IVS Memorandum 2008-014v01

26 August 2008

“VLBI2010 Digital Back End (DBE) Requirements”

Bill Petrachenko

V2C Memo: VLBI2010 Digital Back End (DBE) requirements

Bill Petrachenko
August 26, 2008

Introduction

As a starting point for discussion, I present and discuss here (without going into details) the functions that I think should be in the VLBI2010 DBE. I tried to be somewhat aware of what might be feasible in real hardware but still include everything I thought useful.

Anti-alias filter

To handle the standard VLBI2010 1 GHz band, the DBE anti-alias filters should be 1.024 GHz in width. Whether the first, second or higher Nyquist zone is used will depend on the details of the preceding down-conversion system and the bandwidth of the sampler. It is understood that customized filters may be necessary at some sites to eliminate large sources of RFI in the band.

Sampler

To handle the standard VLBI2010 1 GHz band and at least the second Nyquist zone, the sampler clock rate and bandwidth should be at least 2.048 GHz. Wider bandwidths typical of modern samplers will enable even higher Nyquist zones, which may be desirable.

Since the VLBI signals will be re-quantized to 2-bits after the PPF&FFT (see next section), it is necessary that the sampler resolution be larger than 2-bits to avoid a significant second quantization loss. In the absence of RFI, a practical sampler resolution might be 4-bits, with each additional bit providing 6 dB of headroom for RFI. At least 8 bits of resolution are recommended for the VLBI2010 sampler, although higher would be preferred.

Poli-Phase Filter plus Fast Fourier Transform (PPF&FFT)

The main parameter to specify for the PPF&FFT is the sub-band bandwidth. For numerical efficiency of the FFT it needs to be a power-of-two sub-multiple of the sample clock frequency. Furthermore, for efficient RFI excision and thorough spectral monitoring, narrower sub-bands are better. Although the final decision is somewhat arbitrary and may depend on downstream computations, I think bandwidths as low as 1 MHz are not out of the question.

Polarization Conversion

Although polarization processing may best be done post-correlation, if good enough amplitude and phase control can be achieved, another option is to convert linear polarization digitally to circular polarization in the DBE.

Re-Quantizer

For efficient transmission to the correlator, the PPF&FFT output needs to be re-quantized to 1 or 2 bits. For optimal 2-bit data, the +/- voltage thresholds should be set near 1 sigma. Since performance varies slowly with threshold, wired thresholds may be adequate under many circumstances. However, use of sub-band specific thresholds will be more robust under sub-optimal conditions and should be implemented if possible. Threshold values can be determined from sub-band power monitoring or stream statistics. [These are both discussed in later DQA/calibration sections.]

Corner Turner

For each output clock pulse of the PPF&FFT, data is grouped naturally as a set of complex data points, one per sub-band. However, distribution to correlation resources is done most efficiently if data are re-grouped into continuous streams for each sub-band. This is referred to as corner turning and is efficiently implemented in the DBE.

Sub-band Selection

Not necessarily all sub-bands are output from the DBE to be transmitted to the correlator. Sub-band selection should be arbitrary to allow adaptability to changing band optimization schemes, RFI environments, etc.

Burst Acquisition

In order to minimize acquisition time, data needs to be acquired at as high a rate as possible. These bursts of data need to be buffered so they can be transmitted at a rate matched to the storage media while slewing to the next source. The DBE may be a convenient location for the buffering.

DQA and Calibration

The main DQA and calibration functions are to monitor power levels, assess digitization thresholds, detect phase calibration signals, measure cable delay, and perform spectral and temporal decompositions to evaluate RFI and electronic transfer functions. Each function can be performed most naturally at different points in the signal processing chain, e.g. prior to PPF&FFT, between PPF&FFT and re-quantization, or after re-quantization.

- **DQA and Calibration functions prior to PPF&FFT**.

