Notes from AGATA LLP Interfaces meeting held at CSNSM Orsay on 23rd June 2004

Present: Dino Bazzacco, Marco Bellato, Lounis Benallegue**,** Piotr Bednarczyk, Damiano Bortolato, Francois Bourgeois, Pierre Edelbruck, Xavier Grave, Roberto Isocrate, Ian Lazarus, Sebastien Lhenoret, Patrice Medina, Christophe Oziol, George Pascovici, Vic Pucknell, Alberto Pullia, Henning Schaffner, Bruno Travers, Christian Webber, Dave Wells, Mirek Zieblinski, *other people were present too.*

Note: all task numbers refer to the list we produced at the 24th March meeting (summary is in Appendix A). These notes follow the chronological order of discussion in the meeting, not the task number sequence.

i) Preamplifiers (Tasks 1 and 2)

Alberto Pullia made a presentation about the preamplifier making 4 main points of which 3 were discussed:

- fast reset mechanism control
- Inhibit
- Energy Range
- Baseline shift

The preamp SHDN control can be used to either permanently disable fast reset (SHDN= high) or run it under preamplifier control (SHDN = low). It can be used in a third mode whereby it is switched from disabled (high) to enabled (low) whenever the digitiser detects an over-range condition, going back to disabled (high) when the digitiser detects ADC values lower than some predefined threshold. It was agreed that the 4 clock (50ns) latency in the ADC isn't significant and that the digitiser for the demonstrator should be designed so that the SHDN signal for the core can be driven either from slow control (always on/off) or from the FPGA which serialises that ADC data (preamp reset under digitiser control).

The presentation contained some preliminary results about the Time Over Threshold (TOT) measurement of the Inhibit pulse for Pion energy measurement. The measurements suggest that using an ADC to sample the Inhibit pulse transitions allows interpolation to determine a good threshold point on each edge and therefore an improved timing performance. It was pointed out that the galvanic isolation will have an effect on the thresholds and should be taken into account in these tests. The comparison of the interpolated method with the results obtained by simply sampling the presence or absence of Inhibit using a logic gate sampled at 100MHz has not yet been made. As noted at the previous meeting, the digitiser has no concept of timestamps, triggers or events, so any measurement of the Inhibit width in the digitiser involves a change of concept and some new way to correlate the processed Inhibit data with the raw ADC samples sent to pre-processing. Patrice Medina offered a simple vernier method (split 10ns steps into 4x 2.5ns) as an improvement on the counter without adding an extra ADC. This method still needs an extra data path.

ACTION Patrice Medina and digitiser team to calculate the cost and power budget of an extra fibre link to pre-processing from the core card and also the cost of an additional sampling ADC for Inhibit. The pre-processor cost will also need to be increased to receive the new fibre link.

ACTION Alberto Pullia to compare the results of the "free" method of measuring Inhibit (put Inhibit status into ADC D15 bit and send it to pre-processing to count the width in 10ns steps) with results from some sort of processing (ADC plus interpolation) in the digitiser.

The results of the tests and the associated price increase will be presented to the AMB if they improve significantly on the 0.2 to 2% resolution quoted at the last meeting (and subsequently agreed to be sufficient by AMB- task 2: email from Juergen Eberth confirmed this) using the simple counter method.

Inhibit TOT test results using a random pulser were presented and it was shown that a correction can be made for the contribution from tail of the previous pulse, arriving at the same resolution as with a fixed rate pulser. The need for such an algorithm was identified at the previous meeting (task 1b) and it has since been shown to work by Alberto Pullia's tests.

Energy range was not discussed.

Baseline shift of the core preamplifier is expected to be negative rather than positive at high rates because the core preamp is AC coupled. This is not a problem for the digitiser which allows full bipolar operation including offset control in either positive or negative sense. It was noted that this can be tested at Cologne using the prototype detector and a high rate source.
ACTION George Pascovici will test the offset behaviour of the prototype detector at 50kHz counting rates.

Use of Inhibit on segments- it was agreed that the segments will send the shared inhibit signal with the ADC data in bit 15 (top bit) and pre-processing can use ADC overflow to decide which segment is really being reset (this doesn't work if 2 of the segments are reset very close together, but in this case the OR of the Inhibits will distort the Inhibit pulse width anyway, so the information is already lost). The ADC/vernier method is not under consideration for segments because of the extra cost and the fact that the information is also available on the core (assuming that the pion interacts in a single segment only as was stated in the March meeting).

ii) Trigger Protocol (task 9)

It was agreed to make the trigger decisions in the GTS rather than replicating the logic in each of the segment and core mezzanines. Therefore the timestamp will not be continuously broadcast by the GTS. Instead it will be used as a data label and sent with the local trigger (and again with the validation or reject signal). Note that the centralised decision mechanism means that explicit rejections must be sent for events which are not validated. An event number is required by the software so will be sent with validations as well as the timestamp value. The revised system is sketched on the next page.

