Notes from AGATA ADP meeting held at Legnaro on 1st December 2004

Present: Dino Bazzacco*, Marco Bellato*, Damiano Bortolato*, Pierre Edelbruck*, Andres Gadea, Xavier Grave*, Roberto Isocrate*, Ian Lazarus*, Denis Linget* Patrice Medina, George Pascovici*, Alberto Pullia, Ur, Christian Webber*. (* indicates people present for further informal discussions in Padova on 2nd December)

***Note**: these notes do not repeat information from the presentations made in the meeting. Presentations are available separately from the Padova AGATA web site at http://agata.pd.infn.it/meetings.html*

1) Introduction (Dino Bazzacco)

The presentation included an overview of the AGATA Data Processing teams which are now under the responsibility of Dino Bazzacco. The AMB will review the project status (specification, timescale and cost) on 10th December 2004 and information from this meeting will form part of that review. As a result of organisational changes at KFA Juelich, Werner Gast is no longer able to take an active role in AGATA and consequently the pre-processing algorithms activity needs to be re-considered (by the AMB)- should it remain as a team? If so, who should lead it?

An important change for the demonstrator is that the maximum singles rate (counting rate in 1 crystal) which must be handled has been reduced to 10kHz from 50kHz. This was decided in recognition of the problems of doing real time PSA on such a large amount of data. The pre-processing hardware must still be capable of handling the full (50kHz) rate.

2) Review of Actions from 23rd June LLP Interfaces meeting

(a) Time over threshold energy measurement for pions in preamplifier- it has now been agreed that a simple 10ns sampling of the preamplifier's Inhibit signal will give sufficient resolution. There is no need to add extra interpolation hardware or data links to improve further. The sampled Inhibit signal will be transferred from digitiser to pre-processor in bit D15 of the 16 bit ADC data words.

(b) The preamplifier output offset behaviour of the prototype detector at 50kHz counting rates has not yet been measured for an AGATA detector. This will be tested at Liverpool (**ACTION Ian Lazarus**). Tests have been performed in Italy with a Clover detector (AC coupled, normally –ve pulses) and the baseline shift was seen to go positive (opposite sense to pulses) with a time constant typical of the HV circuit as seen when ramping the HV up or down.

(c) Offset algorithms. This work is not complete and the question remains open of how to adjust for offsets and match the ADC input range to the preamplifier's output range. Patrice Medina proposed to overcome latency of links from digitiser to pre-processing by performing simple offset control in the digitiser using ADC under/over range. This would remove 1 link between the segments and the pre-processor. The proposal was not accepted because the algorithm is not yet decided. Also it was noted that the offset control is to be adjusted very slowly so latency is not an issue except for pulsed beams.

(d) Interaction of GTS trigger with the pre processing was discussed in June and the scheme proposed has been implemented.

(e) System latency results were presented during Marco Bellato's GTS talk.

(f) The use of Xilinx cores for slow control is no longer an issue. The digitiser uses €50 Xport chips and Xilinx have a free 10/100 Ethernet IP core. Orsay already have the Xilinx Gbit Ethernet IP. No one else needs Ethernet IP. Patrice Medina has found a true project-wide licence for Gbit Ethernet IP which could be bought for €16k if it becomes necessary in the future (this is from a 3rd party supplier of IP for Xilinx FPGAs).

(g) System C- a CVS repository has been created at Orsay. Ways of using it were discussed later (2nd Dec).

(h) Xilinx Trace port- the conclusion is that the trace port will share a bus with the DAQ data bus, so either one or the other is operative. This saves space but includes the trace port functionality.

(i) CPCI extenders have not yet been discussed- **ACTION Marco Bellato and Pierre Edelbruck** to discuss whether Marco's CPCI extenders are compatible with the backplane TCLK port.

(j) FPGA dimensioning in the pre-processing hardware has been done using information from the MWD implementations at Daresbury/MPI, IReS, INFN, Orsay and Orsay (none of these is adaptive).

