-out board), see following sections, could lead to a pure software level 1 trigger allowing more flexibility and selectivity in the triggered data flow. The new protocol will allow the information technology (IT) server infrastructure to be reduced, whilst having higher band-width data transport.

The slow control is performed via IPbus [14] through STARE. Additionally there is the power supply unit (PACE-PSU), able to distribute the needed voltages and currents to all the boards and the cooling and enclosure mechanics.

In the following sections, each component of the electronic chain are described in detail.
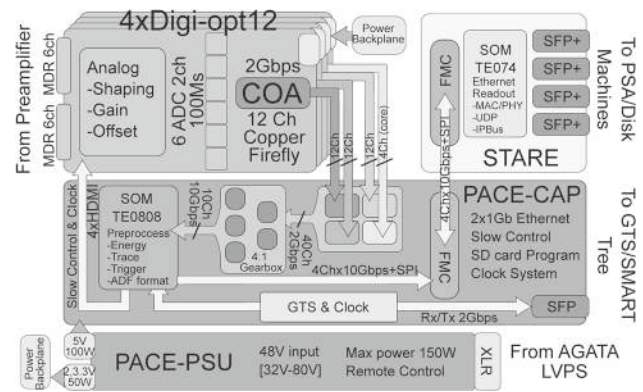
Fig. 1 Scheme of the AGATA Phase 2 electronics architecture. The detector signals arrive via MDR-26 cables to the Digi-Opt12 digitiser boards, the 40 sampled signals are readout at 2 Gbps via Firefly cables to the CAP pre-processing board and aggregated using time multiplexing into 10 lines at 8 Gbps, that are sent to the pre-processing TE0808 SOM. The accepted event buffers are then transfers via the FMC connector [15] to the STARE board to be sent to the processing computer farm. All boards are supplied with the necessary voltages and currents by the PACE-PSU (Power Supply Unit). The physical distribution of the boards is discussed in Sect. 7

4 Digitiser DIGI-OPT12

The pre-amplifiers signals are captured with custom digitiser cards designated "Digi-Opt12" (see Fig. 2), configured in 12-channel size for segment electrode readout, or 4-channel size for core electrode readout. The latter card reads in parallel the "core" pre-amplifier signals as provided in two different keV-to-V sensitivities for noise optimisation in a wide energy range, and two control signals for optional Time-Over-Threshold analysis [8]. The core-configured Digi-Opt12 card provides also a fast analogue output signal that can be used to build an AGATA trigger when coupling to other detectors. Each AGATA encapsulated detector is sampled by a set of digitiser cards, consisting of three segment-configured Digi-Opt12 cards and one core-configured Digi-Opt12 card.

4.1 Digi-Opt12 Core and Segment Board

The Digi-Opt12 cards where originally based on the ADC1413D dual-channel, 14-bit, 100MS/s flash ADC (Analog to Digital Converter) chip by NXP semiconductors, but this chip was officially declared end-of-life by the manufacturer, and thus a development has been started, and is currently ongoing, to adapt the Digi-Opt12 design to next-generation differential flash ADCs.

After investigating different solutions available on the market, the ADC32J44 ADC from Texas instruments was chosen. This chip, available in different frequency options, is fully compatible with the JESD204A protocol used by the AGATA pre-processing electronics cards. A much stricter

Fig. 2 Photograph of the new Digi-Opt12 v. 3.7 card with two ADC32J44 chips installed

range for the common-mode voltage of the input differential signal required a careful re-design of the conditioning electronics in the analogue signal path. Namely the common-mode voltage-range is as narrow as $0.95\,\mathrm{V} \pm 25\,\mathrm{mV}$ for the new ADC where it was as wide as from 1.1 to 2 V for the old ADC. A smart re-design of the analogue input stage makes the common-mode voltage at the ADC input no longer dependent on the one of the input signal [16]. Operational amplifiers for the signal conditioning are chosen with low noise and wide bandwidth so as to not be limiting factors. Remotely controlled coarse gain setting is implemented in the card for 7 MeV or 20 MeV energy ranges. A fine gain adjustment in 0.5 dB steps is available through the built-in gain control in the ADC chip. Although the new ADC has the option to use an internal dither signal for SFDR (Spurious Free Dynamic Range) optimisation, this turns out not to be useful for nuclear spectroscopy applications and is turned off by the card initialisation script so as to increase the effective resolution in terms of SNR (Signal to Noise Ratio) and ENOB (Equivalent Number Of Bits). The preliminary value obtained for the ENOB is 11.70 bit in the 20 MeV range and 11.42 bit in the 7 MeV range, very close to the ADC specification of 11.7 bit, and a little better than that featured by the previous-generation cards. As is well known the spectroscopy measurement quality is not only affected by the ADC ENOB but also by integral (INL) and differential (DNL) non-linearities,
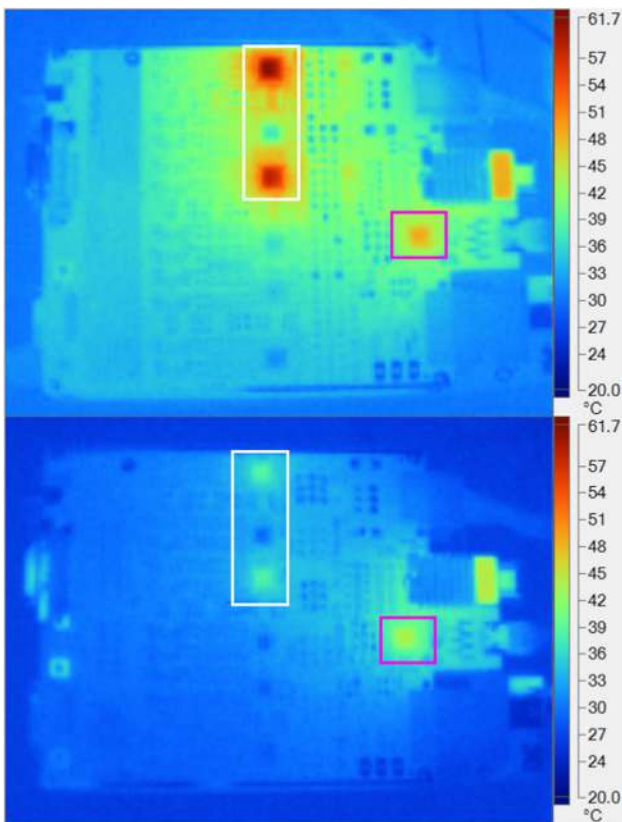
especially in case of high counting-rate. According to the datasheet, the older ADC1413D chip features a INL of ± 5 LSB and a DNL of -0.95 to $+0.95$ LSB, with typical values in the ± 0.5 LSB. The new ADC32J44 has a much better INL of ± 1.5 LSB and a better DNL of ± 0.3 LSB, which is very appropriate for the specific application. The new ADC's internal register scheme is not compatible with the previous one. For this reason a new slow-control system is being written. The essential functions have already been implemented and used for a first characterisation of the new DIGI-OPT12 cards. A smart algorithm is being developed for automatic selection of the proper scripts for the previous- and next-generation cards. This allows the co-existence of original and legacy boards in the detector array, even inside the same digitiser module. The cards, originally designed to host a Snap12 optical transceiver [17] are now wire connected to the back-end boards by a custom designed high-frequency multi ribbon cable. The distinction between "segment" cards (12 - channel) and "core" cards is maintained for full compatibility with the existing infrastructure.

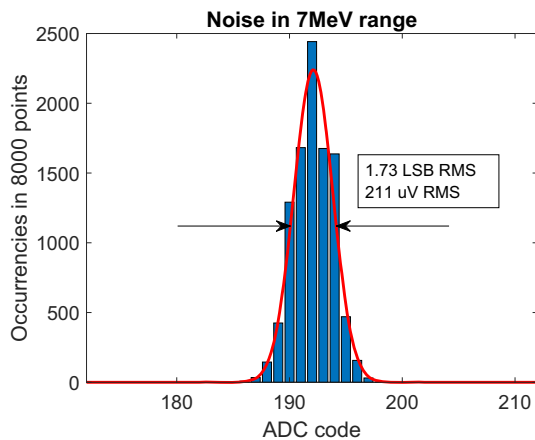### 4.2 Performance figures of the Digi-Opt12 boards

Both the old and the new A/D converters feature 2V (p-p) differential input voltage and 72.5 dB SNR (Signal-to-Noise Ratio), but the new ADC consumes much less power, i.e. 400 mW (cf. 1150 mW) @100MS/s. The reduced power consumption allows for operation with less ventilation. Nevertheless the currently used liquid cooling infrastructure will be kept in order to guarantee a high degree of thermal stability in long-term acquisitions. The first prototypes of the new Digi-Opt12 cards, as re-engineered to use the ADC32J44 chips, are currently under characterisation. Thermographic images of previous-generation and new-generation core-configured cards are shown in Fig. 3. The dramatic improvement in terms of operating temperature is evident.

A characterisation of the ADC noise has been performed acquiring 10k-points baseline signals and putting the output raw ADC codes into a histogram. The result is plotted in Fig. 4. With a fluctuation equal to 1.73 LSB r.m.s. the new ADC is demonstrated to be compliant with the strict noise specifications of AGATA. In Fig. 5, 1 MeV equivalent signals acquired with the new board prototypes are plotted. The eye diagram of the 2 Gb/s serialised signal seen at the ADC output is shown in Fig. 6.
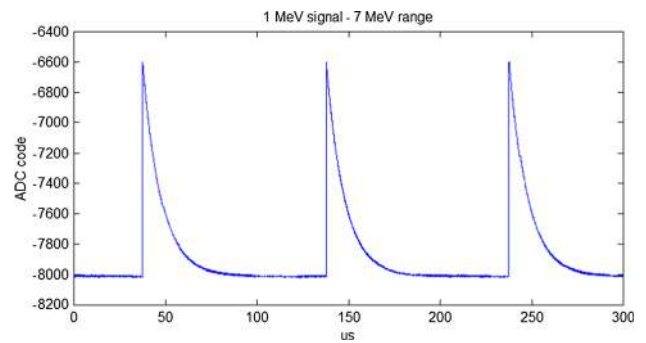
**Fig. 5** Signals equivalent to 1 MeV acquired with the Digi-Opt12 boards equipped with ADC32J44 chips

