Summary of Discussions Held During the July 10-14, 2000 Workweek in Socorro Regarding the EVLA-WIDAR Correlator

*NRC-EVLA Memo# 005 – **DRAFT***

Brent Carlson, August 22, 2000

ABSTRACT

This memo summarizes the discussions that occurred during a series of meetings that were held during the week of July 10, 2000 in Socorro between NRC (Carlson) and NRAO (Clark, Perley, Rupen, Ulvestad, Escoffier, Hankins, Blachman, Butler, Walker, Romney, Napier, et al) personnel. These were technical meetings and were used to discuss and refine the specifications and signal processing of the proposed EVLA-WIDAR correlator. At NRAO's request, this memo also contains a record of discussions within NRAO regarding the correlator design, and NRC's response to them.

Table of Contents

1 Pulsar Processing Issues

- Phased-VLA output must be able to be routed to the Baseline Boards for autocorrelation processing. If this complicates the design too much, NRAO could develop a separate autocorrelation board that performs this task.

- Cross-correlation processing where the phased-VLA (or phased sub-arrays) looks like a single antenna (i.e. phased-VLA output feeds into the correlator as if it were another antenna) would be nice if it does not complicate the design.

  Action: NRC will endeavor to design the Phasing Board so that both of the above requirements can be met as long as the complication and extra cost is not too severe.

- NRAO would like to have 2048 pulsar phase bins if possible. In all previous correlator requirements documents, a maximum of 1024 bins was desired.

  Action: NRC will consider 2048 bins but notes that this may require prohibitively large memory on the Baseline Board. Cost and design complication may be a significant factor here.

- The Baseline Board design and firmware must be capable of reducing the number of lags dumped from each correlator chip so that faster dumping can be obtained with reduced spectral resolution.

  Comment: The current design of DUMPTRIG and the Baseline Board supports this. The dump rate will be determined by the speed of the hardware readout controller. It is important to note that the minimum correlator granularity is 128 lags per cross-correlation product. It may be possible to reduce this for fast dumping, but there is no guarantee that it will be possible.

- Very fast real-time dumping (without phase binning) is desired with dump times down to ~25 microseconds. This is mostly for phased-VLA autocorrelation and probably is not required in interferometer modes.

  Action: NRC will endeavor to design the Baseline Board and readout controller so that nothing, other than data volume bottlenecks, limits the minimum dump time. Meeting the 25 microsecond requirement may require more than one hardware readout controller and therefore this is a cost issue.

- Very fast dumping with phase binning is desired with dump times down to ~25 microseconds. This is required for phased-VLA autocorrelation and interferometer cross-correlation.

Action: NRC will endeavor to design the Baseline Board to meet this requirement. There may be an additional cost since additional readout controllers are most likely required to meet the spec. NRAO stated that 125 microsecond dumping is probably sufficient in inteferometer mode.

- It is highly desirable to bunch phase bins across only part of the pulse period so that the phase bins are optimally placed (say, on the pulse).

  Action: The DUMPTRIG protocol supports this requirement. Actually enabling this capability is a matter of the design of the DUMPTRIG generation circuitry on the Station Board and is not foreseen to be a problem with excessive cost implications.

- Pulsar gating is a requirement for the correlator. It was decided that one independently programmable pulsar timer is required for each baseband (i.e. two per Station Board) and that multiple (one per sub-band) independent epochs and gate widths be generated from each timer.

  Action: NRC will include this functionality in the design of the system.

- It is required that the phased-VLA digital output be able to be connected to external pulsar processing machines.

  Comment: This requirement is inherently met in the current Phasing Board design.

2 Radar Mode Issues

- NRC wanted to know the minimum integration time requirement when operating in radar mode. In this mode, ~30 kHz bandwidth with ≤ 1 Hz resolution bandwidth is required. This mode will normally be achieved with recirculation which imposes restrictions on the minimum integration time. NRAO stated that the minimum integration time is ~3 seconds and also that the radar signal is a CW rather than a chirp.