Pcal detection. A very efficient time-domain approach can be applied at this point to extract all Pcal tones simultaneously. It requires that the Pcal signal repeat every μs (or repeat at some other regular interval) in the sampled data (which in turn requires that the effective LO and Pcal frequencies be separate by an even MHz). Since the Pcal signal repeats every μs, it is possible to average the signal at individual sample points over the 1 μs interval. The FFT of the averaged signal then extracts the Pcal amplitude and phase for each tone. The averaged signal can be written

$$\overline{S}_i = \frac{1}{M} \cdot \sum_{j}^{M} S_{ij} \qquad (1)$$

where

S_{ij} is the raw sampler output for the i^{th} sample in the j^{th} μs

$\overline{S}_i$ is the average of the i^{th} sample in the μs

M is the number of μs over which the average is taken.

Cable delay detection. The above Pcal detection process enables the following method of cable delay detection. This method requires that the Pcal signal be generated in the control room and split. One of the signals is transmitted by fiber to the antenna focus for normal injection into the astronomical signal and the other is injected into the signal in the control room just prior to the splitter before the UDC's. Two pulses then appear in the averaged signal in equation 1. Their time separation is, to first order, twice the cable delay. However, the process is not perfect. A second fiber path is required to transmit the Pcal signal to the focus and its delay variations may be different from the signal cable. However, if all fibers are in the same bundle, perhaps the delays will be roughly the same. A second drawback is that the astronomical signal path includes additional devices. Half of the delay variations in these devices will be bundled into the cable delay measurement. However, these delays are not likely to correlate with antenna direction and hence should not bias the position estimate of the site. They are more likely to be clock-like and estimated along with the clock, unfortunately biasing estimates of precise time transfer. In fact, nothing addition needs to be done in the DBE to enable this method of cable delay detection. However, it is an additional reason to implement the Pcal method described in the previous section.

Power monitoring. Power detection of the input signal is required both for radiometry and for setting the sampler input power to near an optimal level. Since the front end noise diode will be switched on and off throughout each second (perhaps at 80 Hz), it is necessary that power levels be detected synchronous with the on/off signal. This can be done either by reading accumulators at the on/off rate or by accumulating into separate registers. Power is calculated in the usual way, i.e.

$$P = \frac{1}{N} \cdot \sum_{k}^{N} S_k^2 - \left(\frac{1}{N} \cdot \sum_{k}^{N} S_k \right)^2 \tag{2}$$

[Note: A means should be available external to the DBE to switch off the input power to enable the determination of a zero power level.]

There may be advantage to detecting power at each sample within a μs as with the Pcal detection method, i.e.

$$P_i = \frac{1}{M} \cdot \sum_{j}^{M} S_{ij}^2 - \left(\frac{1}{M} \cdot \sum_{j}^{M} S_{ij} \right)^2 \tag{3}$$

If the power level is found to be lower at the peak of the Pcal pulse, this could indicate that the Pcal pulse is saturating.

Autocorrelation. Spectral decomposition will provide useful information about RFI sources and aggregate instrumental band shaping. However, if the resolution of the PPF&FFT is large enough, additional spectral information may not be needed. Autocorrelation is likely to be costly in resources for an FPGA, but perhaps an existing high performance custom correlator chip could be added to do this function.

- **DQA and calibration (for each sub-band) after PPF&FFT and before re-quantization**

Power monitoring. Power detection, according to equation 2 (where S_k are now the sub-band outputs) of each sub-band output is required to set the re-quantization thresholds for that sub-band and to provide spectral information about the input sampled signal (assuming autocorrelation is not included in the DBE design). Perhaps a gate would be useful to cut out pulsed RFI. Detection synchronous with the noise cal on/off signal is required.

Pcal detection. If it is expected that tones will generally occur at arbitrary frequencies (and hence be undetectable prior to the PPF&FFT), a number (one or more) of standard VLBI tone detectors will be required per sub-band

- **DQA and calibration (for selected sub-bands, one or more) after PPF&FFT and before re-quantization**

Time-binned power detection. To gain information about pulsed RFI, power should be detected into higher resolution time bins.

Pcal detection. If Pcal detection at arbitrary frequencies is not included for each sub-band, it should be included (several tones) for the selected sub-band.