Fig. 3 Comparison of the operating temperatures of the "core" Digi-Opt12 components. The previous version is shown in the upper part and the new version is in the lower part. The new ADCs (inside the white rectangles) are even cooler than the clock distribution chip inside the magenta square
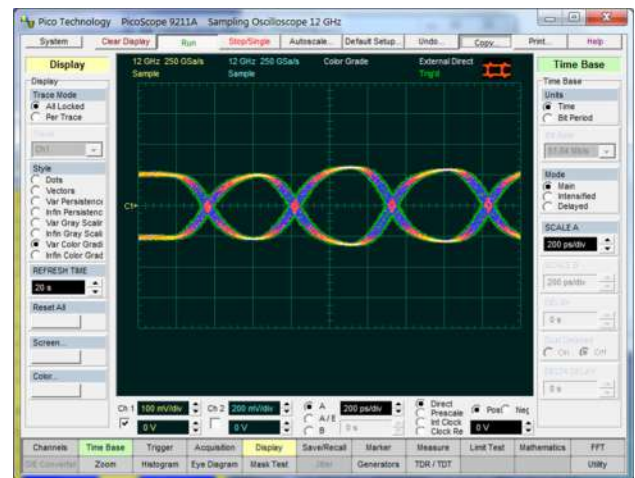
Fig. 6 Eye diagram of the high-frequency digital ADC output as captured with a PICOSCOPE 9211A

Fig. 4 Histogram plot of the baseline values in a 10k points signal

5 The pre-processing board PACE-CAP and associated firmware

The pre-processing board PACE-CAP functions as the motherboard of the system, housing the pre-processing FPGA board and the STARE board. It connects to the GTS system and manages slow control, clock distribution, and is responsible for all timing synchronization. Moreover, it handles the 4 to 1 channel aggregation procedure and manages the readout to STARE via four 10 Gbps Aurora links. This section will provide a brief overview of the PACE-CAP board's hardware, along with a description of the System on a Module (SOM) FPGA firmware and its performance.
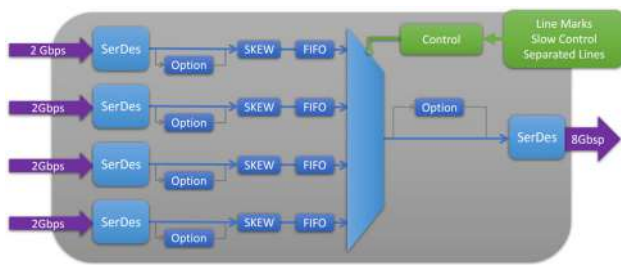
Fig. 7 Conceptual drawing of the time domain multiplexing link aggregation (IDM) concept built in the receiver section of PACE

## 5.1 Matching the bandwidth and FPGA capabilities with time multiplexing concept

The intrinsic characteristic of the digitalised data for AGATA, at the nominal 100 Msps sampling rate, requires a 2 Gbps readout link, with a total of 76 Gbps for the 38 pre-amplifier channels per segmented encapsulated detector. For the transceivers populating the modern FPGA's 2 Gbps is underrated and, furthermore, there is a strong dependency of the FPGA cost with the number of transceivers. The mid-range FPGAs with 10–20 transceivers has sufficient resources to perform the tasks of the AGATA pre-processing. Moreover, on the small-size SOM devices, the number of high speed channels is also critical due to the limited number of connectors available in the mezzanine board. These two reasons are sufficient to consider a solution to approach the real total bandwidth of this mid-range FPGAs with an AGATA digitized aggregated bandwidth. The concept selected to face it is by using a Time Domain Multiplexing link aggregation gearbox, with a 4 to 1 rate. The implementation relies on the commercial Texas Instrument TLK10000 devices. These devices are designed to work in pairs to aggregate data in one of them, send the data through a higher speed network and dis-aggregate data in the receiving second device. Our development adapted this concept to receive the high-speed link in an FPGA with an associated firmware that will recover data, splitting it on the original 4 links. On the AGATA electronics in particular, the JESD204 protocol is used, so this is the protocol use for the development [18]. To prove the feasibility of this solution a prototype called IDM was built [19] in a mezzanine form factor, to be installed in an evaluation board with the selected pre-processing FPGA SOM. The concept is presented in Fig. 7. The mezzanine was done as an independent device that is initialised automatically through a small form, low cost FPGA. The data arriving from each of the four Digi-Opt12 digitiser boards was sent through a Samtec Firefly 12-line connector [20], that allows flexibility on selecting optical or copper cables. After successful test of this proof-of-concept in 2019, the same hardware architecture was implemented, in the receiver section, in the PACE-CAP motherboard.

5.2 The PACE-CAP board design

The PACE-CAP motherboard takes the IDM in the receiver section and in addition incorporates all functionalities needed for the AGATA Phase 2 electronics. In this new board, the aggregation system is implemented with the TLK10002 [21] and the copper firefly has been selected to connect to Digi-Opt12. In order to use Firefly cables with the Digi-Opt12 boards it was necessary to develop a transceiver-size adaptor called COA to interface the Digi-Opt12 optical socket to Samtec firefly connector. After passing through the IDM section, the 10 aggregated input data lines at 10 Gbps lines are connected to the SOM board Trenz TE0808-05-BBE21-A [22] that houses a Zynq Ultrascale+ XCZU15EG-1FFVC900E. Correspondingly, 10 GTH transceivers of the FPGA de-serialise the data and the firmware recovers the 38 JESD204 channels, corresponding to the segment and core ADC data, decoded and ready to be processed. One important task of the AGATA pre-processing is the generation of the local trigger (level 0 trigger), that contributes to the global trigger, as well as the full system synchronisation. This task is taken by the GTS part of the hardware and firmware. The GTS tree has an alignment protocol to ensure that the system clock is distributed synchronously to all the ADCs and pre-processing systems of the array [23]. The PACE board is connected to the GTS tree through an SFP optical transceiver, where the global clock is received, the triggered events are sent and global trigger replies received. The link provides a common clock to all the system that is cleaned and fine-tuned during the alignment protocol in the PACE Clock & GTS hardware section. This section of the board also provides 3 independent PLL systems, one for the GTS, one for the SOM and one for the aggregator system. Following the pre-processing on the SOM FPGA, all the events validated by the GTS second level trigger, are sent for readout via the Ethernet STARE board, plugged into PACE-CAP as a mezzanine though an FMC Vita 57.1 [15] four 10 Gbps links using the Aurora protocol [24]. While the slow control commands are passed to PACE-CAP via an IPbus firmware implemented in STARE, the PACE board incorporates two small FPGA Lattice Machxo3 [25] in order to perform the local slow control, monitor and power the rest of the system. One of the small FPGAs is for the slow control and has access to all the hardware bus lines, the other one is for power supply management and start-up of all the system, that is managed autonomously. The latter includes the control of the power supply unit board, in such a way that the power of all the system can be remotely controlled. These two FPGAs have a power supply independent of the rest of the system. In addition, the board incorporates two Ethernet connections, one for the FPGA fabric, devoted to the IPbus control in case the STARE board is not present. The second one is for the multi-processor system (MOSOC) present in the Zynq Ultrascale+ FPGA and
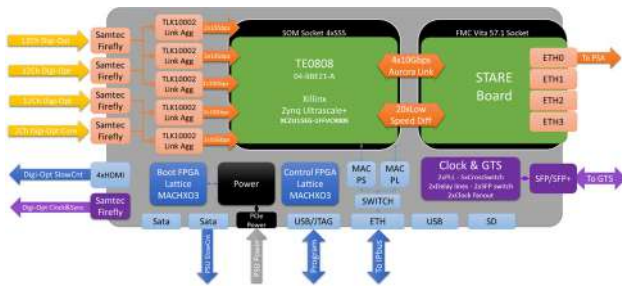
Fig. 8 The PACE pre-processing board block diagram
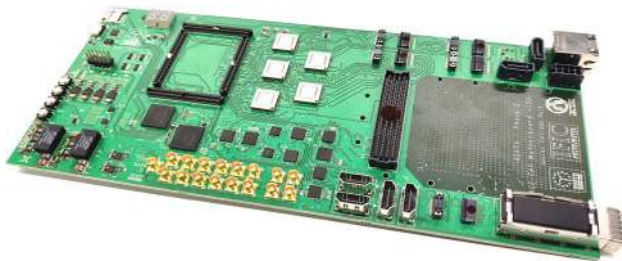
Fig. 9 Pre-production PACE board implementation (without the commercial SOM board)

is used for remote controlling the MPSOC and to perform the GTS alignment software procedure. Furthermore, there is an SD socket for the MPSOC programming, a USB-C for debugging purposes and the JTAG/UART socket. The board is manufactured in 12 layers and the design included a thorough simulation of the high-speed differential links. Figure 8 shows the PACE board layout and Fig. 9 show the picture of the pre-production version of the board.

5.3 Processing firmware

The pre-processing firmware has three main goals: (1) energy determination and trace construction, (2) local trigger generation and global trigger management and (3) data formatting and readout.

All the channels are processed in parallel by a process module called Datapath, that is identical whether it corresponds to a core channel or a segment channel and is synchronous with the 100MHz system clock from the GTS. The data path employs a digital Constant Fraction Discriminator to identify if a signal exceeds the trigger threshold. Simultaneously, it processes the Moving Window Deconvolution (MWD) trapezoidal filter, which has a maximum length of 20 μs. In the event of a triggering signal, a trigger request is sent to the rest of the PACE system involved and to the GTS, the energy is calculated from the MWD trapezoidal filter and with a linear interpolation the time is extracted using the crossing point of the dCFD within the 10 ns width of the samples. The firmware has four data store modules, i.e.:

(i) Event module: includes a complex memory that stores all the information each time a trigger request is sent from the processing module selected for triggering, normally from the core. It also communicates with the GTS system via the "GTS leaf" module and, in case the event is validated, the corresponding memory contents are sent to the readout engine in charge. In case of rejection or time-out the memory contents are deleted and the slot is free for a new event. In order to preserve the data during the request-validation/rejection/time-out cycle of the GTS, there is an array of such memories that stores up to eight events. The data sent to the readout engine is packed in standard AGATA Data Frame format (ADF) [1,26] with 100 samples per channel and event. It is possible to extend the traces sending 200 samples per channel and event. On these packages we also find the calculated energy from the trapezoidal filter of all the channels, the timestamp, the fine value of event time from the linear interpolator on the core and the type of event sent by the packager

(ii) Long traces module: Long traces are segments of the data stream used for diagnostic or response function calibration. There are two type of long traces, the standard up to 4000 samples (≈ 4μs) and the very long trace stored in the DDR memory. The first one is constantly stored for all 38 channels and can be sent, within a single ADF packet, upon request, via the slow control, to the readout engines. The second one is storing constantly the previously selected channels into DDR memory and can be sent to the readout engine in multiple packet protocol.

(iii) Spectra module: its task is accumulating the energy spectra of every channel since the last reset, called by the user via the slow control. The readout of any spectra, in a single ADF packet, can be requested via the slow control.

(iv) The monitor module: This module stores a full memory of the internal selected data. It can be sent to the readout engine in an extended multiple packet protocol.

In order to perform the readout, four readout engines have been created in the firmware, each one is connected, via one of the 10 Gbps Aurora links, to STARE, that is in charge of sending the data packages via Ethernet. These readout engines can select which of the memory blocks are reading and combine the packets into the readout data flow. As mentioned before, all packages are formatted in ADF or multiple package ADF based. It includes ADF IDLE frame, Error frame and System Off frame that are sent automatically depending on configuration and status of the system.