3) Request from Ancillary Detector group for fast trigger output

At a recent ancillary detector group meeting there was a discussion about triggering during which the need for a fast multiplicity style trigger from the AGATA demonstrator was identified. The need arises in cases where the ancillary detector counting rate is very high, so a gamma-ancillary trigger is needed to cut the counting rate in the ancillary device. Normally this cut is applied at an early stage (before ADC gating) and ideally after 500ns to 1us. The AGATA system is pipelined which means that there are built-in latencies. For AGATA these latencies are handled by buffer memories creating effective digital delays. However, most ancillary detectors (apart from GREAT with its triggerless TDR system) require a real-time hardware trigger. The latency in AGATA means that the trigger cannot be produced at the GTS root node in less than 6 or 7us. This is too late for the ancillary coincidence trigger and so options for earlier triggers from AGATA to ancillaries were discussed. The only possibility is to provide a multiplicity output (logic signal) from each crystal based on some sort of digital trigger algorithm. Several algorithms have been implemented in FPGAs by members of the pre-processing team in IReS (CFD), Padova (SCC) and Daresbury (CFD) and also in DSP at Cologne. These algorithms have a latency too (around 500ns) due to pipelines internally. There is a further latency of around 800ns introduced by the serial links from digitiser to pre-processor (from serialiser, deserialiser and buffer memory), so implementing the algorithm in the pre-processing would generate a delay (latency) of at least 1.1us. Alternatively the algorithm could be implemented in the digitiser with an output port could be on the digitiser itself. The possibility of using a second preamplifier output and a separate electronics chain was briefly discussed but received little support.

There was a discussion about the necessary timing accuracy of the algorithm (bearing in mind the intrinsic timing properties of Ge and also coincidence windows in use today for such triggers) and it was agreed that

50ns FWHM is good enough. The low energy efficiency should be maximised since the timing requirement is not very tight.

There was also a suggestion from Andres Gadea that the requirement for a fast trigger might now continue from the demonstrator into the full AGATA. This led to a debate over triggering philosophies and there was no definite conclusion. Specific examples of ancillary detectors where the fast hardware trigger will be necessary were requested.

The discussion continued on 2nd December where the suggestion was made by Dino Bazzacco to send the trigger over a direct fibre (no serdes) from digitiser to pre-processor for connection to a front panel output at the pre-processing. Marco Bellato expressed concern about the impact that the fast trigger has on the operation of AGATA's own trigger, for example the concept of trigger partitioning is undermined by a common fast trigger. It was proposed that the fast trigger should be software controlled so that it can be disabled.

Summary of conclusions on fast triggers:

1) Discussion is not complete
2) Details of specific cases (counting rates etc.) are needed to clarify the requirement for fast triggers in the full AGATA.
3) It was agreed to include an output connection from the digitiser and the pre-processing from which a fast trigger could be sent. A third option was proposed during discussions the following day of generating the trigger in the digitiser, but sending it (via a direct fibre) to the pre-processing to be connected to the ancillary detector.
4) With the connections in place as described above, it is not necessary today to decide whether the fast trigger algorithm will be implemented in the pre-processing or the digitiser. So this decision was deferred. (Note AMB have since decided that the trigger algorithm will run in the digitiser.)
5) The generation of a trigger from the logic signals (and conforming to AGATA's grounding scheme) will be the responsibility of the ancillary detector group.
6) Trigger specifications agreed: Maximum latency (total) 500ns (The latency specification was relaxed to 1us in discussions after the meeting and in the AMB.) Timing of 50ns FWHM or better.
7) A slow control connection is needed to adjust the CFD threshold and to enable/disable fast trigger outputs.

4) Digitiser presentation (Patrice Medina)

The digitiser presentation is available separately at *http://agata.pd.infn.it/meetings.html* so only discussions and additional information are noted here.

The **weight** of the revised (2 module) mechanics for a single crystal is planned to be only 25Kg (rather than 45Kg for the original device).