The core and segment mezzanines are now simpler because they don't have to keep track of time. They can only measure time indirectly by using the local trigger as a time reference and measuring relative to the local trigger. A consequence of this is that the local trigger must be distributed to all mezzanines with low skew and a fixed relationship to the 100MHz system clock. The actual value of the timestamp must be captured by the GTS with low jitter too, but distribution to mezzanines of the timestamp value is not time-critical because the essential time information is carried in the local trigger signal. There was a discussion of how we cope with overlapping events and it was agreed that the situation where local trigger #2 is coincident with validation #1 would cause contention on the 8 bit bus used to transmit timestamps from GTS to other mezzanines. So it was decided to use 2 different 8 bit buses for local trigger and for validation. The inter-trigger time was also considered to ensure that we have time to transmit a timestamp + event number for the shortest inter-trigger gap. For physics reasons the triggers can't meaningfully come closer than 1us, so this means there is plenty of time to transmit 6 timestamp bytes and 3 (or 4?) bytes of event number.

The question of time-ordering the validations was discussed. The local triggers will be time ordered, but it cannot be 100% guaranteed (yet) that the GTS will send the validations (or rejections) in strict time order. This is, however, a design objective and will be implemented if possible in GTS.

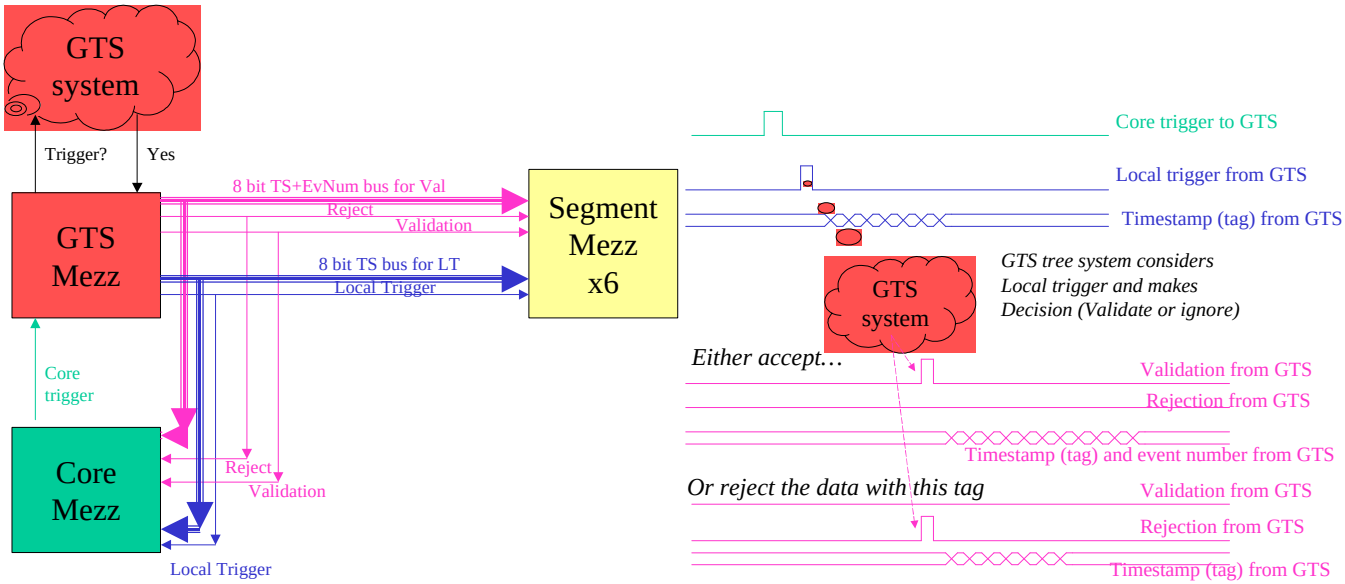

Diagrams showing operation of the pre-processor triggering system.

iii) System Latency (task 6)