In order to interact with the GTS tree, validate the local triggered events and recover the system clock, there is a GTS section in the firmware. This section includes a GTS leaf that communicates with the GTS tree and with the firmware event memory section and provides the global validation and rejections, aligns the global timestamp, monitors the count-
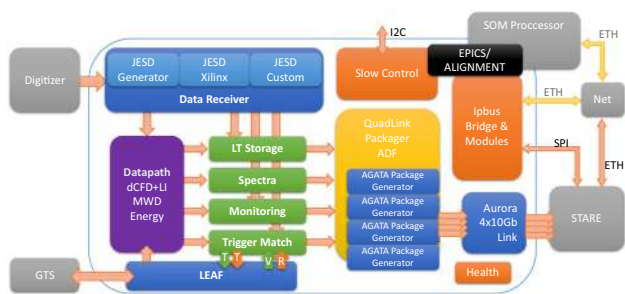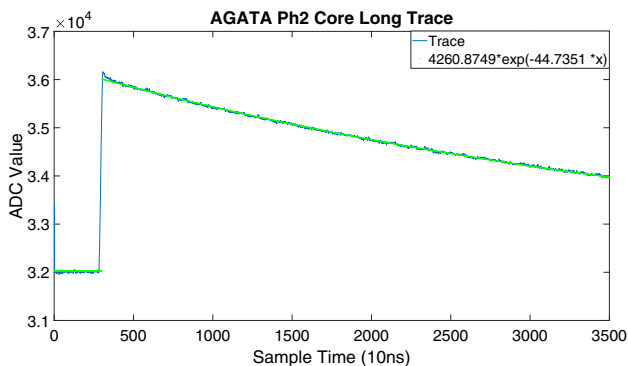
Fig. 10 The PACE firmware modules

Fig. 11 First PACE readout of a 35 μs AGATA core signal trace with the fit of the pre-amplifier response function

Fig. 12 The STARE conceptual design

6 Serial transfer and readout through ethernet STARE board: hardware and firmware

The STARE board is a FMC mezzanine mounted on the PACE-CAP pre-processing board and designed to do the readout using fast Ethernet protocol at high bandwidth. The STARE mezzanine compiles the generated data from the PACE-CAP board to a standard UDP Ethernet data format at 10 Gbps. The board is designed latency-free between the incoming data and the outgoing data toward the data acquisition workstations. In order to improve the counting rate capabilities of AGATA, or the possibility of a pure software trigger, or improving its sensitivity by using more complex algorithm to process the Pulse Shape Analysis, the STARE board should have the possibility of multiple data lanes configuration. Such requirements could not be found on the shelf at the time of the start of the upgrade and triggered a specific development described in this section.

6.1 STARE hardware design

This is the result of a long research and development program started in 2016 [27]. The processed data from the PACE-CAP are sent to the STARE over a Xilinx AXI stream AURORA interface, using up to 4 independent 10 Gbps transceiver lines. Each line is processed independently inside the STARE and the data is formatted to be transferred through 4 independent 10 Gbps transceivers to servers or computer farm using UDP protocol. Each of the 4 UDP channels can send data to several different servers. The STARE module is equipped with external memory to allow re-transmission of any packages not correctly received by the HTC server farm. Figure 12 shows the STARE conceptual design and functionality.

6.2 STARE board description

The STARE board is composed of a TE0841 SOM from Trenz Electronics [28], an FMC connector and four SFP+ 10 Gigabits transceivers. It contains 2 EEPROMs, a high

ing rates and implements the "backpressure" control. Backpressure refers to a mechanism used already in phase 0 and 1 which inhibits temporarily the trigger as soon as, for any detector, the data acquisition or data processing is unable, to consume the data from the readout hardware.

The last element of the firmware is the slow control, that is done through the IPbus implemented in STARE via a slow speed differential connection in the FMC. There is also the capability to control PACE in standalone mode by IPbus. There is an IP block on every relevant firmware module with the required set of registers. There exists a specific IPbus module that translates the communication with the slow control FPGA outside to control the Hardware. Figure 10 describes the PACE firmware components.

5.4 Performance figures

To date, the PACE board is at its last stage of development and the pre-production boards have been produced and tested at INFN-LNL over the full readout chain using an AGATA detector. In Fig. 11, the very first 35 μs ($\approx$8 kB) long trace, captured by the Long Trace Memory module and readout through Ethernet, is shown. The measured decay time from the pre-amplifier matches the expected value. The total firmware utilization is less than 50% of the device, however, a higher percentage of memory resources are allocated (80%).
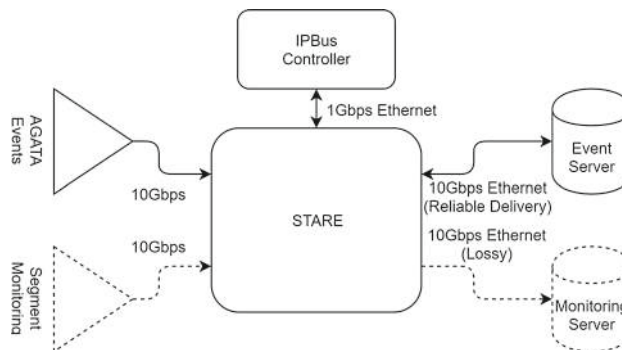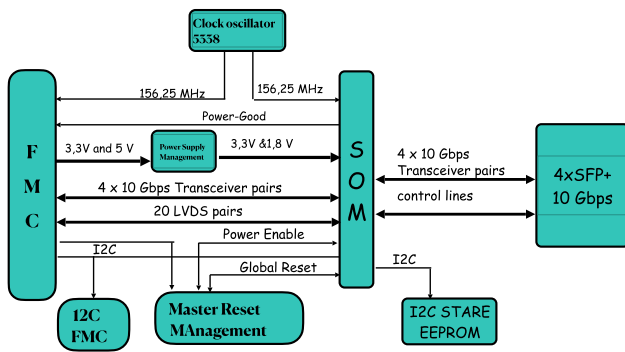
Fig. 13 STARE electrical and interconnection block diagram

Fig. 14 STARE power supply management and secured power up sequence

precision oscillator and clock generator. The STARE design respects all electrical requirements described in the FMC Vita standard [15], although it does not fully comply with the mechanical requirements because of the SOM and the 4 SFP+ connectors dimensions, which makes the height of the STARE over 10 mm, which is the requirement of the Vita FMC standard. The length of the STARE board will also be longer than the VITA recommended rules. The PCB design will have a rugged metallic border to ensure thermal conductivity since the SOM and SFP transceivers temperature might increase during full speed transmission. Figure 13 show the electrical and interconnection diagram.

6.3 STARE prototype board design

The STARE board prototype has extra functions which makes the STARE development and debugging easier. These functions are the Power supply management and the Jtag interfaces. These will not be used in the production board. The power supply management on the STARE board is very complex due to the power up sequence for the SOM board and the SFP+ transceivers. Figure 14 shows the block diagram of the STARE power supply management. The external power supply input has active multiplexers to switch automatically from either the external input or the FMC input. The STARE can be used with or without any carrier board for development necessity. Moreover, the JTAG programming will use different modes of connections (FMC JTAG, simple Xilinx HX Programmer or a JTAG SMT2 module). Figure 15 shows the STARE prototype fully populated. The STARE prototype was manufactured and a full validation tests was peformed over 400h with continuous data transfer and without error (see Sect. 6.7 for more details). In the production phase the JTAG programming and the power supply will be implemented through the FMC connector only. Moreover a test bench is designed to be used to validate the STARE production and reduce the maintenance time. The test bench is controlled by IPbus interface. A software package has been developed to make fully automatic tests of the STARE board.
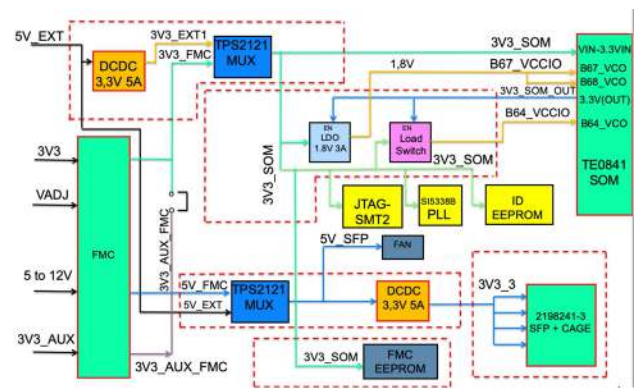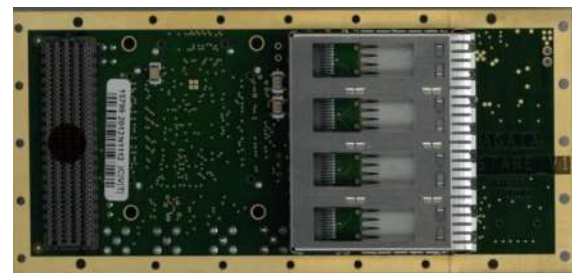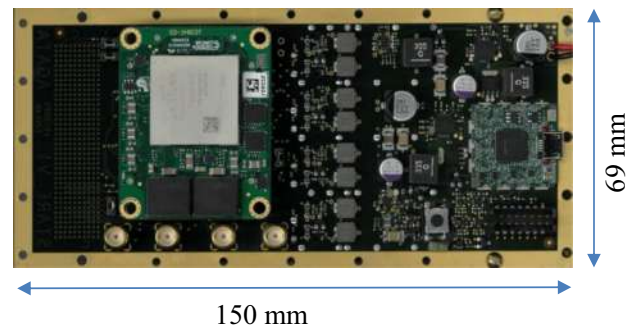
**Fig. 15** STARE Prototype: top side (top) and bottom side (bottom)

The package includes a specific programming section to configure STARE so that it can be immediately installed in the PACE-CAP board ready to be connected to the different networks and servers (Ipbus and UDP servers). Figure 16 shows the STARE test bench ready to test the STARE production boards.

