# Notes from AGATA LLP Interfaces meeting held at CSNSM Orsay on 24th March 2004

**Present:** Marco Bellato, Lounis Benallegue**,** Damiano Bortolato, Patrick Coleman-Smith, Pierre Edelbruck, Werner Gast (morning only), Xavier Grave(morning only), Roberto Isocrate, Ian Lazarus, Sebastien Lhenoret, Christophe Oziol, Alberto Pullia, Cayetano Santos, Bruno Travers, Dirk Weisshar. (Reiner Kruecken, Juergen Eberth and Dino Bazzacco also attended for short periods.)

## Introduction and presentations

After a short introduction from Reiner Kruecken there was a series of presentations about the preamplifier (Alberto Pullia) the Digitiser (Patrick Coleman-Smith) the Pre-Processing (Sebastien Lhenoret), the GTS (Marco Bellato) and System Simulation (Marco Bellato). These presentations are included in the appendix at the end of these notes.

During and after the presentations many questions were discussed and later some people agreed to continue working on the answers to the questions. These notes are an attempt to record the important points from the discussions and our decisions during and after the presentations.

The names against the tasks are the ones we identified in the meeting. If other people want to contribute too then they should contact the people doing the job and offer help. I realised that we forgot to make anyone responsible for starting the discussion within each task. Please don't wait for someone else to do it!

**The timescale agreed was that where possible (within the limitations imposed by meetings of other groups) the tasks set out in the rest of the document will be completed (or showing significant progress) within 1 month.**

The remaining part of these notes describes the background, discussion and tasks allocated for these ten points:

1) Preamplifier: Inhibit output.
2) ADC Input offset: usage and control.
3) How do we synchronise (line up) all the channels?- dealing with link latencies.
4) Control of the preamplifier's pulser
5) Digitiser and pre-processor spare channels
6) Slow Control Connections
7) How to identify correlation between local triggers and validations
8) Data buffering in the pre-processing
9) System simulation model
10) GTS Mezzanine Pinout

I have added a couple of comments (*in purple italics*) which came to mind while writing the notes.

## 1) Preamplifier: Inhibit output.

**Background**
The preamplifier first stage is designed so that it will not saturate even with 50MeV pion events occurring at a rate of 10kHZ. However the second stage can saturate either due to these 50MeV events or a sustained high rate of lower energy events. So the preamps have a mechanism for fast discharge of the second stage using a switched current source, controlled by a voltage comparator which compares the output to a predefined threshold. The signal which enables the switched current source is sent back to the digitiser and labelled "Inhibit". By measuring the duration of Inhibit the amplitude of the saturating pulse can be measured. This combination of current source and time-over-threshold-counter is effectively a Wilkinson ADC and so can be used to measure energy. Tests with a pulser show that the Inhibit pulse has an intrinsic noise jitter of about 4ns rms and a width in the range from 1-10us for a typical energy range with good linearity.

**Discussion**
At the moment the preamp sends Inhibit to the digitiser, and receives an enable/shutdown signal from the ADC to shutdown or enable the fast reset comparator. The digitiser system sends the Inhibit bit to the pre-processor in the unused (D15) data bit with the ADC data samples (14 bits of data, 1 bit of ADC Over-flow, 1 bit Inhibit). It is proposed to measure Inhibit only on the core contact, not the segments. In fact the segments are grouped so that 3 segments share a single Inhibit, so such a method can't be used anyway. It was noted during the discussion that the pion's 50MeV will usually be localised in 1 segment.

Two possible changes were discussed: firstly to add a TDC in the digitiser using a frequency higher than 100MHz to measure the width of the Inhibit to the best available accuracy. This was rejected for 2 reasons- firstly there is no spare capacity in the data links to the pre-processor so a new fibre link would be needed and secondly it is not obvious how this data could be exactly time-correlated with the ADC data streams. Instead the Inhibit signal will be sampled every 10ns and sent with each ADC data sample as previously planned. *(NB there is a latency in sub-ranging FADCs; 4 clocks in AD6645, 8 clocks in AD9245, 14 clocks in the AD9444 and 16 clocks in ADS5500, so the Inhibit appears in the pre-processor's input data stream before the associated ADC data words.)*

The second possible change which was discussed was to move the control of the fast reset into the digitiser. In this case the ADC's full scale is used as the threshold and the ADC's over-range bit is sent back to the preamplifier to switch on the current source which discharges the 2$^{nd}$ stage of the preamplifier without any further check of the preamplifier output voltage. The advantage of this method is that it saves one connection between digitiser and preamp. The disadvantages are that both the start and the end of the preamplifier discharge is delayed by the ADC output latency period (40 to 160ns depending on the ADC selected for the digitiser) and that the preamplifier becomes dependent on the AGATA digitiser for its fast reset operation rather than a self-contained unit useable in other applications. This discussion was not concluded and no decision was taken.