The proposal in the presentation to remove the **offset control link** from the segments was rejected and the links will remain.

The latest **price** estimate is €30k plus taxes for the production devices (more for the prototype). There is a possibility to save money by assembly in-house at IReS rather than using a company. Some test assemblies will be made to develop the necessary skills at IReS. The advantage for AGATA is that the manpower is free.

Grounding was discussed. George Pascovici explained that he would like the digitiser to partition the analogue and digital grounds, coupling with a ferrite bead to filter high frequencies. Alberto Pullia reminded the meeting that the flow of return currents in grounds need to be considered. Patrice Medina showed the circuits and PCB layouts of manufacturer's development boards for several different ADCs, all of which used a single common ground for both analogue and digital parts. This is in accord with the Analogue Devices AD6645 data sheet which recommends a single ground (not a split ground plane) for best performance. The AD6645 data sheet also recommends that connections to the ground plane are made in such a way that return currents are drawn away from the ADC, not flowing under it.

It was decided that:

1) The digitiser grounds will be connected in accordance with the AD6645 data sheet, i.e. a single ground plane and care will be taken to ensure that connections to the outside are made such that return currents are directed away from the ADC (and its analogue input stage).
2) The mesh ground system which was proposed and accepted at the October 2004 Infrastructure meeting will be implemented initially for each full triple cluster in the demonstrator. This will simplify both the preamplifier and the digitiser by removing the opto-isolators and their associated power supplies.

ACTION Patrice Medina and Alberto Pullia will check the signal levels of all signals between the preamplifier and the digitiser (some, e.g. Inhibit and shutdown, use –5V and ground for logic signals).

Resolution.

Dino Bazzacco expressed his concern about low energy resolution in digital systems and declared a wish for a resolution of 1keV at 100keV (later in the discussion it was noted that the AGATA detector specification calls for 1.1keV at low energy). Patrice Medina reported measuring 0.9keV from a planar Ge detector at low energy with a TNT card adjusted for 0-10MeV full scale. The same detector was measured at 0.8keV FWHM using analogue electronics with 4us shaping constant. He also reported that the effective number of bits in the TNT2 cards depends on the dynamic range: the normal 0-5MeV range is measured at between 11.6 and 11.8 whereas the ENOB for the full 0-20MeV range is only 11.2 bits.

During the discussion the size of the input signal (from the preamplifier) was considered, and the matching of this signal to the ADC's input range is critical for good resolution at low energy. The issue of improving ADC performance at low energy by adding gain was discussed. In principle the gain should be added as early as possible in the system (preamplifier) but the preamplifier output range is limited by the ±6V supplies and the non-rail-rail buffer op amps (LM6172) in the preamps. Gain at the input to the digitiser doesn't improve the signal-noise ratio but can reduce ADC quantisation noise. It was agreed that we should compile a unified table detailing the preamplifier output signal characteristics (amplitude, gain, offset, range, rise time, fall time) and the digitiser input stage characteristics (dynamic range, offset, gain adjustment range etc.). These figures exist in the 2 specifications, but have not been written down alongside each other for easy comparison and checking.

ACTION George Pascovici and Patrice Medina to make a unified table of preamp/digitiser analogue signal characteristics and meaning (MeV equivalence).

As a result of the mismatch between the preamp output range and the digitiser input range it is necessary to control the gain and offset of the digitiser input. There is provision in the digitiser to control both these

parameters. There is also the provision in the preamplifier of the over-range recovery feature (fast reset). As noted in earlier meetings, Alberto Pullia reminded us again that offset adjustment and over-range fast reset are mutually incompatible- only 1 mechanism should be enabled at once. It was noted again that the core signal (which is AC coupled) will change offset only a little at high counting rates, but the DC coupled segments are susceptible to changes in offset due to rate. The effect in segments is reduced because their counting rates are lower than the core. Alberto Pullia described a student project at Milan for a negative feedback loop to control baseline. A proportional integral is made of the difference between target baseline and ADC output and sent back via a DAC to be added to the ADC input. The DAC value is also entered into a FIFO which is arranged such that the DAC data arrives at the output at the same time as the ADC data to which it was applied. A digital subtraction is performed in a manner analogous to sliding scale correction. As a result the effective ADC width can be increased by several bits.