6.4 STARE firmware design

The STARE firmware block diagram is shown in Fig. 17. It is divided into 5 modules: the AURORA interface, the package generator, the IPBus CORE, the UDP interface and the Data generator. Each module contains an IPBus slave providing slow control and status monitoring. The Aurora interface receives data from the PACE module through the FPGA multi-Gigabit Transceivers. The event buffer size can be con-

**Fig. 16** STARE Prototype: top side (top) and bottom side (bottom)
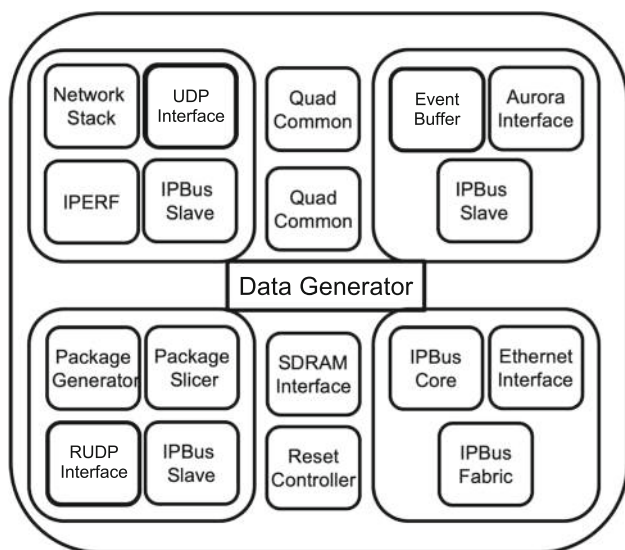
Fig. 17 The STARE firmware block diagram

trolled via Slow Control. It must be a multiple of 64 bytes. For AGATA the default event size is 8 kB. The Aurora interface stores the data into double toggle FIFOs so it can handle the 10 Gbps rate. The first Fifo stores the first event and while sending data to the package slicer the following buffer stores the event into the second FIFO. The package generator combines data buffers from either the package slicer or from the RUDP interface (in case the received data from the server is lost). The package Slicer takes the event buffer received from each FIFO and splices it into 8192 bytes size packets to be

sent over UDP interface using the Network Stack. The UDP Interface block compiles the spliced packet to be compatible with the UDP data format. It converts the data buffer into a jumbo frame of 8192 bytes, which is compatible with the AGATA event size. The UDP interface adds specific protocol data to each packet and transmits the packet to the server. The IPbus interface is composed of an IPbus core that interfaces the STARE to the IPbus server network through the Ethernet interface, and several Slave IPs which are developed independently inside each firmware block. Finally the data generator module simulates experimental data that are sent, for testing purposes, to the data acquisition servers. There are four different types of sub data generators inside the firmware. The simple 8 kB counter sub generator, used to validate the data transfer from the STARE to the computer server at 10 Gbps. Continuous data is sent at full speed and checked on the receiver side one by one and whenever an error occurs it increments a counter. See Sect. 6.7 for the error count during testing. Three more sub data generators were designed to simulate AGATA events in order to validate the DAQ software receiver which was developed for the new electronics.

6.5 Reliable UDP 10 Gbps transfer SRP (selective repeat protocol)

One of the requirements for the STARE board is that all the data received from the pre-processing must be delivered to the server. As the transport protocol used is UDP, this functionality is not guaranteed by default and must therefore be added. The local network interconnection between the STARE and the server is not a dedicated connection. A stare can be connected through a high speed switch to one or more servers, this depends on the data flow and the server capacity to hold the online data analysis. It is not a point to point connection and hence data losses can appear if the server gets busy and some data get lost. It is better to resend a lost data than having a deeper buffer receiver, because the experiment counting rates are high (50 KHz) and each event transmits today 8 kbytes, and in the future this number might increase thus creating a difficulty in the receiver buffer depth. The other widely used transport protocol TCP does provide guaranteed delivery but during the R and D phase using the iperf program on the Virtex7 evaluation board the data transfer rate cannot exceed 2 Gbps. Thus, an explanation for why UDP was chosen anyway. The UDP interface is able to utilise the full bandwidth for a single session. However, the drawback is the need of a supplementary protocol to guarantee delivery. This will slightly decrease the available bandwidth and required a development time. For the STARE project, the requirement is that all data has to be delivered and only delivered once. This means that lost data should be re-sent and data that is sent twice should be discarded. Several protocols were studied and finally the Selective Repeat Protocol [29,30] was
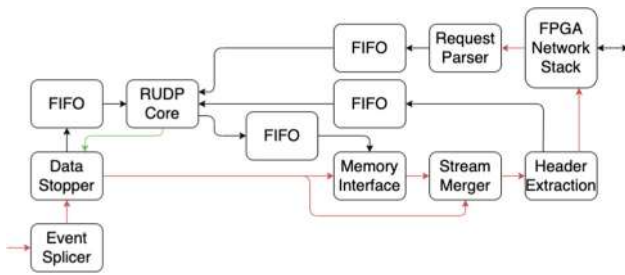
Fig. 18 Selective repeat protocol block diagram. Figure taken from Ref. [27]
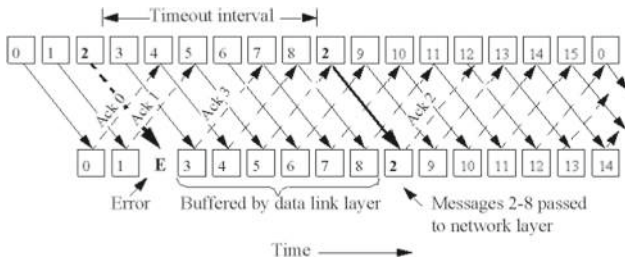
Fig. 19 Selective repeat protocol mechanism with intelligent handshake process, to minimise time shift, and synchronise quickly. Figure taken from Ref. [27]

chosen, since it has proven to be the more efficient option. This however comes at the cost of increased complexity in the firmware. The block diagram of the Selective Repeat protocol is shown in Figs. 18 and 19 shows schematically how operates [31].

When using the Selective Repeat protocol, only the lost frame will have to be sent back. This results in a better utilisation of the available bandwidth, especially if the loss rate is high. It also means that the memory interface must allow single frames to be read back for re-transmission. As each frame is treated individually, it also means that a timeout mechanism for each frame is needed. This increases the implementation complexity on the sending side. The core of the Selective Repeat protocol implementation is the RUDP Core. RUDP is short for Reliable UDP which is what the Selective Repeat protocol provides. The RUDP Core keeps track of all the active frames and generates timeouts used to re-send frames in case an acknowledge is not received. When a frame is allowed to pass through the data stopper the frame is marked as active in the RUDP Core. The next step in the frame lifecycle is for the timeout countdown to start. This happens when the frame has passed through the Header Extraction module. In this way the timeout won't be affected by frames read back from the memory delaying the frame transmission. The timer is now running. If the acknowledgement frame is received before the timeout occurs the frame is marked as inactive and the RUDP Core is finished with the frame. If the acknowledgement frame isn't received in time the frame times out. When this happens, a request is sent to
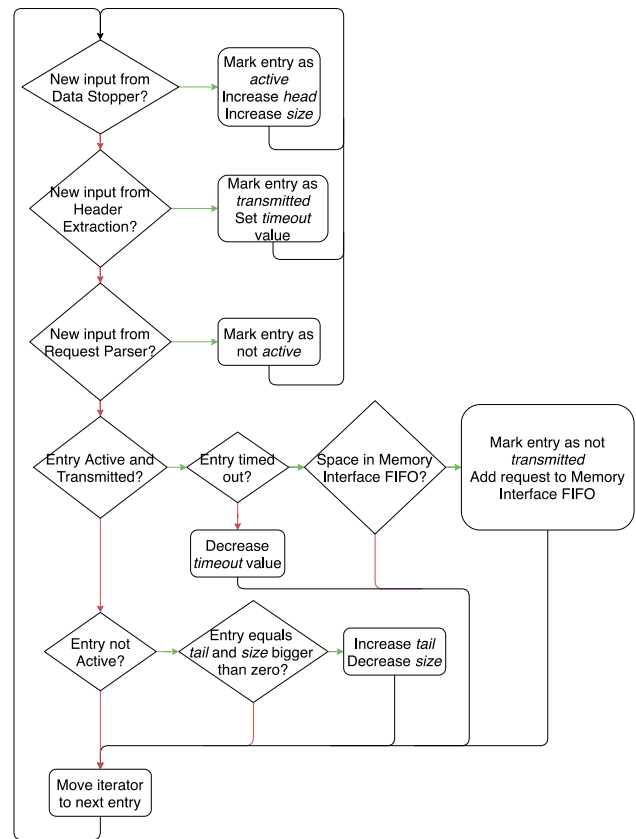

Fig. 20 RUDP core flow chart. Figure taken from Ref. [27]

the Memory Interface block triggering the frame to be read from memory. When the frame passes through the Header Extraction module the timeout is started again. This process will continue until the acknowledgement frame is received. The flow described is shown in Fig. 20. The RUDP IP implementation is not included in the current firmware it will be included as an upgrade in the future, because the priority is to deliver to the project a first version with the software developped on the receiver side. It is only developped and tested in the evaluation board. There is lot of work to be done in the future to include the RUDP.

6.6 Control and monitoring

All control and monitoring of the STARE is done by the IPbus interface. All the STARE parameters are software controlled. In addition, the STARE firmware has an automatic setup interface using IPbus and a dedicated Controller to correctly reset, initialise and make the STARE ready to receive data without any external or manual commands. All the startup procedure are automatically launched through Ipbus. The CAP IPbus interface is developed inside the STARE firmware and converts each ipbus command into an SPI format sent to the CAP board.The slow control software will have access to
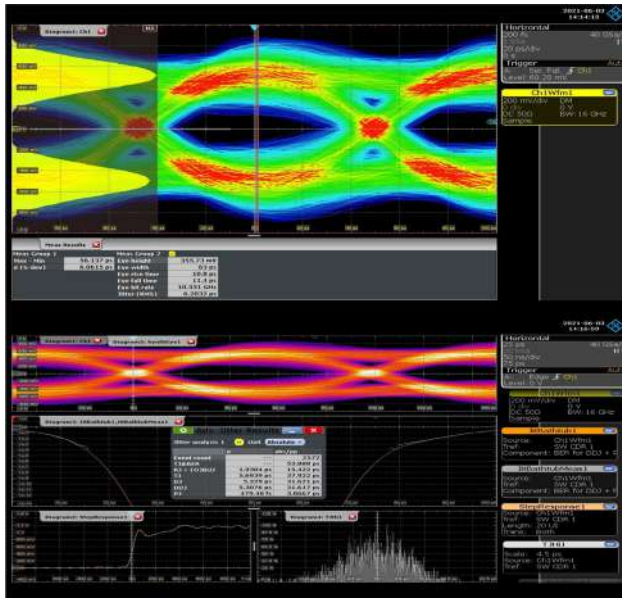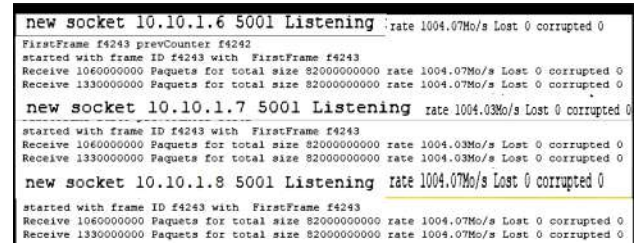
Fig. 21 STARE 10 Gbps Eye diagram measurements The top eye is the online eye display and the bottom figure is the processed eye from sampling signal

Fig. 22 Continuous 30 Gbps data transfer rate results from 1 STARE over 3 days

continuous monitoring of the STARE FPGA temperatures, and the Aurora interface synchronisation status as well as the UDP interface modules. Error status can be monitored to see if there are any transfer problems by configuring the data generators and by spying on the error counter registers or the error data registers. The SOM power supplies can be measured and monitored through SC and alarms can be activated in case of a power supply failure in order not to harm any components.