There were 2 tasks allocated as a result of this discussion:

**Task 1**
**(Alberto Pullia, Dirk Weisshar, Werner Gast, Patrice Medina, Ian Lazarus, Denis Linget)**

a) Discussion of the implementation of the fast reset (whether it is to be controlled locally in the preamplifier or remotely by the digitiser's ADC) and whether there is any advantage in timing Inhibit on segments as well as the core (not possible with shared inhibit).
b) Discussion with pre-processing algorithms team about implementing an algorithm which adds the result of measuring the reset pulse width (either from Inhibit or ADC over-range) to the ADC value prior to saturation. This algorithm must take account of the ADC output latency and make appropriate corrections. The size (processing power) of FPGAs in the pre-processing should be reviewed to see if the new algorithm fits in the planned devices.

**Task 2**
**(Juergen Eberth)**
Discuss with AMB whether the proposed energy resolution for pions of between 0.2 and 2% (depending on pulse width and hence the energy) using the "free" Inhibit width measurement is good enough. If not then additional money must be spent to add a TDC and a readout fibre to each digitiser.

## 2) ADC Input offset: usage and control.

**Background**
The preamplifier output range is set to 100mV/MeV which means that with no baseline offset the digitiser ADC can accept 10MeV full scale. This range is reduced by baseline shifts caused by preamp output pulses sitting on the tails of previous pulses. The digitiser has 2 mechanisms to ensure that the input stays in the ADC's operating range: first a coarse gain switch (x1 or x 0.25 attenuation) and secondly a DAC which can shift the baseline to compensate for rate-related baseline shifts. The preamplifier's fast reset mechanism (see previous section) can be regarded as another way to regulate the baseline shift in cases with extremely high energy-rate products.

**Discussion**
Questions concerning the offset DAC were discussed-

- who controls it?
- what is the algorithm for setting the DAC?
- what effects do changes in the baseline offset have on the signal processing?
- how does it inter-work with the preamplifier fast reset
- how fast does it need to be updated (i.e. what fibre bandwidth is needed assuming it is externally controlled by the pre-processing)

The answers we came up with were:

- Q Who controls it?
- *A The digitiser offset DAC must be controlled by the pre-processor.*

- Q What is the algorithm for setting the DAC?
- *A We don't know- this is a question for the pre-processing algorithm group (see Task 3). Its not as simple as converting an average baseline value to a DAC setting to shift the ADC input because of*

*situations such as pulsed beams where the rate can change very fast and also because the rate of change of the offset is important- this is explored more in the next question.*

- Q What effects do changes in the baseline offset have on the signal processing?
- *A If the baseline is changed by the offset DAC during the MWD processing time then it will affect the energy resolution because new values and stored values will have different offsets and so differences between samples will be increased or decreased by the change in offset. In this case we would need to calibrate the offset DAC. With a DC input, values would be set in the DAC by the pre-processor and the changes in the ADC code noted. This calibration allows the DAC setting to be corrected in a way analogous to a sliding scale correction. This is certainly not trivial on a sample by sample basis over all possible DAC values and needs to take in to account the ADC output latency. However if the rate of change is slow compared to the MWD processing then the effect is one of changing the pole zero circuit to give an apparent increase or decrease in the preamplifier output decay time depending on whether the offset is added or subtracted. This will affect the signal processing in that the PZ correction will need to be adapted to account for the offset change.*

- Q How does it inter-work with the preamplifier fast reset
- *A As noted in the initial discussions, the fast reset and the digitiser offset should not be used simultaneously. It was noted that high energy pulses from pions will not be a problem in all experiments so the preamplifier's fast reset is not required in all experiments.*

