

Correlator meeting, Arpad's room, October 11 2012, 11.00 am
Present: Jonathan, Des, Jintao, Harro

agenda

actions:

Jonathan: test fix of Erlang performance problems reading UDP packets
(on hold)

Jonathan + Salvatore: figure out how to make packet spacing adjustable
(on hold)

Harro: save N11L4 scan 11 reference data to LTO4
stays: still waiting for LTO4 tapes

topics:

updates

aob

next meeting

-updates

* unofficial action items

Des mailed around the URL of his VDIF timestamping memo. Jintao mailed pulsar-gating memo #5 for discussion.

* control system

Des has made a lot of progress on the control system. It is now at the point where it must be known where all the data needs to be sent. The buffer addressing scheme (ip-address/port number) has been changed lately but the documentation about it is not up-to-date yet. Jonathan promises to inform Des as soon as it's updated. Des also requests more N11L4 data from Harro.

action Jonathan: inform Des when buffer addressing documentation is updated

action Harro: create more N11L4 VDIF data

* firmware

Jonathan got data out of the UniBoard and captured into datafiles with byteswapped data. One station seems to produce good data, the other products are wrong. Discovered an error in the back node validity handling. In the front node also problems with the validity handling were discovered. Back node should hopefully be fixed. There's still a problem with the validities when switching data streams in the front node. Jonathan has an idea what the problem might be. Finally the amount of stored validity bits needs to be doubled: at the moment only one validity bit is kept for both polarizations. (This is connected to the change in input buffer addressing/layout). This may affect timing since the memory blocks are far apart.

The output of the front node in simulation was giving all idle pattern but in hardware it seems to work. Possibly a problem in the high speed transceiver reset. It was discovered that the reset pulse was too short compared to the transceiver documentation specified value. Coming to think about it: glitches in hardware transceiver reset may have been observed. The reset pulse should now be long enough.

Jonathan reports that Salvatore is running back node simulations with the new DDR controller (because of updated Quartus version to 11.1 (or 2)) and is seeing some timing issues. He is trying different placements and rearranging 'logic lock' regions - which can have an effect on the placement/routing and thus timing.

* pulsar gating

Jintao produced a memo (#5) on the pulsar gating, investigating the worst case pulsar period variation for binary pulsars. The memo seems to indicate that once per minute is not enough, even for the slowest varying pulsars: after 30 seconds the accumulated error on the linear approximation goes outside the acceptable bin range of + or - 1. The fast varying pulsars are - obviously - even worse. However, one second update period with linear approximation is certainly good enough. Refer to memo #5 for the actual error numbers.

The coefficients can be pre-computed or computed on the fly - the tempo2 software (which generates the coefficients) is fast enough to generate them within one second.

* offline software

Harro has made good progress with j2ms2 - a measurement set with the data captured by Jonathan was generated. Labels are still incorrect but that will be fixed soon; the focus had been on the mechanics of getting the data correctly from file to measurement set. Some plotting tools in Python need to be developed. Des suggests we should ask Mike Sipior if building a full casa install is an option; it may have ample plotting capabilities but we don't know: we haven't actually used casa for years.

action Harro: talk to Mike Sipior re. casa

-aob

Jonathan remarked he read Harro's timing document and said it is ok - what's described is what we currently do. Harro will upload it to the wiki.

action Harro: upload the timing/synchronization document to the wiki.

-next meeting

Thursday 18th October 11.00 am, after jive coffee