Correlator meeting, Arpad's room, September 20 2012, 11.00 am
Present: Jonathan, Des, Jintao, Salvatore, 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: write up a small doc about timing mechanism
in progress, remains

Jintao: write up a document discussing the different ways in which pulsar gating coefficients and gating waveforms can be transmitted to the uniboard
done remove

Des/Harro: save N11L4 scan 11 reference data from Mark5 diskpack to file
We now have two scans (scan 11 + scan 12) for 13 stations on file.
Comment Des: store converted-to-VDIF data as well?
stays but edited: Harro save raw data to LTO4 tape

topics:

new FFT
control system
pulsar gating
timing of correlator
actual tests
timeline
offline software

aob

next meeting

-new FFT

Jonathan: new design synthesized but not all components have their interfaces into the new filterbank changed accordingly (e.g. packet_rx). The filterbank fits easily and the design is currently being simulated. Results expected this week. Jonathan skeches a picture of the VHDL modules and their interrelations and illustrates which parts need changing. Also the latency of the new filterbank must be measured - some important signals between modules go past the filterbank but must arrive at the same time as the data that has gone through the filterbank.

-control system

Jonathan: the input buffer is now divided into four segments/buffers, each representing 1 second's worth of data. The size of the segments will depend on the resolution and the sampling rate of the incoming data, both as programmed by the control software. The amount of segments will be fixed to four for all sampling rates so even when doing very low data rate observations there is stll only room for four seconds of data. Making this more flexible requires too much work at this point (significantly reworking packet_rx) so it will be deferred to a later time - the incoming VDIF packets are not written back-to-back yet. The correlator supports handling data of different input resolution simultaneously but not different sampling rates.

The FIFO changes have been implemented. The 'full' bit remains even though we know there is a fixed amount of segments: it is crucial for correct operation. The definition of the contents of the FIFO are now officially changed to hold the amout of integrations the correlator can process from the corresponding buffer/segment. Processing starts when the amount of elements in the FIFO is greater than one.

The start time of the correlator triggered a fundamental discussion. At the moment this is just one 32-bit number, the VDIF second value. This is definitely not good enough. Each VDIF timestamp consists of a reference epoch + number of seconds since that epoch. The problem lies therein that (1) stations are free to choose their reference epoch and (2) within a timespan of 34 years there are 68 epoch + seconds-since-epoch combinations that represent the same actual UT second. The VDIF timestamps of the incoming data (and the starttime) should be normalized to a reference epoch of choice. This is not trivial as leap seconds between VDIF reference epochs should be taken into consideration as well.

Action Des: come up with a plan of how to deal with timestamps in the correlator

Des: created the CCS component that, given stations, vex subbands and desired polarizations queries the correlator control database and divides the data sources over the JUC inputs and the correlator product numbers to enable. Close to integrating with startup + configuring the JUC itself.

-pulsar gating (PG)

Jintao: Found an article dating from 2000 in which people at the VLA demonstrate pulsar gating using a PC to feed the phase predictions and gating waveforms into the correlator.

Jintao's document which lists three possible PG approaches is discussed. We still lean towards the PC-based solution as it is the most flexible and yielding the best science case. This method would add 256Mbps of input data to each front node for the current design.

Salvatore remarks that adding logic for pulsar gating support (extra 10GigE interface, FIFO, DDR controller) may well make the design not fit on the front node and possibly results in having to drop number of stations processed in one front node in favour of PG. Concluded it is at the moment impossible to see which solution would actually *fit* in the front node design.

Action Jintao, Salvatore, Jonathan: produce a small report investigating the logic use of the three PG scenarios and see if they will fit or what has to be deleted to make it fit on the front node

Action Arpad: schedule PG meeting with sfxc people (Aard? or both?) discussing Jintao's three PG scenarios to comment on their science/applicability/comparativeness to sfxc.

-timing of correlator

Salvatore: the designs seem to meet their timing ok

-actual tests

Jonathan: using the Erlang data sender manually data was sent to and succesfully received by the JUC. After the FIFO was programmed with a number data was sent to the back nodes. The data capturer was not set up so no output was captured to disk. The test will be repeated with the capturer enabled.

-offline software

Harro: working on j2ms2 to produce measurement sets out of JUC data. Still working on data structures to build up the mappings and organizing the data. The decoding of JUC output data packets triggered another discussion: what does the timestamp in the output actual mean? The timestamp is built out of two values: an integer second part and a framenumber. Turns out that at the moment it is completely useless; it is the VDIF start second and a VDIF framenumber. The latter is completely unrelated to the timestamp of the current integration.

After some discussion it was felt that with arbitrary integration times a physical encoded timestamp (either by samplenumber within second (which is unpredictable on account of delay compensation) or otherwise) would be impractical. Given that an integration is made out of an integral number of FFT segments and each FFT segment is an integral number of samples the integration length can be computed at high precision by CCS.

It was decided that the integer second field should go and that the framenr becomes the integration number, counting from 0 at the nominal correlator start time and then upwards. 32 bits of integration number yields, at the shortest integration time of 20ms, 23860 hours of continuous integration in one scan without overflowing. The VDIF epoch + start time should also be carried in the meta data because j2ms2 needs to know the origin of the integration counter.

-aob

nothing

-next meeting

Thursday 27th september 11.00 am, after jive coffee