- Q How fast does the offset need to be updated (i.e. what fibre bandwidth is needed assuming it is externally controlled by the pre-processing)
- *A We didn't come up with an algorithm to control the offset adjustment, so it was not possible to calculate the required update rate and hence the required fibre bandwidth. The decision to continue to control the offset from the pre-processing does, however, confirm that fibre links are needed.*

**Task 3**
**(Alberto Pullia, Werner Gast, Patrick Coleman-Smith, Sebastien Lhenoret)**
Discuss with pre-processing algorithms team about implementing an algorithm to control the digitiser's input offset (assuming that it doesn't operate at the same time as the preamplifier reset). This algorithm must take account of the effects of changing the baseline on the other pre-processing algorithms and also consider both fast baseline changes (e.g. pulsed beam) as well as gradual changes. Consideration should also be given to the impact on Pulse Shape analysis (e.g. does the PSA need to know the offset setting and when it changed?) The result of this discussion will affect the selection of the fibre link from pre-processing to digitiser. The size (processing power) of FPGAs in the pre-processing should be reviewed to see if the input offset algorithm fits in the planned devices or not.

## 3) How do we synchronise (line up) all the channels?- dealing with link latencies.

**Background**
During the digitiser and pre-processing presentations the question of synchronisation and link latencies was raised several times. The problem is that as a fast serial link is started up, the receiver must lock onto the incoming data stream, deciding where a word starts (which is bit 0?) and this takes some time.

Different links will take different times to do this. In fact the same link will take different times to do this on different restarts. This startup variation means that different links will have different latencies- some will have locked to the clock earlier than others, so will send their data out earlier. The latency affects the links in both directions (to/from GTS; from pre-processing to digitiser and from digitiser to pre-processing). It must be calibrated out dynamically (because it changes every time the link is started its no good to re-use old values).

**Discussion**
How do we calibrate the link latencies? It was noted that a return path is needed from the signal source to the receiver and back again in order to calibrate the latency- only the transmitter can know when it sent the signal, so it must also be the place to receive the returned copy of the signal and count the time difference. Such return paths are built into the system. In normal operation they are not set up to loop back calibration pulses, so a special setup/training mode will be required. For example the pre-processing can send a command with the clock signal to the pre-processor which can be distributed within the digitiser to all data paths (if the digitiser is in calibration mode). The pre-processor would wait to see when it comes back down each of the data links from each segment and the core, noting the differences. In normal operation these calculated latency variations will be used to set different latency compensation delays in the pre-processing input buffers so that all ADC sample data streams are time-aligned across the whole crystal before processing them. After the latency compensation delays all data relating to a certain instance in time (when a physics event occurs) will emerge from the delays in parallel, ready for processing together.

Finer alignment than 10ns would require the introduction of new timing verniers in the digitiser to phase-shift the ADC clocks. This is not included in the present specification and would incur additional cost and manpower to implement it. The question was raised as to whether sub-clock alignment is required by PSA (or even if it is possible to do).

**Task 5**
**(Ian Lazarus)**
Ask PSA group (Thorsten Kroell) whether sub-clock (better than 10ns) time alignment is necessary for PSA and if so how important is it and what time alignment should we aim for. Note that anything better than 10ns requires new money and extra manpower!

**Task 6**
**(Damiano Bortolato, Patrick Coleman-Smith)**
Produce a short proposal for how to do the training and calibration for link latencies and clocks for the digitiser, pre-processing and GTS.

## 4) Control of the preamplifier's pulser

**Background**
The preamplifier for the core segment has a built-in pulser which couples capacitively through the detector to all 36 segments, allowing the testing of back segments which would see very low count rates from a radioactive source on their crystal's front face. The amplitude and start time of this pulse

are controlled by the digitiser cards in response to commands sent via the slow control. What is not clear is whether there is a need in the digitiser to convert a single slow control command to a sequence of commands to the pulser (for example to slowly increment amplitude over 100 000 pulses to make a ramp to test linearity). It is simpler just to make a 1-1 correlation between the slow control and pulser commands. However the example of the ramp would need 200 000 slow control cycles to change the amplitude 100 000 times and to trigger the pulser 100 000 times. The possibility of firing multiple pulsers simultaneously to test timing was also raised.