  Comment: It appears that with recirculation active it is possible to achieve close to this minimum integration time using one sub-band correlator—freeing the other 15 sub-band correlators to be used for the simultaneous wideband continuum measurement. However, depending on the design of the recirculation buffer, an SNR loss of about 6.5% at the edge lag channels may be incurred (using 240k memory size yielding an integration time of 3.84 seconds—refer to NRC-EVLA Memo# 004) since it may be necessary to use single-port memory for speed and capacity (the SNR loss is eliminated if dual-port simultaneous read/write memory can be used). NRC will endeavor to design the recirculation buffer so this SNR loss is not incurred provided it is not a significant additional cost and technology is sufficiently capable (i.e. large dual-port, high-speed

memories, or some sort of alternating read/write scheme when the burst duty cycle is < 50%).

- NRAO requested that the sub-band digital radar mode signal (from the Station Board) be available via a dedicated front panel or rear connector. This will allow external computers to capture and process the data if necessary.

  Action: NRC will build this capability into the Station Board. There will be one connector for each baseband (two per Station Board). Each connector will include a clock at the sub-band sample rate, 4 data lines, and some sort of timing signal such as TIMECODE or perhaps just a 1 second time tick. Signal timing is TBD.

- NRC wanted to know what the maximum strength of the narrow band radar-mode signal is. The signal strength can affect the design of the FIR filters—requiring the carrying/accumulation of many more bits in the adder tree than with just noise because the adder tree, in this case, sees the results of convolving a narrow bandpass with a narrow signal. NRAO stated that the narrow radar signal will be a maximum of ~50% of the total power in a 30 kHz wide band.

  Action: NRC will ensure that the FIR filter is designed to be able to handle this type of signal.

3 Phased-VLA VLBI Issues

- NRAO stated that it is desirable to have at least 16 Phasing Boards—each one producing phased-VLA output for one sub-band—and provision for more (up to the maximum available) if possible.

  Action: The number of phasing boards that are supported in the correlator directly impacts the "Sub-band Distributor Backplane" design since it is this backplane that provides access to station data needed by the phasing subsystem. NRC will endeavor to design the Sub-band Distributor Backplane so that the phasing subsystem has access to a minimum of 16 sub-bands (to support 16 Phasing Boards) but with a goal of providing access to all. The current cost estimate (NRC-EVLA Memo# 001) is for 8 Phasing Boards. NRC will consider delivering 16 Phasing Boards provided it is within the estimated cost envelope.

- Currently, (according to NRC-EVLA Memo# 001), the Phasing Board phases antennas in bunches of 4 before going to the final sub-array adders (and each of the antennas can be inhibited from being added in the bunch). NRAO does not see this as a problematic limitation but asks NRC to consider removing this restriction.

  Action: NRC will consider removing this restriction in the design of the Phasing Board.

- To facilitate recording narrower (sub-)sub-bands while still retaining the total bandwidth available with the delivered Phasing Boards, NRAO would like NRC to consider putting FIR filters on the Phasing Board before final outputs. Each of these FIRs should have access to all phased sub-array outputs.

  Action: The current conceived design of the FIR filter (NRC-EVLA Memo# 003) supports this functionality and NRC will endeavor to include this capability in the design. The number of filters that will be available is determined by cost and available board space.

- NRAO stated that the nominal "Sub-band Distributor Backplane" design is such that 2 sub-bands are available from each baseband. With 8 basebands, this leads to a total of 16 sub-bands.

  Action: This configuration requirement will be included in the design. (Indeed, this requirement may be met by just duplicating the outputs of two connectors that go to the sub-band correlators!)

- NRAO requested that NRC consider using digital single-sideband mixers (DSSB) on the output of the Phasing Board so that arbitrary placement of narrow sub-(sub-)bands can be achieved rather than requiring placement within strict integer slots.

  Action: NRC will consider including this requirement in the design of the Phasing Board since the proposed design of the FIR filter, and the fact that the data is not time division demultiplexed at this point makes it seem feasible to do this. The number of DSSB mixers/filters that can be provided on each Phasing Board is subject to cost and board space. Also, the optimum placement of the filters (i.e. before or after phasing) requires further investigation.