6.7 STARE prototype validation and results

The production of ten STARE prototypes has been completed and the boards have been qualified and validated using a new software readout SQM developed by the AGATA Data Flow team for the HTC server. The eye-diagram test for the high speed transfer lines was validated using the Rhode & Schwartz 16 Ghz bandwidth 60 Gbps sampling rate oscilloscope. The eye diagram in Fig. 21 shows the time for a digital signal to be logic 1 or logic 0 at a certain frequency. In the case of the STARE the signal in the eye diagram has a signal width of 50 ps and the time which is stable @ logic 1 or 0 is about 80%, which is close to the theoretical values. The 20% are for the rise time and the fall time. The measurements shows jitter < 20 ps and a transmission Bit Error Rate (BER) of 10^{-12}. Presently, the production of STARE boards is ongoing, 150 modules will be delivered within the fourth quarter of 2023.

The UDP 10 Gbps transfer rate was validated with the server receiver software. The first tests used only 1 optical

fibre connection at 10 Gbps. A full test program, synchronised with the STARE data generator allowing the receiver to know what the transmitter is sending, was designed to check the data transfer during several days. There were several long time tests which were made and one of them is chosen here as an example using three STARE fibre optic outputs simultaneously connected via Ethernet to three different servers. The result proves that the STARE can transfer about 30 Gbps in parallel (3×10 Gbps) with loses of less than 2.10^{-6}. Figure 22 shows a 3 days test made with 0 losses. 972 Tbytes were tested by the 3 servers @ 10 Gbps on each server. The RUDP firmware will be implemented after the validation of the production boards inside the AGATA electronics.

7 Mechanics and infrastructure

To maximize solid angle coverage, the AGATA structure leaves little room to place electronics close to the array, except for the very front-end pre-amplifier part that needs to work at cryogenic temperature to minimise noise, and the warm part of the pre-amplifier, that needs to be in the as close as possible, which are are located inside and outside the detector cryostat, respectively. The electronics of previous phases of AGATA have used standard buses, like ATCA [32], for the pre-processing electronics. However, for phase 2, the AGATA collaboration decided to move to a more compact and power efficient ad-hoc design rather than commercial bus solutions, at least for the customised parts. In order to host these boards, a new mechanical infrastructure was conceived using commercial racks that allows the new electronics to be placed close to the array in a more compact configuration. At the same time, the pre-processing hardware, which had previously been placed in dedicated servers, is now placed in the same enclosure as the Digi-Opt12 boards, together with a new reduced-size power supply. The complete set of electronics placed inside a closed box contains the set of four Digi-Opt12 cards, the pre-processing board (PACE-CAP) with the Ethernet readout board (STARE) and the power supply (PACE PSU). In order to dissipate the heat produced by them, aluminium cooling blocks, in contact to specific hot points, and
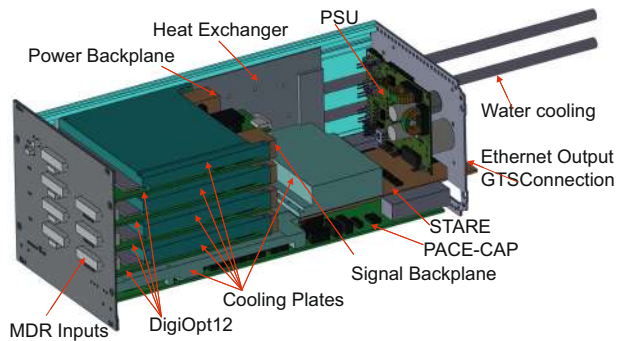
Fig. 23 Conceptual design of the AGATA phase 2 mechanics
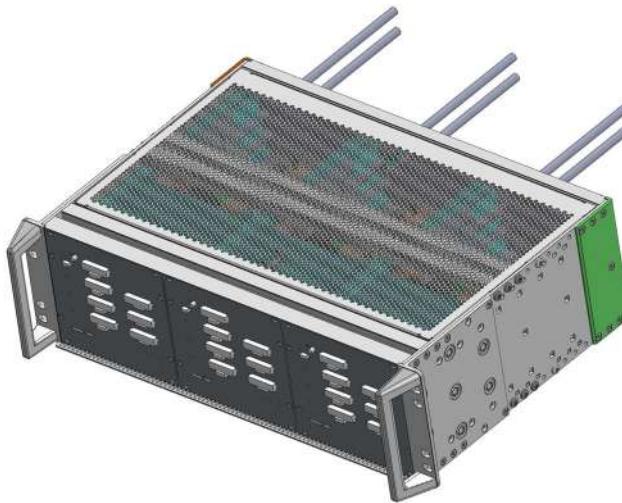
Fig. 24 Design of the crate with the AGATA phase 2 electronics for an AGATA triple cluster [11]

attached to each one of the electronic boards with thermally conductive silicon pads, are used. A heat exchanger, connected to the water cooling circuit available at the host laboratories, is used to refrigerate the aluminium blocks.

Figure 23 shows a schematic view of the new AGATA box for one AGATA crystal. As shown, two back-planes connect the four Digi-Opt12 boards providing power and connection between them. Additionally, an auxiliary connector in the signal back-plane allows for an alternative access through SPI and I2C protocols to the Digi-Opt12 boards. Digi-Opt12 board outputs connect to the PACE-CAP board using FireFly cables avoiding the use of optical transceivers. The configuration and the slow control of the Digi-Opt12 boards are achieved by the PACE-CAP board through HDMI cables.

Power for the complete system is provided by the PSU module which takes a 48 V, 3.5 Amps input and delivers + 5V to PACE-CAP using a standard PCIx connector and + 3.3V and + 2V to Digi-Opt12 boards through a connection to the power back-plane. The PACE-PSU has power monitoring features to control and regulate the DC/DC converter output voltages as well as reading thermal sensors of the sys-

tem. All these features can be accessed in a standalone way through an I2C connector or directly from PACE-CAP board using a dedicated SATA connector. A 19 in. 3U europacPRO rugged rack from nVent Schroff GmbH. allocates three boxes of electronics in a compact system Fig. 24. The XLR-type power connection, the Ethernet optical data outputs, the connection for the Global Trigger System (GTS) (see Sect. 8), the IPBus connector for the slow control (see Sect. 9) and the water pipes of each box are all accessible at the back of the rack.

8 AGATA synchronisation and trigger systems, from GTS to SMART

Since the AGATA demonstrator phase, the 100 MHz clock distribution for time stamping as well as the global trigger decision have been performed by the GTS system [1,6]. Each AGATA crystal runs independently and the pre-processing card, called *leaf* in the GTS system, uses the time-stamp and sends trigger requests for validation via the GTS tree to the ROOT mezzanine, acting simultaneously as concentrator and clock reference (see Fig. 25). The ROOT mezzanine can be connected to a Trigger Processor assigning validation or rejection responses to each *leaf* for readout based on multiplicity or coincidence user defined conditions in given time stamp windows. During the demonstration phase, the AGATA trigger processor was limited to 40 leaves including ancillary detectors and a maximum of 255 nodes (*leaf* +FIFO+ROOT). After twenty years of operation, the GTS architecture in AGATA, based on a customized protocol, has reached many limits in term of number of channels, compatibility's with modern FPGA and possible developments of the protocol. The present system can not be expanded beyond 1π (45 crystals) with a limited number of ancillaries. For phase 2, the collaboration is preparing two major upgrades. The first one is the development of a trigger processor able to manage more than 40 *leaves* using either the GTS or the new protocol. The second development is a new system using modern FPGA's and hardware, much more integrated, with a new standardized package protocol for trigger and time stamping called SMART. The fundamental concepts of GTS remain with advantages of a lighter protocol compatible with modern FPGA's and expandable to other electronic system used by complementary detectors. In this section these two upgrades are described.

8.1 The 255 IP GTS trigger processor module

The GTS processor module, later referred to as TP, is a triggering module based on the Xilinx Virtex-7 VC707 development board. It gathers all the incoming trigger requests from the GTS tree into partitions, then decides which ones to
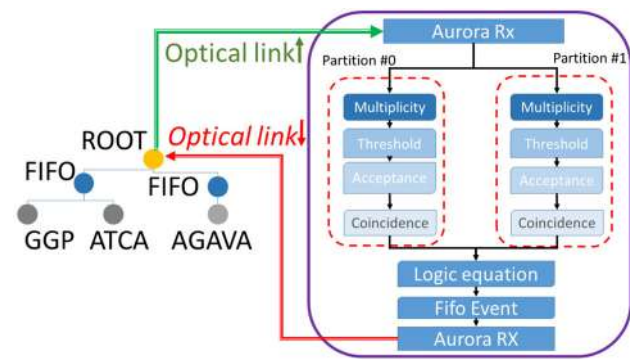
Fig. 25 The GTS Tree overview and Trigger Processor flowchart. The gray dots represent the *leaf* mezzanine generating requests. The blue dots represent the GTS FIFO aggregator mezzanines. The yellow dot represents the ROOT mezzanine. All these mezzanines are connected by optical link in blue
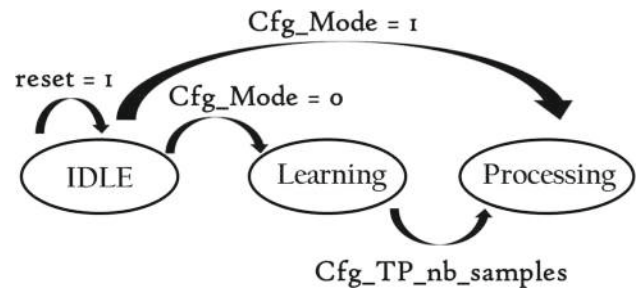

Fig. 26 `TS_48b_reference` state machine

Fig. 27 Diagram showing the assignment of the channels to one or more of the eight partitions available

accept or reject depending on a user defined configuration. The fundamental I/O and protocols used are identical to the phase 1 board [1].

The trigger processor is connected at the top of the GTS Tree after the ROOT mezzanine (see Fig. 25). The new TP handles up to 256 channels or *Leaf*. The 256 channels can be dispatched over 8 partitions. Each channel can be assigned to one or several partitions. Multiplicity threshold within a partition and coincidence between partitions based on their relative TS difference are configurable. The general architecture is based on Aurora Rx input GTS frames dispatched in the 8 partitions. After processing as described in Fig. 25, output messages are returned to leaves through Aurora TX. The GTS Trigger Processor platform is configured by software using the IPBus protocol.

In the following subsections, the detailed design is reviewed.