The work on calibrating out the latency is not yet finished (task 6). Damiano Bortolato presented a scheme whereby the GTS rocket io link is set to loop back at the far end and is restarted many times until the measured link latency for TX + RX is found to be either a minimum or a maximum value according to a calibration giving the distribution of all possible total latency values. It is assumed that the extreme minimum and maximum values occur when both TX and RX links are at their minimum or maximum latency simultaneously. The assumption that the links are symmetrical is also made (initial tests by Marco Bellato suggest that this is a valid assumption; more work is being undertaken). No results are yet obtained, but by restarting the links continuously for a period of a few seconds it is estimated that the value can be measured ±100 or 200ps. Further work is taking place to test this idea

The possibility of using the core pulser was also discussed and it was agreed that this will provide a useful test of the calibrated latencies- if the signals arising from a test pulse to the detector core can be seen simultaneously at the preprocessing then the link latency calibration will be assumed correct. (There will be some additional errors due to differences between segments in detector cabling, segment capacitance etc which will prevent the inputs being exactly coincident even if the links are perfectly compensated for latency errors.)

Marco Bellato raised the question of temperature effects in the fibres- the jitter and propagation delay are both affected by temperature. Typically the propagation delay changes by 60ps per degree per km. It was agreed that the first people to have a suitable fibre will make temperature/timing tests.

Patrice Medina described the clock distribution within the digitiser- there is a fixed path length to each FPGA using a Lattice clock distribution chip (50ps max jitter).

It was concluded that work will continue on all 3 strands: the extremes method from Damiano Bortelato for latency calibration, the track matching and adjustment in the digitiser described by Patrice Medina and the independent verification using the core pulser.

iv) PSA Time Alignment (task 5)

Ian Lazarus reported Thorsten Kroell's personal answer (he's no longer the PSA team leader) that getting all samples aligned within a single 10ns clock is probably enough since PSA will re-align all inputs against a simulated pulse sampled at 1GHz.

Dino Bazzacco reported that the reformed PSA team was not planning to meet until the early autumn. It was agreed that we need an answer before that to avoid design delays. Dino suggested that a figure of 5ns alignment window across a crystal would be a suitable design target for the latency calibration and delay matching described above (iii).

v) Slow Control for digitiser (task 8) and pre-processing

It was agreed that the digitiser's slow control will no longer be connected via fibres connected to the pre-processing. Instead the Ethernet will be directly connected to the digitiser for slow control. The standalone USB ports for lab testing will be retained.

There was a proposal from Christophe Oziol for a pre-processing slow control interface. The idea is that on the carrier the Ethernet connection (RJ45) will be connected to a media converter and then, using Media Independent Interface (MII), connected to a software switch in the FPGA power pc on the carrier card. From the carrier card, MII links will connect the Ethernet to the other mezzanines. In this way the high current external interface for Ethernet is only done once. The software switch is a standard switch implemented by Xavier Grave within his own Linux implementation (this software is completed). This mode of operation was agreed by everyone to be a good idea and was adopted.

The use of Ethernet cores for Xilinx was discussed- it was reported that the cores purchased by Orsay are on a pre-project basis (i.e. available outside Orsay). Ian Lazarus has been told by UK's Xilinx distributor that cores are purchased for use on a site (for any project, but only 1 site). Marco Bellato uses CERN cores for CERN projects by connecting remotely to CERN computers (i.e. is "virtually" in CERN). It was noted that the EDK software includes free power pc code for Ethernet. The question about cores needs more work:
ACTION Marco Bellato to check with Italian Xilinx contact whether we can buy cores on per project basis.
ACTION Pierre Edelbruck to check whether Orsay core licence allows use outside Orsay but within the AGATA project.
ACTION Ian Lazarus to check with UK Xilinx contacts whether we can buy cores on per project basis.

vi) Spare Channels (task 7)

During a recent digitiser team meeting it was agreed that the digitiser team would propose the following: the core spare channel will be normally used to transmit data to pre-processing using a different gain range. The pre-processor will select whichever input is appropriate (largest unsaturated input) in normal operation. If either link fails then the other will be used as the only core link and, if necessary, its gain will be switched to a new value. The segment spare channel will not be a fibre (only 6 fibres will be used), but will be

implemented by multiplexers in the serialiser FPGAs such that any of the 6 ADCs can be replaced by the spare ADC channel just by changing a multiplexer setting. The use of the segment spare channel will therefore become invisible to the pre-processing. This proposal was accepted.

vii) System C (task 12)

The Orsay groups have implemented system C models and offered to make available a CVS repository for shared system C code to ensure that we all have access to the same codes. This is subject to satisfactory arrangements about offsite access being made between Pierre Edelbruck and the Orsay software people. The CVS offer was accepted and details of how to use it will be sent out by Pierre when its all arranged.

viii) Offset Algorithm (task 3)