- It was agreed that NRAO would be responsible for the design of the VLBI recorder interface that selects and formats one or more Phasing Board outputs for data recording. This interface could conform to the newly developed VSI-H interface, VLBA interface, or both. NRAO stated that they do not want to have to handle TIMECODE if the data is going to a VLBA recorder "output box"—the instrumental delay through the phasing subsystem and recorder interface will simply be calibrated. Everyone agreed that this instrumental delay should be constant even through power cycling.

4 FIR Filter Issues

- Programming time for all FIR filters in the entire correlator should be about 1 second.

  Action: This requirement should easily be met without having to have dedicated shadow memory on the FIR filter chip or separate memory on the Station Board. It is not clear whether this includes on-the-fly calculation of FIR filter coefficients or not. If it does, then FIR filter frequency design methods using FFTs will most likely have to be employed since sophisticated general purpose FIR filter design

software is often very slow. Frequency design methods have been used successfully in the investigation of the WIDAR technique.

- The FIR filter should, as a minimum, be a total of 1024 taps with two stages (i.e. 512 taps per stage as outlined in NRC-EVLA Memo# 003). Also, simulation has shown, that at least 12 bits should be available out of the 1st stage lookup table. With this minimum configuration it will be necessary to have a separate radar-mode filter with 1024 taps per stage. If a chip can be made with 1024 taps per stage, then a separate radar-mode filter will not be required.

  Action: NRC will endeavor to design the filter to meet the minimum requirement. For cost reasons, the chip will probably be an ASIC—although a full custom device will provide superior performance (more taps). If additional interested funding partners can be found, then it may be affordable to build a full custom device.

- The issue of FIR filter bandpass equalization was raised by NRC—that is, using the FIR filter to compensate for non-uniform gain in the analog baseband. NRAO stated that this probably is not a useful function for interferometer mode operation—although it could be very useful for phased-VLA operation.

  Comment: The FIR filters are capable of performing this function and so this is purely a software design issue and is therefore at NRAO's discretion.

5 Recirculation Discussion with Ray Escoffier

An informal discussion with Ray Escoffier was held regarding the recirculation design outlined in NRC-EVLA Memo# 004. A synopsis of the discussions is as follows:

- Ray confirmed that the size of recirculation memory is determined by the readout time of the correlator chips. In the memo, a size of 2 Mbytes (i.e. two buffers of 2 Mbytes each) was estimated to be sufficient (readout time of ~8 msec).

- In the memo, it was stated that there is an SNR loss (greatest at the edge lag channels) if recirculation is used depending on the number of lags to be synthesized and the size of the memory. Ray pointed out that the SNR loss does not occur if the buffers are written to and read from at the same time since, once half full, the read pointer advances at the same rate as the write pointer. However, this operation requires dual-port memory and currently (and perhaps in the foreseeable future) dual-port (synchronous static) memory is not large enough nor fast enough to do this. The dual buffer, single port memory proposed to be used in the memo does have the capacity and speed to be used in this application—but if used it will incur the stated SNR loss.

  *Action: NRC will monitor dual-port technology and use it if it is cost effective so that the SNR loss will not occur. This may be aided by decreasing the dump time of the correlator chip—requiring less RAM. The dump time decrease *may* be a*

natural consequence of meeting high performance pulsar requirements. The current plan, however, is to used DDR SDRAM.

- The no SNR loss design could be achieved with dual, single-port memory buffers if there is burst dead time. That is, recirculate less than the factor that is possible when synthesizing a large number of lags. In this case, the controller alternates reading and writing of data to effectively provide dual-port operation—but the number of lags that can be synthesized will be less than otherwise could be available.

  Action: If single-port memory must be used in the implementation of recirculation, NRC will endeavor to include this (programmable) capability in the design of the recirculation controller.

6 Delay and Phase Models

- NRAO stated that they wish to use the program "CALC" to generate real-time delay and phase models for the correlator. That is, the correlator will get 0^{th} and 1^{st} order (point-slope) floating-point coefficients for station baseband delay and station sub-band phase every 50 milliseconds. NRAO will develop the software to do this and NRC software will be the recipient of these models. NRAO requested that the NRC software maintain a very small queue of models (few hundred milliseconds to a maximum of ~1 second) so that source changes could be effected very quickly.