The **DNL** of the AD6645 ADCs was discussed. Patrice Medina has measured the DNL in TNT2 cards and found the same structure that has been observed by Martin Lauer and by Ian Lazarus in separate tests. This is a problem caused by imperfect matching of the sub-ranging stages and, according to Martin Lauer's work, can be corrected by adding 2 more fractional bits (making the ADC words 16 bits) using a lookup table to replace the 2 lsbs with 4 bits. The lookup table is addressed from ADC data lines D8..5 which are the cause of the DNL structure. Ian Lazarus reported that tests suggest that each ADC will need its own lookup table. So a calibration is needed with ADC DNL correction factors stored in the database for each detector/digitiser set along with such things as measured decay time constant for preamps, gain matching coefficient(s) etc..

5) Pre-processing presentation (Ian Lazarus)

The pre-processor presentation is available separately at *http://agata.pd.infn.it/meetings.html* so only discussions and additional information are noted here.

The exact nature of the pulse shape analysis is not yet known (hardware or software) but the default position is to use a PC farm unless some hardware solution is proven better. Therefore the pre-processing output stage is to be designed in such a way as to be useful to a PC farm.

The format of the pre-processing cards was discussed. Initially the plan was to mount mezzanines on a Compact PCI (cPCI) hardware format for initial testing and then migrate later to ATCA-AS (PCI Express) if the technology matured sufficiently. However the protracted debates over such things as costs have delayed things to the point where we have already passed the mid 2004 decision point about using ATCA. So there was a discussion (started on 1st Dec, continued on 2nd) about the relative merits of building a cPCI card prior to an ATCA card. Ian Lazarus identified the risks of switching to ATCA as:

- technology maturity- is ATCA with PCI Express likely to be widely used long lived and sensibly priced?
- "right first time"- cPCI is a known format in which we have experience. ATCA is newer and involves a learning phase which increases the probability of a design mistake (and hence a rework with extra cost and delay)
- Delay the time scale for delivery of an ATCA card will be longer than for cPCI because of the learning phase.

The advantages of switching now are that we need only 1 development and can concentrate our limited resources on the version required for the demonstrator- the cPCI version was a temporary solution whose time seems to have already passed. ATCA allows various protocols to be used from simple Gbit Ethernet through to PCI Express and offers good flexibility for the PSA input stage. Technically it does everything CPCI does and a lot more.

A quick survey of the technology tells us that there are several suppliers of ATCA crates and some switch cards are starting to come to market. Commitment from the telecoms industry is growing and has shown itself in several large orders. PCI Express cores are available now from Xilinx as well as a development board in ATCA format from Xilinx and Avnet. The Advanced Switching features are not yet delivered but the signs are good.

An initial rough timescale was set out leading to an ATCA prototype carrier ready to start testing around end 2005 (cPCI would have been mid 2005). The tests would take place up to mid 2006 with production boards delivered by the end of 2006. This rough timescale (and a final decision about switching to ATCA) need more thought and a small meeting of the pre-processing team was planned for the following week in CERN to discuss technical details.

In any case the core and segment mezzanines would be ready for testing in mid 2005 and the GTS mezzanine in early 2005 (hardware only- no VHDL code).

The possibility of small changes to the pre-processing was mentioned (e.g. using 1000ns traces instead of 600ns in the pre-processing's 20us data buffer) and it was agreed that the design will be flexible enough to cover such requests and that where possible memories will be dimensioned generously (if it can be done without increasing the cost!)

6) GTS presentation (Marco Bellato)

The GTS presentation is available separately at *http://agata.pd.infn.it/meetings.html* so only discussions and additional information are noted here.