It was agreed that the offset will not be adjusted on a sample by sample basis. Instead a slow adjustment will be used at a rate such that it doesn't affect the MWD algorithm noticeably. The exact rate is to be determined after tests to see how fast we can adjust a baseline offset without affecting MWD.
ACTION Sebastien Lhenoret, Patrice Medina, Patrick Coleman-Smith to find out how fast we can adjust the baseline without affecting the MWD and consequently the specification of the fibre link we need.

ix) Pre-processor pinout and connections (task 11)

The TCLK port was discussed. Christophe Oziol proposed a scheme using the top (user defined) part of the PICMG 2.16 (CPCI) crate. Boards would be mounted on the back (behind the backplane) with bidirectional LVDS buffers and connected together with either a ribbon cable or a special PCB. This technique is necessary because of front panel space limitations. There was some discussion about whether we can send 100MHz signals through the backplane connectors (and down a ribbon cable). It was noted that a Gbit ethernet can be run over those connectors so this suggests that 100MHz operation should be OK.

The idea of broadcast commands (transmitted and received via trigger root node) was discussed. This has replaced the dedicated serial i/o links we originally discussed for inter-board connections. In principle the same function can be performed by the broadcast links. Examples of use are mainly to do with transmission from the root node to all cards; the up-link to the root node might not be needed but will be left in for the short term at least.

x) GTS Status

The GTS PCB will be 16 layers and the detailed design will be completed in a couple of weeks. All components are bought. However the pinout can still be updated because the interface is implemented in FPGAs. The mezzanine needs only a 3.3V supply, deriving all other supplies locally. The estimated maximum current needed by GTS is 6A. Only 7.6A in total is allowed from CPCI and the remaining 1.6A will not be sufficient for the core mezzanine and the carrier. However a further 5A is available from the 5V CPCI supply. It was noted that the fibre transceivers are very sensitive to power supply quality.

xi) JTAG Connection

A JTAG connection is available via the PICMG 2.16 backplane but its not yet clear if there is space on the carrier card to use it. It was agreed that JTAG connections are desirable if possible.

xii) Agilent/chipscope Trace port

The GTS mezzanine has no room for these connections and so the possibility of mounting them on the carrier board was discussed. Its not clear that there is space on the carrier either. A bus switch could help ease the

space problem and Bruno Travers agreed to investigate this. Ideally at least 1 FPGA per mezzanine (GTS, Core, Segment) should have this connection available.
ACTION Bruno Travers will look at the possibility of using bus switches for trace port connections.

xiii) CPCI Extenders
Marco Bellato has a CPCI extender but the Orsay groups need to buy one. However, its not clear how we can use an extender if the TCLK cards are mounted behind the backplane. Marco will send Pierre details of the extender owned by Padova.

xiv) Prices
There was a discussion about revising the prices as requested by AMB.

Albert Pullia reported that the preamp price has already been reviewed and discussed in the detector group.

The Digitiser price is being reviewed by Patrice Medina (including now a separate estimate of the extra cost to include an extra fibre link if required for measuring the Inhibit signal width more accurately).

At the moment there is no reason to change the pre-processing prices from those in the December 2003 estimate. However, the Orsay groups are now recalculating the prices in detail and revised figures will be available in a couple of weeks. The potential new fibre from the digitiser to pre-processor and the additional TCLK backplane cards might add to the cost, but these items will be reviewed alongside the overall price review. Another unknown is the FPGA cost in the pre-processing because the algorithms are not yet finalised. This will be tackled firstly by asking the pre-processing algorithms group (via Werner Gast) for estimates of FPGA resources needed for algorithms and secondly by making a list of FPGA resources needed for all algorithm implementations made by DL/MPI, IReS, Orsay, Padova/LNL for MWD and trigger algorithms. These will then be expanded to include adaptive operation of MWD and new algorithms (e.g. pion energy measurement) before confirming FPGA sizes.

ACTION Ian Lazarus to contact pre-processing algorithms group leader (Werner Gast) to ask for their estimates of resources need for algorithms in AGATA pre-processing.
ACTION Patrice Medina, Roberto Isocrate, Sebastien Lhenoret (and anyone else with MWD/trigger algorithm implementations) to send details of FPGA resource useage to Ian Lazarus.

Appendix A- Summary task list from 24th March 2004 LLP Interfaces meeting

Preamplifier: Inhibit output.

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.

ADC Input offset: usage and control.

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.

There is no task 4.

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

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.

Digitiser and pre-processor spare channels

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..)

Slow Control Connections

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.

How to identify correlation between local triggers and validations

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.

Data buffering in the pre-processing

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).

System simulation model

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.

GTS Mezzanine Pinout

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.