8.1.1 The 48 bit timestamp reference generator (`TS_48b_reference`)

This module generates a local 48-bit timestamp with a 10.00 ns resolution using the 100 MHz GTS global clock. It can be configured to start in *manual* or *automatic* mode. The manual mode is selected through the `cfg_Mode` signal (`cfg_Mode = 1`) and the `TS_48b_reference` starts after reset of its internal timestamp counter with the value given by the `cfg_delay` register set by the user. In automatic mode, selected through the `cfg_Mode` signal (`cfg_Mode = 0`), the `TS_48b_reference` starts after reset with a learning phase during which it tries to synchronise with the global 48-bit timestamp of the GTS tree (see Fig. 26). For that purpose, all the GTS leaves that are part of the GTS tree should be up and running during the GTS TP learning phase. Presently, the learning phase stops when the TP has collected 100 requests of all defined leaves. It
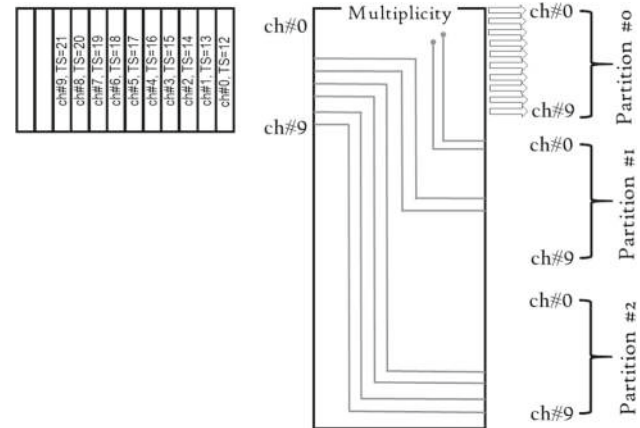
selects the shortest time stamp as local reference. This procedure guarantees that the system collects the oldest request possible from the complete system.

8.1.2 The multiplicity filter trigger

The multiplicity module processes incoming trigger requests timestamps, calculates the partition multiplicity depending on channels assignment and compares it with the partition threshold defined by the user. If the threshold condition is not met, the received trigger requests are rejected. Otherwise, the acceptance and coincidence windows are generated for the concerned partition. As mentioned in the GTS trigger processor architecture overview section, the current version of the platform handles up to 8 partitions which means that any channel out of the 256 supported by the GTS communication protocol can be assigned to one or more partitions.

Figure 27 show an example of this channel assignment feature. As a reminder, the channel assignment is a static configuration defined by the user during software setup. Changing channel assignment supposes a software platform reset. If the multiplicity threshold condition is met, an acceptance window is generated allowing late trigger requests part of the same event to be accepted.
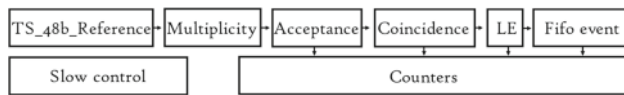
Fig. 28 Detailed view of the Trigger Processor global flowchart

8.1.3 The coincidence window trigger

If an acceptance window is generated for a specific partition, then a coincidence window is generated for the concerned partition. This coincidence window is generated with a certain width and delay defined by the user through software.

8.1.4 The logic equation

The logic equation module analyses coincidence windows generated and compares them with the logic equation configured by the user. If the condition set by the user is met, then all the pending trigger requests are accepted and sent back through the GTS tree to the GTS leaves. Otherwise, they are rejected.

8.1.5 Fifo event

The Fifo event module stores incoming GTS requests through the Aurora Rx interfaces, then generates the GTS validate/reject response requests. It works as follows: as soon as a trigger request is received on the Aurora Rx interface, it is stored into the fifo event and an analysis window is opened. The TP_GUI software defines the analysis window width through the CFG_TIMEOUT register with respect to the following requirement.

During the analysis window, if the logic equation condition is met, then all the stored requests are validated and sent back to the GTS tree through the Aurora Tx interface. If the logic equation is not met during the analysis window, the fifo event content is flushed and reject responses are generated.

The global flowchart is shown in Fig. 28.

8.2 The SMART protocol

SMART is the newly developed trigger and clock distribution to be implemented in the second half of the Phase 2 of AGATA. It keeps the functionality of the previous GTS system using a new and lighter, but robust protocol, adapted to small packets of data, and makes use of an industrial standard for the hardware. When completed the SMART system will replace the GTS system for AGATA without changing the pre-processing infrastructure. The global architecture remains the same with a high level *tree* gathering the local requests from the channels (*leaf* or *end − point*) via *ROUTER's* to a *HUB* providing a Time Stamp in all cases. When the trigger feature is used, the trigger decision is sent
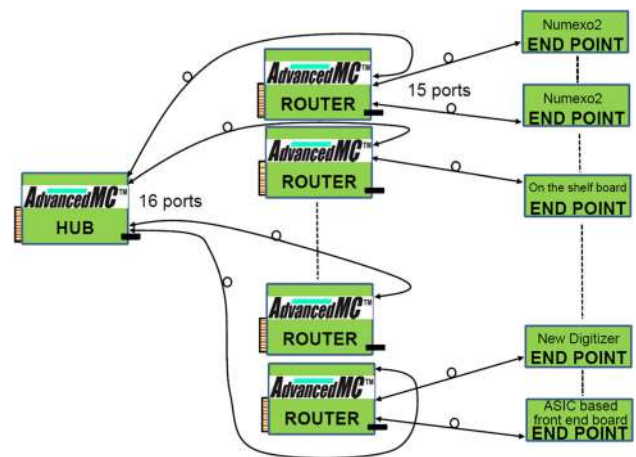

Fig. 29 The proposed SMART architecture for time stamping and trigger decision

back (routed via messages) to the *leaves* similarly to the GTS system. The basic features of the 256 IP Trigger processor described in the previous section will be re-used in a more modern SOM module. SMART will provide a 48 bits/10 ns time stamp and a 32 bit event number using a unique 100 MHz clock for synchronisation. The transmission between the HUB and the ROUTER is 4 Gb/s with a Payload Data Rate of 400 MB/s. The transmission between the ROUTER (ex-FIFO) and the digitisers will be 2 Gbps with 8 B/10 B encoding towards Digitisers and boards to synchronise with a Payload Data Rate of 200 MB/s. The connectivity and data transmission can be SFP or QSFP connectors up to 7 ms with passive Direct Attach Cables characterised at 10 Gbps or optical transceivers, fibre and/or copper cords up to 150 ms with OM3/OM4 fibres. The general layout of the SMART system is presented in Fig. 29. As shown in the figure, a key feature is that the END_POINT firmware can be adapted to a large variety of digitiser or pre-processing boards. The feature makes SMART a possible universal system for clock AND trigger distribution for heterogeneous front-end electronics.

The chosen standard for the SMART hardware is the μTCA for the crates with double/full size AMC form factor (carrier) in "dual star" shelves. For the FPGA and processing hardware for the HUB (former ROOT) and ROUTER's, an integration of "System On Module"durable industrial solution, reducing costs, engineering time and design errors is proposed. The architecture is built with Xilinx Zynq FPGA (ARM processing/16 Multi Gigabit Transceivers). The detailed description of the SMART AMC board is shown in Fig. 30.

With 15 ports (3 in the case of GTS) per ROUTER, up to 240 Digitisers or Leaves can be synchronised by 17 AMC's housed in 2 μTCA shelves which is equivalent to 120 former GTS V3 mezzanine. Upgrading to the latest Kintex UltraSCALE+ FPGA (Xilinx KU19P) will lead to a sub-
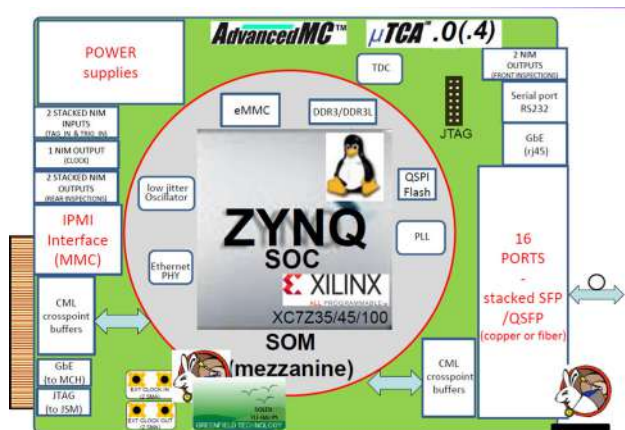
Fig. 30 The SMART module architecture

stantial increase of the incoming channels, this new HUB called "SMART MCH" will give the possibility to connect 32 ROUTER boards. In this case, up to 480 End Points (Digitisers) could be synchronised.

8.3 Interfacing AGATA with other detectors using SMART

The coupling of AGATA to complementary detectors has played a major role in spectroscopic investigations in the last decade. AGATA was coupled to the PRISMA or VAMOS++ magnetic spectrometer [33,34], large volume scintillators [35–37], heavy-ions detectors like DANTE, LYCCA or TRACE [38,39], fast-timing array [40], the neutron detectors NEDA [41], the charged particle arrays DIAMANT [42] and MUGAST [43] and LEPS detectors. So far, the time stamping and trigger validation was performed using the GTS system (see Sect. 2 and introduction to Sect. 8). The AGAVA VME board [1] ensured the distribution of the AGATA current time stamp value to the complementary detectors data flow and sent trigger request from the complementary detectors, allowing them to participate in the global trigger decisions, to reduce the AGATA readout rate and increase sensitivity. In the next AGATA phase, the SMART system will be used to implement a similar concept. AGATA will operate at different facilities coupled to local instruments but and to detectors, constructed by other collaborations, having specific front-electronic.

An upgraded version of the AGAVA board for SMART is foreseen and, since it is not expected anymore, the interaction will be mainly with instrumentation based on the VME standard readout and it will be necessary to incorporate more flexibility. This is related in particular for instruments that use digital sampling electronics. It is expected that the direct implementation of the SMART protocol into the complementary instruments front-end electronics will be used, as in the case of NEDA.