  Action: NRC will design the station control software accordingly. Also, the "DELAYMOD" and "PHASEMOD" signals in the correlator support arbitrary update rates and can therefore easily handle models every 50 milliseconds.

- NRC wondered if there were any requirements for real-time delay models into the correlator such as from a WVR for high-frequency observing. NRAO stated that any real-time models will be merged with geometric models before being fed to the correlator so this not a concern for the correlator.

7 Correlator Monitor, Control, Software, and Testing Issues

- NRAO stated that they want to dynamically schedule the operation of the EVLA to make optimum use of telescope time given the current conditions (weather, time-of-day, array configuration, observing queue etc.). Observing schedule changes could be in response to a human request or could be automated with computers making the scheduling decisions. Thus, the correlator must be able to be quickly reconfigured within 1 second or so. This means that any data or control queues should be short and reconfiguration should not require much pre-trigger or setup time. This requires a high-performance control-computing environment with dedicated LANs that are not slowed down with other non-real-time, non-critical data transfers.

  Comment: NRC will provide a PC-based computing back-end for correlator control and data handling. PCs were chosen to capitalize on the commodity nature of the computers: leading edge performance can be purchased for very low

cost because PCs have been engineered for the mass market. To meet performance requirements, all "hard real-time" processing will occur on the Station and Baseline Boards—the PCs will handle "quasi-real-time" computing. Short queue buffering on the hardware boards will ensure that the PCs do not have to immediately respond to hardware requests.

- Some time was spent discussing what the interface between the PCs and the correlator hardware boards would be. In NRC-EVLA Memo# 001, NRC suggested that this be some high-speed SCSI interface and that the memory on the boards look like a "RAM disk". It was agreed that SCSI may not be appropriate (although it hasn't been totally ruled out) for this interface because of the seemingly high data transfer protocol complexity. Both NRAO and NRC agreed that, in any case, there should be two interfaces between the PCs and the hardware boards: one interface is for high-speed data transfer, and the other interface is for monitor and control.

  Comment: NRAO and NRC will continue to consider various interfaces for the PC to hardware board connection. The important principles are that it should be an industry standard interface that can be added to a PC by simply purchasing a mass-produced interface card, and that the interface should support data transfers on the order of ~100 Mbytes per second so that there is ample room for future performance upgrades (i.e. by replacing the generic PCs with faster ones that will surely continue to become available).

- The choice of operating system (OS) for the PC back-end is not entirely clear. Both NRAO and NRC have considerable experience with VxWorks however, use of a generic OS like Linux (favoured by NRC) has some advantages such as its ease of programming, its non-proprietary nature, its portability, and its longevity. Documented source code for Linux is available, virtually guaranteeing complete control of the computing environment in perpetuity—something that cannot be said for VxWorks. Additionally, the "hard" real-time performance requirements of the PC back-end is not as stringent as is traditionally the case since performance, and the consequent design of the system, requires dedicated hardware on-board the Station and Baseline Boards to handle real-time functions.

  Comment. NRC will provide a PC back-end with the Linux operating system. If NRAO wishes to have VxWorks, then they will have to incur the cost of associated development licenses and target licenses. NRAO and NRC will continue to consider the choice of OS and the choice may not be made until the underlying hardware (correlator hardware daughter board and interface) is more concretely defined.

- NRC will provide a "logical correlator" software interface to the correlator. At NRAO's request, the interface will be/use CORBA and NRC agrees that this is a good choice. The software that NRC provides consists of two "layers" and a separate set of test functions that plugs into the top layer. The bottom layer consists of the hardware device drivers. This software provides a convenient interface to write and/or read specific hardware locations to perform specific functions. The second

(top) layer is used to allow high-level software to logically configure and monitor hardware and read out data. Depending upon the underlying hardware, the bottom layer could be in the PC or it could be on the Station/Baseline Board's daughter board. The test function software is used for exercising correlator functions before higher-level software is available. The test software will be delivered with the system, but eventually all of its functions should be available in higher-level software.