PPF&FFT. If additional spectral information is required, the selected sub-band can be expanded in a second level PPF&FFT.

- **DQA and calibration (for each sub-band) after re-quantization**

Stream statistics. The number of data points at each re-quantized state is counted.

Appendix A

Below is an email from Ed Himwich and Brian Corey discussing further details of digital back-end radiometry.

From: E. Himwich, B. Corey

Re: Digital back-end radiometry

Since there are currently at least two independent Digital Back-end developments for VLBI underway right now, we wanted to start a discussion for development of a specification for how radiometry should be supported by the hardware. Hopefully this discussion is occurring early enough that we can avoid having any really necessary features missing from the new systems that are currently being developed. To seed the discussion, we are proposing a draft specification. It is based on our experience with the more traditional back-ends and on discussions with some of you. If we have radically underestimated what can easily be accomplished in a digital world, we hope someone will pipe up with better information. We always want to do things better as long as it is practical. Conversely, we also hope to hear if anything is likely to be impractical. Please let us know if we have left anything important out. Our goal is to develop a consensus specification that can reasonably be achieved and will yield high-precision radiometric data sufficient for high-quality astronomical VLBI imaging.

Our basic suggestion is that we should be able to measure Tsys with a precision of 0.3%. Assuming a noise adding radiometry cal that is about

10% of system temperature, this implies a precision of ~0.03% in the basic power measurements. By "precision" here we mean both repeatability under identical operating conditions and relative accuracy (independent of the constant scale factor relating output counts to input power) between measurements at different operating points, as occurs when Tsys or the system gain changes. Any overall, consistent scale error in the measured Tsys values should be less than 1%.

Based on these assumptions and practical operating considerations, we propose the following specifications:

1. Linearity: 0.3% or better over operating range (see item 3). Linearity is defined by the difference in the read-out when the same step in input power is applied at different operating points. This should hold for all sizes of steps of input power.

2. Quantization: should be at least a factor of two more precise than either the time-bandwidth product of the detector or 0.03% of Tsys, whichever is smaller.

3. Range: The range should be sufficient to cover input power ranges of Tsys to twenty times Tsys, where Tsys is the lowest normal operating signal level of the system. This allows you for example to compare sky and warm load measurements.

4. Zero level: If the detector can have an offset greater than 0.3% of Tsys, a method needs to be provided to measure it at the 0.3% level. The zero level should be stable to less than 0.03% of Tsys on time scales of seconds, and to less than 0.3% on time scales of hours.

5. Device integration periods available: one second, but others may be useful too. For synchronous detection, longer integration periods may be desirable.

6. Traceability: For any AGC stages in the system, it must be possible to know the actual values in use well enough so that the final constructed read-out precision is no worse than the above specified precisions. In particular the overall precision of 0.3% in Tsys must be met. This may imply making the changes in AGC synchronous with the detector integration periods.

7. Read-out frequency: It should be possible to read the integrators (and any additional required information such as AGC values) once per integration without putting a significant load on the device communication bandwidth along with any other normal processing that occurs. The read-out should require a small enough fraction of the integration period so that it can be easily started and completed within one integration under normal operation using the timing information in item (8) below. If possible the device (or set of devices if results from more than one device are returned together) should return a 4 or more bit sample number counter (arbitrary initialization) that can be used to verify that read-outs nearby in time came from different integration periods. It may be desirable to have a mode that will allow all the detectors to be read back with one request and return values from the same integration.

8. Timing: It should be possible to determine where in the integration cycle the system is with enough precision so that the control computer can synchronize itself with the update cycle. This could be done either by a separate query to the detectors or by having the detectors return the information about when the response was sent for each sample or group of samples. The information should allow the timing to be determined at the millisecond level. e.g. the number of milliseconds into a one-second cycle can be returned in a 16 bit integer. It is desirable to be able to synchronize the cycle to an external 1 PPS.

9. Parallel integration: It is desirable to at least have a mode, if it is not the only mode, where all devices in the system integrate with the same start time.

10. Synchronous detection: It is desirable to have synchronous detection using a noise diode modulated on and off by a square wave.