**Discussion**
It was agreed that the digitiser core card will provide all the connections internally to implement state machines, such as the ramp, if necessary. This requires signal paths from slow control to FPGA and from FPGA to pulser control signals.
Testing timing with the pulser is not so easy because of the link latency problems described previously. The link latencies are corrected at the input to the pre-processor such that the signals are in line there. But in the digitisers the latency is not corrected, so the variation will result in a spread of pulse times even if a global start command is issued via the GTS global command link to all pre-processors. So making the detector pulsers useful for timing is not possible without spending extra money to provide a way to calibrate latencies in the digitisers. It was agreed that this decision to use the pulsers for energy only (not timing) can be reconsidered if a real need (backed up by extra money!) is identified. What is required is one more fibre from the digitiser to the pre-processor which is tightly coupled to the fibre issuing global commands and clocks. The new feedback fibre would be connected either to loop-back in the transceiver receiving the global clock and commands (no latency) or to loop back using the output of the global clock receiver (with latency) so that the pre-processor can measure the difference.

**No tasks identified.**

## 5) Digitiser and pre-processor spare channels

**Background**
Some redundancy is built into both the digitiser and the pre-processor so that faulty channels can be switched out and replaced by spares. For the core, the digitiser has 1 spare channel which is switched by enabling/disabling input buffers before the ADCs. There was some discussion about whether the spare channel needs its own fibre link or whether one fibre can be switched in the serialiser FPGA to accept either of 2 ADCs as its data source. For the segments, each group of 6 outputs runs down a 12 way ribbon fibre and one extra digitiser channel is provided to drive a $7^{th}$ fibre (5 unused). Analogue inputs are switched to the spare channel using a switched input buffer.

**Discussion**
It was agreed that the core is important for triggering and so it should have 2 complete data paths through the digitiser card and 2 fibres to the pre-processing. It was decided that the pre-processing will receive both core fibres and will use FPGA reconfiguration to choose which is processed. Similarly for the segments, FPGA reconfiguration will be used. In the case of segments there is one FPGA per 2 inputs (3 FPGAs per group of 6)- the spare channel can only be connected to 1 FPGA, so if the failed channel is in one of the other 2 FPGAs then the FPGA with the spare input must process 3 inputs

instead of 2 after reconfiguration. Therefore this FPGA will need to be resized (note that FPGA dimensioning needs to be re-checked anyway as a result of the offset algorithm (point 2).

**Task 7**
**(Patrick Coleman-Smith and Sebastien Lhenoret)**
Consider the use of spare channels in the digitiser and the pre-processor and check that what is proposed here will really work for switching spare channels (FPGA reconfiguration, software reconfiguration, database reloads for detector-specific parameters etc..)

## 6) Slow Control Connections

**Background**
During the discussion about the digitiser it became clear that we have not yet decided what happens about the slow control connection, in particular where it comes from. In the pre-processor the GLP's slow control Ethernet is connected to the carriers via RJ45 front panel connections and connected to an embedded PowerPC in the Xilinx FPGAs. An option proposed for the digitiser is to use a small external PC with Ethernet connection which drives several USB connections to the digitisers. This solution requires PCs to be bought (probably shared between several digitisers, so the cost would be low but care would need to be taken with earthing). Another option for the digitiser is to copy the method used in the pre-processor (RJ45 connector and PowerPC in the Xilinx FPGA). This requires isolated fibre Ethernet. A third option is to have a bidirectional slow control fibre link between pre-processor and digitiser (no Ethernet port on digitiser).

**Discussion**
The options described above were discussed a little and no firm conclusions reached. Input from the GLP group (in charge of slow control) is also needed.

**Task 8**
**(Patrick Coleman-Smith, Xavier Grave, Christophe Oziol, Dino Bazzacco (GLP), IRes (Marc Richer?))**
Look at options for connecting slow control to digitiser (direct, via PC or via pre-processing) and make a recommendation as to which should be used.

## 7) How to identify correlation between local triggers and validations

**Background**
The pre-processor interaction with the GTS was an area where we identified some confusion, so we discussed an example where multiple local triggers are generated in the time when we are waiting for a validation for the first. The question is this- which of the local triggers is to have its data selected for readout? The proposed GTS method is that the GTS trigger decision has a latency depending on the physics of the trigger (but limited to 20us by pre-processing buffer depth as agreed in Legnaro in September 2003). Associated with the latency is also an acceptance window. Local triggers are

accepted if they were generated before the validation in the time period defined by the latency period and the acceptance window. To help keep track of this process, the GTS mezzanine will update a full 48 bit timestamp every clock cycle *(how does this get sent to the other carriers?).* When the trigger validation is issued, it will come with an event number (24 bits) which is actually redundant in a timestamped system, but will be useful for diagnostics. It will also be associated with a particular timestamp from which the trigger latency and acceptance window limits are subtracted to identify the range of timestamps for which local triggers are accepted.