- NRC has proposed that system testing be developed within the context of normal observing operation. That is, a "test observation" will be performed just like a normal observation except for a few key words that indicate that a test is being run. The Station Board will be equipped with a test vector generator with an entry point close to where data enters the board from the antennas. When a test observation is run, data from this test vector generator will replace normal antenna data and back-end data analysis software will check for errors and pinpoint hardware faults. NRAO stated that, for maximum system integrity, it is useful to be able to run test observations whenever data from antennas is useless (i.e. antennas are slewing).

  Action: NRC will design the hardware and software to accommodate testing in the manner described above. As an addendum to the discussions, NRC suggests that it is useful to have pseudo-random test vector generators at the antennas (and receivers/checkers on the Station Board) so that when tests are run, the integrity of the fiber-optic connection between the antenna and the Station Board can be checked.

- NRC stated that, depending on available manpower and funding, they would potentially like to be involved in the development of back-end data handling/image processing software. NRAO stated that they are open to this kind of collaboration.

8 Correlator Data Handling Issues

- It was agreed that the back-end handling of the large volumes of data that will be generated by the correlator requires further study. NRC will provide the low-level software to be able to read out data but where it goes from there, where it gets stored, and how and in what form it gets archived are issues that are to be resolved and handled by NRAO.

  Comment: NRC will be responsible for providing the back-end PCs with nominal temporary storage disks. NRAO will be responsible for providing large volume back-up media (RAID drives, tape drives etc.).

- NRC commented that it may be possible for the hardware LTA controller on the Baseline Board to remove the data bias (bias inherent in the data so that ripple counters can be used in the correlator chip accumulators) before accumulation. This bias is not a function of the data or data valid, but rather is directly calculable based on the number of samples that are integrated in a hardware integration time. This should reduce the output data word size and data volume considerably.

- NRAO stated that initial quantizer histograms (state counts) and re-quantizer histograms will be required to be able to perform coarse-quantization corrections (Van Vleck) to the data.

  Action: NRC will provide this capability as part of the Station Board design.

9 Sub-band Stitching and Calibration Issues

- NRAO expressed some concern about the ability to always stitch together the sub-bands so that a seamless wide-band spectrum could be obtained. This concern was based on the perceived need to have a sub-band that is clear of any narrowband signals to perform the stitching. NRC made the point that seamless stitching *does not* require a clear sub-band. Sub-bands can *always* be stitched together if the lag 0 autocorrelation data before requantization is available. The output spectrum can always be normalized to produce the same level as that from a wideband correlator (see NRC-EVLA Memo# 001). If, however, a sub-band is available that is clear of narrowband signals, then the output can be made virtually immune to the effects of some time-variable interference that is in one or more other sub-bands. The immunity is not complete however, since time-variable interference will produce time-variable (initial) quantization noise that will show up in the clear sub-band—although normally this is a very small effect.

- Some time was spent discussing data calibration (i.e. converting normalized correlation coefficients to Janskys on the sky). It would seem that there are three ways that this can be done:

  1. Switching noise diode with analog detector before sampling. This detector would be susceptible to time-variable interference anywhere in the analog band.

  2. Switching noise diode with a synchronous sub-band lag 0 autocorrelator before requantization. Any sub-band clear of time-variable interference could be used for calibration.

  3. Phase-cal (tone comb across the band) injection and digital extraction (possibly in the correlator). The effectiveness of this is not clear and it adds signals to the wideband spectrum that may not be desirable.

  It was decided that one separate FIR filter be available for each baseband (i.e. two per Station Board). This FIR filter would not produce data for correlation, but rather could be independently programmed to find an interference free (arbitrary bandwidth) sub-band. Dual-bin, synchronous (to noise diode switching) lag 0 autocorrelator dumping would then be performed. This data can be used to produce normalized correlation coefficients immune to time-variable interference and to calibrate the data to Janskys.

  Action: NRC will design the FIR filter with two lag0 autocorrelator bins and the ability to dump the data into the bins synchronous to an external or internal

timer. On the Station Board, a separate FIR filter will be available for each baseband to acquire calibration data as described above. NRC will also consider putting phase-cal extractors on the Station Board but only if it can be done with no significant cost and only if there is board space to do it. The number of phase-cal extractors that will be provided is TBD.