9 Control and monitoring

As already mentioned in previous Sects. 5.2, 5.3, 6.3, the IPBus [44] flexible Ethernet-based control system is being used for the slow control and monitoring of the new boards. Specific software is under development to help users to deal with up to 180 cards for the AGATA 4π configuration. A general scheme is displayed in Fig. 31. The core of the architecture is a python server (blue box in Fig. 31) that will run on a Debian Linux platform. The python language was chosen to take advantage of existing code used to develop and test STARE boards prototypes. The server uses the end-user python API for IPBus, provided by the uhal library, for the transactions with the boards (red line in Fig. 31). The Django framework has been selected to generate the web site and manage interactions with a MySQL database used to store necessary information about users, commands and parameter values. Two user interfaces are provided. The most user oriented one consists of an interactive web interface, running in current web browsers, programmed using classical web languages HTML, CSS and Javascript. The software will allow users to perform configurations and actions applied to individual or global sets of STARE boards and present views of the 180 crystals. Direct access to the server through a python console is also available for more expert users. Finally, a REST (REpresentational State Transfer) API is provided to allow software interoperability: this is indeed required to catch actions sent by other user interfaces such as, for instance, the run control which should deliver at least start/stop signals to the cards. For the monitoring of the boards, several systems are to be set up. From the user interfaces, it will be possible to register which specific parameter is to be monitored and at which frequency. Once registered, the python server will automatically request, at the selected period, the parameter and push it in an InfluxDB time series database. Users can then visualise the time evolution of the sampled parameter through Grafana web based dashboards. Control spectra can be built also directly inside the cards. Thanks to the versatility of the ADF frame definition, such spectra can be pushed, if required and needed, into the standard UDP data stream (green lines in Fig. 31). Once they have reached a processing node of the workflow, they can be part of the global monitoring system of the DAQ box (see XREF Advances in Software Developments). A third way to monitor the card is to be realised using another UDP, 10 Gbits/s, output channel of the STARE board. Raw data streams are sent through the python server to a web embedded, oscilloscope like, application able to visualise simultaneously up to 10 channels.
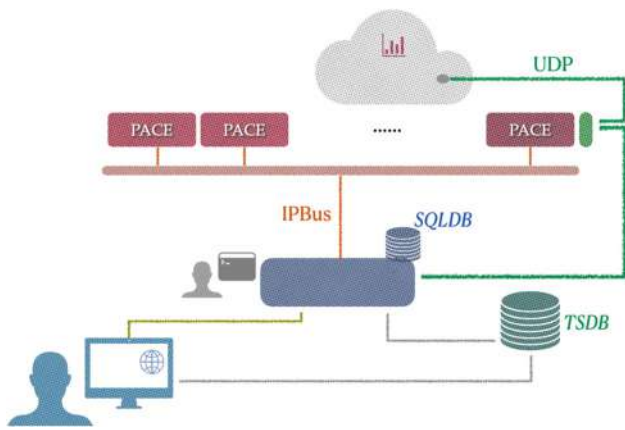
Fig. 31 General slow control and monitoring scheme. The blue box in the center is the python server, core of the whole system. Attached to the server there is an SQL database (SQLDB) to store all the parameters required. Transactions with the PACE boards are realized using the IPBus protocol (red lines). Two user interfaces are shown. The blue one corresponds to the web browser oriented interface while the grey one is for a direct console attached to the server. Monitoring can be performed from the time series database (TSDB) using Grafana dashboards, through in-card produced spectra sent embedded in the UDP data stream to the DAQ box (grey cloud) or through raw data streams feeding an oscilloscope like application
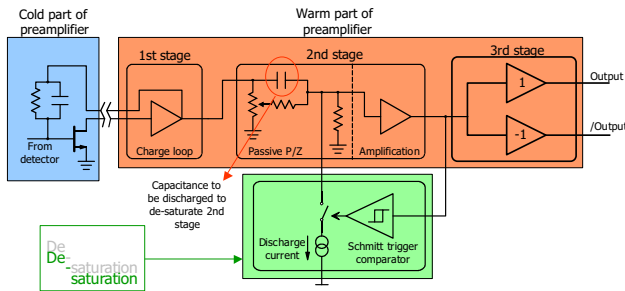
**Fig. 32** Agata pre-amplifier block diagram [45]

10 Outlook: improvements in the pre-processing firmware of AGATA

Several areas were identified for possible firmware improvements, the pre-amplifier pulse reconstruction, the high rate performance, particularly in the case of pileup and the single/multiple interactions identification. When we examine the pre-amplifier block diagram in Fig. 32, it is evident that the transfer function to be expected is of a higher order, due to the 1st (charge loop) and 2nd (Passive P/Z) stages.

We have conducted a preliminary analysis of captured AGATA pre-amplifier output waveforms, focusing on the trailing edge of the pulse. Currently for pulse reconstruction in AGATA, the MWD algorithm used [13] assumes a first-order pre-amplifier transfer function. The offline model that is used to fit the pre-amplifier output is $Ae^{-t/\tau} + B$ (Fig. 33) and is actually a higher order model(1 zero/2 poles) to get a more accurate fit. However, even with this model,
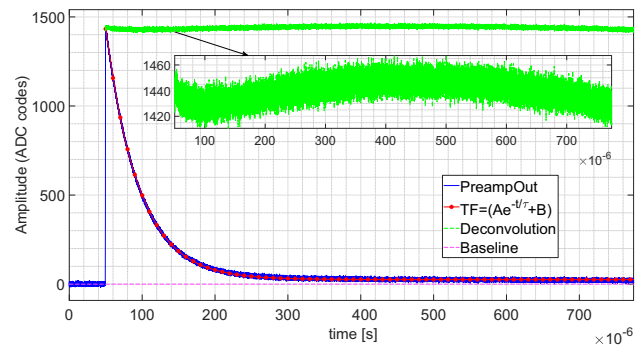
Fig. 33 Pre-amplifier output (dataset 01A/Segment3.dat) shown with its fit ($Ae^{-t/\tau} + B$) and deconvolution (green trace). NRMSE = 97.50%
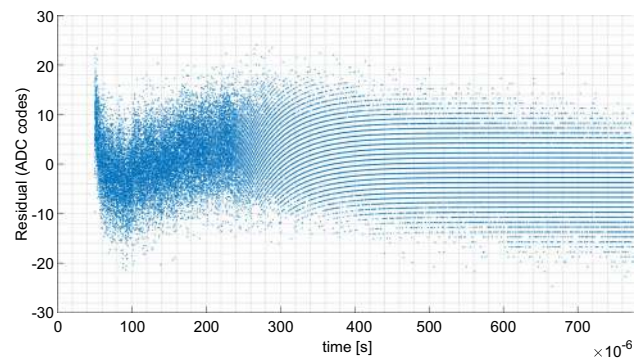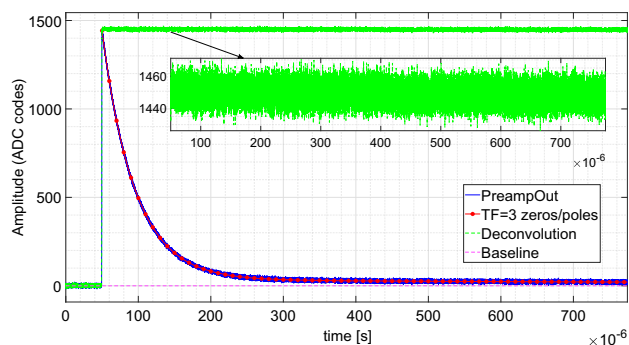

Fig. 34 Residual of the pre-amplifier output (dataset 01A/Segment3.dat) and its fit ($Ae^{-t/\tau} + B$). NRMSE = 97.50%. A pattern is visible at the left hand-side of the plot due to fit mismatches

a pattern is still visible in the residual plot of some of the data sets (Fig. 34). In order to reduce this error, the current practice is to constrain the range of the fit to the start of the decay. The normalised root mean square error (NRMSE) for the analysed data set is 97.50%.

The recent work of Födisch et al. [46] has extended pulse reconstruction to higher order pre-amplifier transfer functions via use of IIR (infinite impulse response) filters that can be implemented in an FPGA. Based on this approach, the transfer function is calculated off-line, using the MATLAB system identification tool. The parameters of this function are then applied to an IIR filter that implements the inverse transfer function, in order to recover the input step to the pre-amplifier.

Our analysis indicates that the pre-amplifier pulse response is better modelled, overall, by a 2nd or 3rd order transfer function (2 or 3 zeros/poles). Figure 35 shows an example of a 3rd order transfer function. There is no visible pattern when inspecting the residual plot (Fig. 36) and as a result the recovered step is flat (Fig. 35, green trace). The pattern with gaps to the right-hand side of the plot (Fig. 36) is due to the ADC quantisation levels. The NRMSE is 97.75%.

Fig. 35 Pre-amplifier output (dataset 01A/Segment3.dat) shown with its calculated higher order transfer function (3 zeros/poles) and deconvolution via an IIR filter (green trace). NRMSE = 97.75%
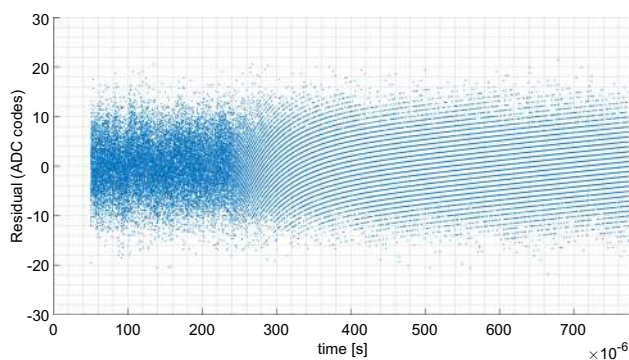
Fig. 36 Residual of the pre-amplifier output (dataset 01A/Segment3.dat) and its calculated higher order transfer function (3 zeros/poles). NRMSE = 97.75%. There is no visible pattern to the left-hand side of the plot, indicating a tighter fit

The next stage of the analysis would be to collect data that contains multiple events per pre-amplifier in order to find out whether the filter coefficients can be further optimised and to verify the models. Since the expected output of the proposed approach is a staircase, further analysis is required to establish whether the algorithms currently used for the trapezoidal filter and the baseline correction are still applicable with this new approach.

Other firmware development areas would be, first the one aiming to process pileup events at higher rates, which are currently ignored. The pre-amplifier output could be delayed (pipelined) in order for the energy processing module to find the time between pulses [47,48]. Once this is known, a dynamic adjustment of the parameters of the trapezoidal filter can take place, which should allow the processing of pulses at higher rates, albeit with lower energy resolution. The second one would aim to distinguish between single and multiple interactions, providing a flag in the FPGA that can be used by the PSA stage [47]. This could be done by analysing the rising edge of the event and looking for changes in the slope that point to collection of charge from different locations in

the detector. This could serve as an interim solution until the multiple interaction analysis is addressed by the PSA itself.

The above firmware upgrades are planned to be implemented and the FPGA resources required for the IIR filter, as reported in [46], in combination with the FPGA used in the PACE-CAP board and its current utilisation, suggest that there are adequate FPGA resources to accommodate the proposed upgrades.

11 Outlook beyond phase 2: digital preamplifier and ASIC technologies

The AGATA collaboration is pursuing electronics R&D towards new concepts and technologies beyond the present Phase 2. The idea, already mentioned in the early AGATA conceptual design, is to bring closer to the HPGe diode the signal pre-amplification and its numerical conversion. An ambitious extension of this idea would be to locate the signal pre-amplification in a high resolution cryogenic ASIC inside the cold part of the detector coupled to a sampling ADC placed in the warm part of the cryostat, in the location presently housing the warm pre-amplifiers. Such development can be only implemented in AGATA together with a redesign of the detector module cryostat, considered as well in our R&D programme beyond phase 2.