The pre-processing will take the ADC samples, extract the energy, the leading edge and perform any other processing. Then the resulting data will be put into a buffer and identified by a tag such as the timestamp of the local trigger which caused them. This is the data buffer which is searched after a trigger validation. *(What happens if the data are not yet ready when validation comes back?).* Data older than the oldest acceptable timestamp are wiped from the buffer (validations will always come in time order and with known latency so old data can be confidently discarded). Data which is too new to be accepted is retained.

**Discussion**
It was noted that for wide acceptance windows it is possible to have multiple local triggers and also for a single local trigger to be in multiple acceptance windows. It is not clear whether data from the local trigger should be discarded from the main buffer when it is written to the readout buffer after validation. It is possible that the same local trigger might be in a following validation too. It is not clear whether in this case data should be duplicated or not in the readout buffer. More discussion on this is needed between GTS, physicists, PSA and the pre-processing. There was also a lot of discussion about how close the local triggers can come because of pileup and the SCC algorithm's deadtime- a minimum gap of 1us was suggested as a realistic gap.

**Task 9**
**(Marco Bellato, Xavier Grave, Ian Lazarus, Sebastien Lhenoret)**
Consider and discuss the proposal presented for data selection based on trigger latency and acceptance windows. Consider how to implement it and look at alternative trigger protocols.

## 8) Data buffering in the pre-processing

**Background**
During the presentation of the pre-processing there was some discussion about the need for a dual ported memory on the carrier board to permit data to be written from mezzanines whilst being read from Compact PCI. The behaviour of this memory, its size and i/o bandwidth requirements can be studied using the System C model of the LLP developed by Marco Bellato.

**Discussion**
It was agreed that more discussion is needed between the pre-processor carrier designers and the people studying the system model (see item 9) to be sure that the output bandwidth and memory size are sufficient and that different options are modelled.

**Task 10**
**(Xavier Grave, Christophe Oziol with system model people (Marco Bellato, Lounis Benallegue, Patrick Coleman-Smith, Christophe Oziol))**
Consider dimensioning of the possible pre-processor carrier memory architectures and the necessary i/o bandwidth (using system model if necessary).

## 9) System simulation model

**Background**
Following earlier meetings a model for part of the LLP system has been developed in SystemC which permits system modelling at the clock cycle level. The model covers the pre-processing and GTS, allowing the study of trigger and readout mechanisms. The work now requires more people to get involved in developing the model's components and studying the effects of alternative component implementations (see for example point 8) and of changes to the system.

**Discussion**
There was some discussion about the availability of software tools (the System C software is free but a viewer needs to be purchased). The possibility of extending the model to cover latencies in the digitiser ADCs and transceivers was discussed.

**Task 11**
**(Christophe Oziol, Marco Bellato, Lounis Benallegue, Patrick Coleman-Smith)**
Continue to develop and test the system model, using it to test effects of changes to the system arising out of today's and any future discussions and to compare alternative implementations.

## 10) GTS Mezzanine Pinout

**Background**
The proposed mezzanine pinout is currently not compatible with the proposed carrier pinout for mezzanines, so some discussion and changes will be needed to converge the 2 proposals. The GTS mezzanine will be used also for the GTS fan-in/fan-out and GTS control functions so that only 1 design is needed although different FPGAs and different numbers of transceivers will be fitted in the different applications. Also the component heights and placements proposed for the mezzanine and carrier need to be checked for compatibility.

**Task 12**
**(Marco Bellato, Denis Linget or Lounis Benallegue, Roberto Isocrate, Christophe Oziol)**
Ensure that the carrier and the GTS mezzanine have compatible pinouts, board size, connector placements and component placements.

## Appendix- Presentations.

Preamplifier- Alberto Pullia

Digitiser- Patrick Coleman-Smith

Pre-Processing- Sebastien Lhenoret

GTS- Marco Bellato

GTS-Pre-processing system model- Marco Bellato.