The GTS mezzanine PCB layout has been redesigned at CERN and a prototype card will be ready in early 2005 for the start of VHDL development by an engineer from Poland working with Marco Bellato. Signal integrity analysis has been performed on the PCB and simulations of the high speed paths look good.

As a result of the link alignment work presented in the talk, the GTS mezzanine as been changed to include switches (with 2ps jitter) for the fibres to allow bypassing of the rocket i/o buffers during link calibration. Equivalent bypass switches should also be added to the rocket i/o transmitters and receivers in the other pre-processing mezzanines and in the digitiser.

A Xilinx/Avnet ATCA board will be used as a rapid test bed for VHDL prototyping for root node and fan-in-out of GTS signals. The GTS mezzanine can be mounted in the user customisable area which has space for 2 CMC format mezzanines.

There was a discussion about how to synchronise and setup the AGATA system and in the course of this the targets proposed by Dino Bazzacco were: alignment within 1 crystal to better than 5ns and alignment between crystals of better than 10 or 20ns. It was noted that the test pulsers in the detectors are likely to have a variation of less than 5ns so could also be used for testing synchronisation. The GTS mezzanine has 2 trigger request inputs from the pre-processing core mezzanines, so 1 can be used for synchronisation. The digitiser could be programmed so that it can be switched (under software control) between setup mode and operational mode. In setup mode, the top ADC data bit would be used for synch (instead of inhibit). The

following system synchronisation paths could, therefore be used to fully characterise the system (see diagram).