Such development requires, firstly, a cryogenic ASIC pre-amplifier development. The interconnection between the germanium segments and the front-end electronics is extremely sensitive to noise. For this reason their length must be minimised. The ideal place to put the new integrated front-end electronics is where the cold FETs are located, around the cold finger next to the capsule contacts. This is the reason why conventional room-temperature solutions are not viable. The goal is thus to realise a monolithic multi-channel low-noise low-power pre-amplifier for high-resolution gamma-ray spectroscopy in the typical energy range of 20 MeV, able to work at 77 K with no sacrifice in terms of performance as compared with current warm pre-amplifiers.

The wiring between the integrated electronics and the warm connectors can transmit heat and warm up both the ASICs and the detector. This is a critical point since it is necessary to keep the HPGe crystal temperature at around 83 K. The ASICs will be connected through wires made of constantan (alloy of copper and nickel) in order to minimise the thermal contacts between the cold detector and the warm part of the cryostat. Such wires have the drawback of having a non-negligible resistance. Special care must be taken in the ASIC design to ensure stability even when powered with such non-ideal supply rails.

The currently used "fast-reset" technology is to be implemented in the ASIC pre-amplifier, which prevents the system from getting paralysed, due to signal saturation, in case

an energetic ion hits the detector. Such technology swiftly resets [49] the pre-amplifier when an over-range event occurs, which reduces the system dead time from several ms to a few μs.

Another innovative aspect of the development is the possible elimination of the bulky discrete 1 GΩ feedback resistor, which would substantially improve system miniaturisation and possibly enhance the energy resolution at long shaping times (6–10 μs). This would not only reduce the white thermal parallel noise but also avoid signal shape distortions due to non-ideal resistor impedance and distributed-line effects [50,51]. The possibility to fine-tune the geometry of the input MOS transistor may allow for an optimised low-noise solution with best capacitive matching. Nevertheless moving the pre-amplifiers from the warm side to the cold side of the detector cryostat is quite challenging. Indeed, the design of a circuit to be operated at liquid nitrogen temperature is critical, since commercial integrated technologies are not validated at cryogenic temperatures.

A new cryogenic ASIC development will have consequences in the design of the analogue to digital conversion, drivers and power supplies. The advantages that a cryogenic front-end ASIC could achieve for the signal acquisition of the detector signals may be lost by an unsuitable ADC device (which needs to be placed inside the cryostat but outside of the cold section) with intolerable noise levels, degrading signal quality, or high-power consumption which will affect the temperature environment. Currently, long cables usually connect the front-end electronics to the ADC boards, which can be replaced once the ADC is placed close to the signal source. However, carefully designed line drivers are needed in order to not decrease the signal-to-noise ratio when sending the digitized signals to the next stage of data acquisition.

Data Availability Statement This manuscript has no associated data or the data will not be deposited. [Authors' comment: This manuscript describes instrumental developments and has no associated data.]

References

1. S. Akkoyun et al., Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip. **668**, 26 (2012). https://doi.org/10.1016/j.nima.2011.11.081
2. A. Goasduff et al., Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip. **1015**, 165753 (2021). https://doi.org/10.1016/j.nima.2021.165753
3. D. Barrientos et al., IEEE Trans. Nucl. Sci. **62**(6), 3134 (2015). https://doi.org/10.1109/TNS.2015.2480243
4. B. Bruyneel, P. Reiter, G. Pascovici, Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip. **569**(3), 764 (2006). https://doi.org/10.1016/j.nima.2006.08.130
5. B. Bruyneel, P. Reiter, G. Pascovici, Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip. **569**(3), 774 (2006). https://doi.org/10.1016/j.nima.2006.08.129
6. M. Bellato et al., J. Instrum. **8**(07), P07003 (2013). https://doi.org/10.1088/1748-0221/8/07/P07003
7. Pci-sig standard peripheral component i/o data transfers specifications. https://pcisig.com/specifications
8. G. Pascovici, A. Pullia, F. Zocca, B. Bruyneel, D. Bazzacco, WSEAS Trans. Circuit Syst. **7**(6), 470–481 (2008)
9. G. Pascovici, A. Pullia, F. Zocca, B. Bruyneel, D. Bazzacco, WSEAS Trans. Circuit Syst. **7**(6), 470–481 (2008)
10. Camera link specification—v2.1. https://www.automate.org/a3-content/vision-standards-camera-link
11. A. Wiens et al., Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip. **618**(1), 223 (2010). https://doi.org/10.1016/j.nima.2010.02.102
12. D. Lewis, in *Proceedings of the DesignCon* (2004)
13. A. Georgiev, W. Gast, IEEE Trans. Nucl. Sci. **40**(4), 770 (1993). https://doi.org/10.1109/23.256659
14. C.G. Larrea, K. Harder, D. Newbold, D. Sankey, A. Rose, A. Thea, T. Williams, J. Instrum. **10**(02), C02019 (2015)
15. ANSI/VITA (VMEbus International Trade Association), *FPGA Mezzanine Card (FMC) connector standard* (2008)
16. A. Pullia, 18th AGATA week and the 2nd Position Sensitive Germanium Detectors (PSeGe) technology and application Workshop, talk 45, slide 6 (2017). https://indico.in2p3.fr/event/14457/
17. Smiths interconnect SNAP12 homepage. https://www.smithsinterconnect.com/products/optical-transceivers/embedded-transceivers/snap12/
18. J. Collado, V. Gonzalez, A. Gadea, Eur. Phys. J. Web Conf. **253**, 01006 (2021). https://doi.org/10.1051/epjconf/202125301006
19. Collado Ruiz, Javier Ph.D. thesis, University of Valencia, Spain (2020). https://roderic.uv.es/handle/10550/73461
20. Firefly^{TM} low profile micro flyover system^{TM} cable assembly. https://www.samtec.com/products/ecue
21. Texas instruments tlk10002 dual-channel 10-gbps multi-rate transceiver. https://www.ti.com/product/TLK10002
22. Te0808 ultrasom+ mpsoc-modul with zynq ultrascale+ xczu15eg-1ffvc900e technical reference manual. https://wiki.trenz-electronic.de/display/PD/TE0808+TRM

23. M. Bellato et al., J. Instrum. **8**, P07003 (2013). https://doi.org/10.1088/1748-0221/8/07/P07003
24. Aurora specifications: Free logicore™ ip design enabling the use of multi-gigabit transceivers for xilinx fpga. https://www.xilinx.com/support/documentation/ip_documentation/aurora_8b10b_protocol_spec_sp002.pdf
25. Lattice machxo3l-2100. https://www.latticesemi.com/en/Products/FPGAandCPLD/MachXO3
26. O.S. et al., Eur. Phys. J. (2023) (**In this issu**e)
27. N. Karkour et al., IEEE Trans. Nucl. Sci. **68**(8), 2005 (2021). https://doi.org/10.1109/TNS.2021.3086684
28. The trenz electronic te0841 technical reference manual. https://wiki.trenz-electronic.de/display/PD/TE0841+TRM
29. R. Benice, A. Frey, IEEE Trans. Commun. Technol. **12**(4), 135 (1964)
30. D.L. Lu, J.F. Chang, IEEE Trans. Commun. **37**(3), 245 (1989)
31. G. Vinther-Jørgensen, Internship report 2 (2019)
32. The advanced telecom computing architecture (advancedtca®) specifications. https://www.picmg.org/wp-content/uploads/PICMG_3_0_Shortform.pdf
33. A. Stefanini, et al., Nucl. Phys. A **701**(120134), 217 (2002). https://doi.org/10.1016/S0375-9474(01)01578-0 (**5th International Conference on Radioactive Nuclear Beams**)
34. M. Rejmund et al., Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip. **646**(1), 184 (2011). https://doi.org/10.1016/j.nima.2011.05.007
35. A. Maj, J. Gaardhøje, A. Ataç, S. Mitarai, J. Nyberg, A. Virtanen, A. Bracco, F. Camera, B. Million, M. Pignanelli, Nucl. Phys. A **571**(1), 185 (1994). https://doi.org/10.1016/0375-9474(94)90347-6
36. H. Wollersheim et al., Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip. **537**(3), 637 (2005). https://doi.org/10.1016/j.nima.2004.08.072
37. A. Maj et al., Acta Phys. Pol. B **40**, 565 (2009)
38. A. Gadea et al., Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip. **654**(1), 88 (2011). https://doi.org/10.1016/j.nima.2011.06.004
39. On the Road to FAIR: 1st Operation of AGATA in PreSPEC at GSI, INPC 2013 – International Nuclear Physics Conference, vol. 66. https://doi.org/10.1051/epjconf/20146602083
40. J.M. Régis, G. Pascovici, J. Jolie, M. Rudigier, Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip. **622**(1), 83 (2010). https://doi.org/10.1016/j.nima.2010.07.047
41. J. Valiente-Dobón et al., Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip. **927**, 81 (2019). https://doi.org/10.1016/j.nima.2019.02.021
42. J. Scheurer et al., Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip. **385**(3), 501 (1997). https://doi.org/10.1016/S0168-9002(96)01038-8
43. M. Assié et al., Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip. **1014**, 165743 (2021). https://doi.org/10.1016/j.nima.2021.165743
44. C. Ghabrous Larrea, K. Harder, D. Newbold, D. Sankey, A. Rose, A. Thea, T. Williams, J. Instrum. **10**(02), C02019 (2015). https://doi.org/10.1088/1748-0221/10/02/C02019
45. A. Pullia, G. Pascovici, B. Cahan, D. Weisshaar, C. Boiano, R. Bassini, M. Petcu, F. Zocca, in *IEEE Symposium Conference Record Nuclear Science 2004*, vol. 3, p. 1411–1414 (2004). https://doi.org/10.1109/NSSMIC.2004.1462504
46. P. Födisch, J. Wohsmann, B. Lange, J. Schönherr, W. Enghardt, P. Kaever, Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip. **830** (2016). https://doi.org/10.1016/j.nima.2016.06.019
47. I. Lazarus. Discussion about UK contribution to AGATA firmware algorithm improvements for LNL phase (2021)
48. M. Lauer, Digital signal processing for segmented HPGe detectors: preprocessing algorithms and pulse shape analysis. Ph.D. thesis, Heidelberg U. (2004)
49. S. Capra, G. Secci, A. Pullia, IEEE Trans. Nucl. Sci. **69**(7), 1757–1764 (2022). https://doi.org/10.1109/TNS.2022.3178760
50. S. Capra, IEEE Trans. Nucl. Sci. **67**(4), 722–731 (2020). https://doi.org/10.1109/TNS.2020.2975311
51. S. Capra, G. Secci and A. Pullia, IEEE Trans. Nucl. Sci. **70**(4), 730–736 (2023). https://doi.org/10.1109/TNS.2023.3259143