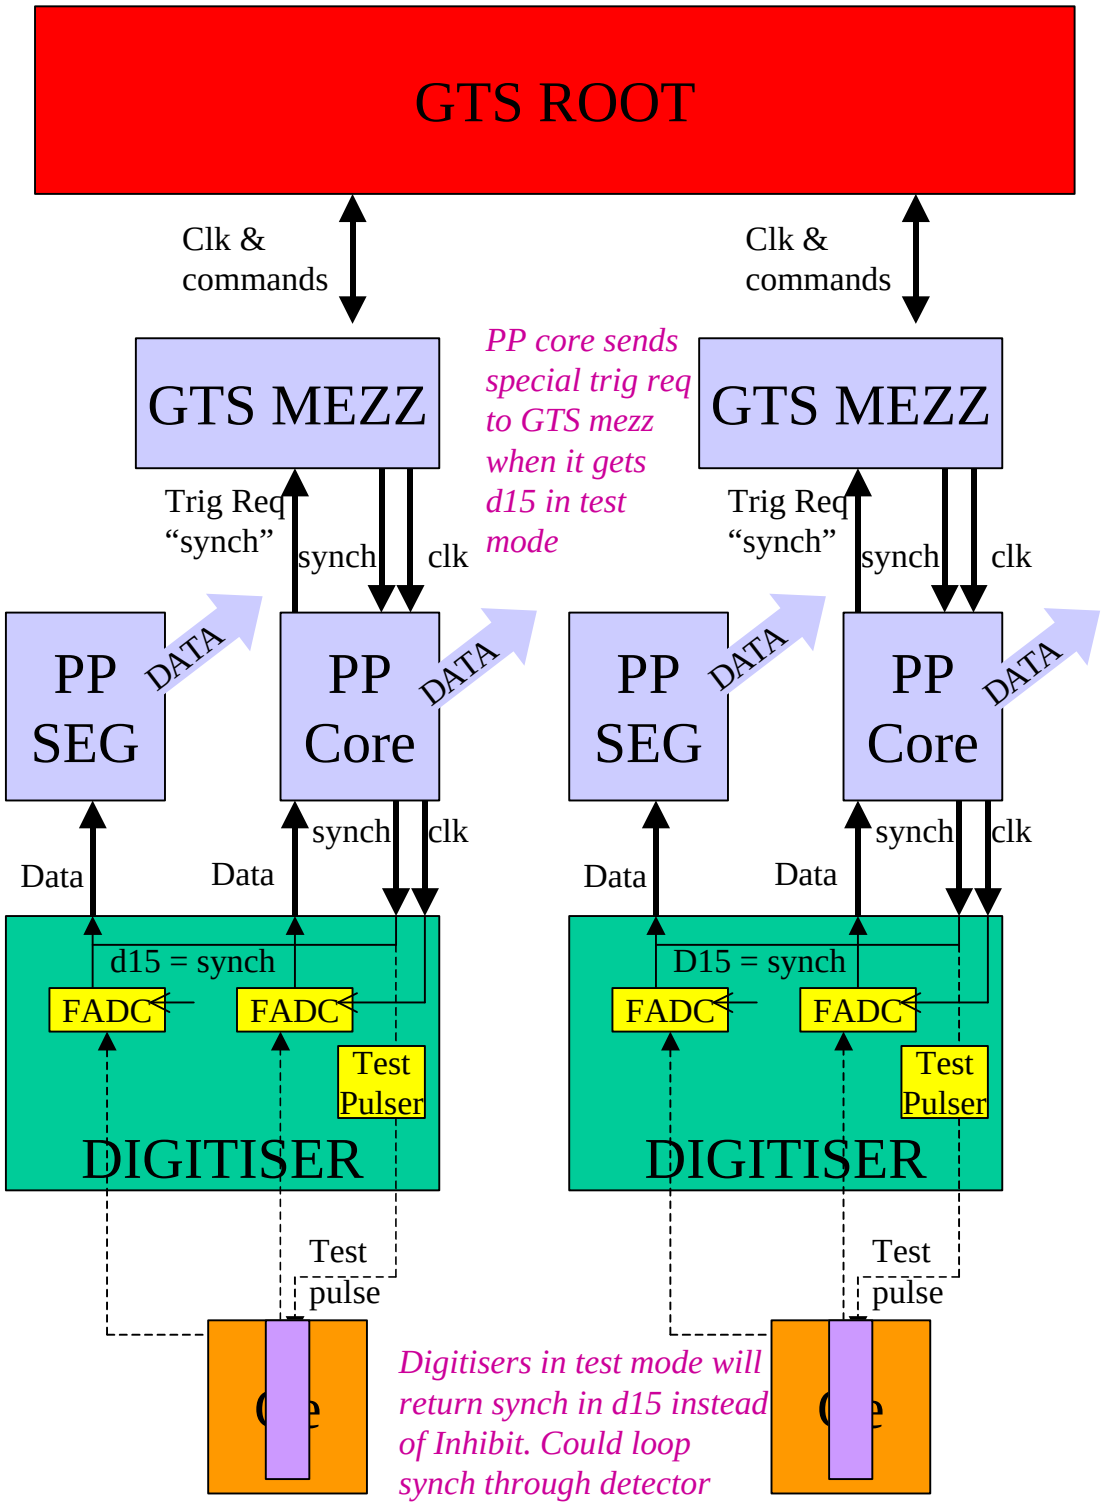

Documentation and AGATA Week

Dino Bazzacco requested that the specifications be updated to reflect the latest changes and with latest costs/timescales. Current versions: Digitiser V3 (Oct 2004), Pre-processing V8 (Oct 2004), GTS Jan 2004 and all require updating.

He also requested that team leaders keep in touch with him and with each other. Each team will need to present the present status during the AGATA week in February 2005.

7) System C (Discussion on 2nd Dec 2004)

Pierre Edelbruck reported that the CVAS repository at Orsay is ready but contains only his bare framework of the system. More detail must be added.

Marco Bellato and Lounis Benallegue have agreed on a unified set of signal names (not the same as in Pierre Edelbruck's model, which needs to be updated).

Christian Webber will start the System C modelling of the GTS and will update Pierre Edelbruck's model to add detail of the GTS.

System C simulations are typically using 10ns (1 clock) step size. The possibility of integrating system C simulation with Monte Carlo simulations of detector response through the preamp response into the system C digitiser was discussed. Whilst this would be nice it was thought unlikely that anyone would have the time to model the full system in this way.