- NRC brought up the issue of the wideband (2 GHz) autocorrelator and how many correlator chips per autocorrelation product will be required to perform the task. It is not possible to produce a seamless wideband auto-power spectrum by concatenating sub-band auto-power spectra because of aliasing. Thus, wideband autocorrelation must be obtained before sub-band filtering. However, since the data is in time demultiplexed form, many autocorrelations (see NRC-EVLA Memo #001) must be performed and this is somewhat of an onerous task. NRC suggested that since the wideband auto-power spectrum is for diagnostic purposes only, that perhaps a lower-SNR auto-power spectrum could be provided with just one correlator chip per result. This would result in a factor of 4 reduction in SNR. NRAO agreed that this is an acceptable trade-off. NRC confirmed that one correlator chip (as currently envisioned: 16, 128-lag cross-correlators, and assuming a time-demux factor of 16) will yield 1024 spectral points across the wideband with an SNR degradation of 4.

  Action: NRC will provide 4 wideband auto-power results per Station Board. Each one will use a minimum of one correlator chip and provide 1024 spectral points with an SNR degradation of 4.

  NRC also pointed out that, depending on the agility of the analog system, it is possible to obtain the wideband auto-power spectrum (with sub-bands) by using two different basebands, each with a different LO offset, but tuned to the same part of the sky. This requires that the correlator chip have the routing resources to perform the necessary cross-correlations.

10 Sub-array Issues

- NRAO requested that the restriction imposed in NRC-EVLA Memo #001 section 3.1.1 be removed. That is, different baselines within the same sub-array should be able to be programmed to do different things.

  Comment: NRC sees no reason why this request cannot be accommodated and will endeavor to design the hardware accordingly.

11 Support for VLBI Correlation

- It has become obvious to many that the architecture of the correlator contains most of the elements of an XF VLBI correlator. For NRAO, this is attractive since the correlator design could be used for a new VLBA correlator should funding for such a system become available in the future. To be fully VLBI ready, the Baseline Board (and mostly the correlator chip) must contain the functions of baseline (fine) delay calculation and tracking, and the so-called phase-modifier calculation and application.

These functions are well documented in [1]. The Station Board contains elements suitable for a VLBI correlator such as the large geometric delay buffer and quantizer statistics accumulation but it is missing elements required for real-time tape recorder control. Presumably, these functions (i.e. VLBI recorder interface) could be plugged into the board in the place of the fiber-optic receiver module as long as it has access to all of the necessary signals on the board.

There are several ways in which VLBI could be done with this correlator:

1. Record (narrow) analog-generated baseband data as is currently done with VLBI systems. Each Station Board would be able to handle 32 of these narrow basebands (or 16 baseband pairs). The FIR filters could be used for narrower spectral-line work or could simply be bypassed. This operating mode would be compatible with all existing VLBI systems. This requires that the "top" of each FIR filter have access to all input data streams since each one will be operated in serial mode (NRC-EVLA Memo# 003).

2. Record wideband demultiplexed data and present it to the correlator on playback as such. The digital sub-sample delay tracking and sub-band FIR filtering would operate as usual. A requantization step is required, but since the correlator is 4 bits, it would incur a negligible SNR loss in addition to the loss from initial 2-bit quantization. In this mode, it is possible to reconstruct the entire wideband cross-power spectrum so that the traditional VLBI problems of "lining up" analog sub-bands is not encountered.

3. Put the Station Board at the VLBI antenna where sub-bands are generated before being recorded. This essentially constructs a digital baseband system but with information (lag 0 autocorrelation before requantization—acquired at the antennas) required to seamlessly reconstruct the wide band. In this case, the Station Board could be used in the correlator, but if so, most of the functions would be bypassed.

   Action: NRC will design the Baseline Board and correlator chip so that they are VLBI compatible (i.e. fine baseline delay and phase modifier). NRC will endeavor to design the Station Board so that it has all of the "hooks" necessary to be configured for any of the above VLBI operating modes and so that a VLBI recorder interface can be plugged into the daughter-board space where the fiber-optic receiver board is normally installed.

12 RFI Mitigation

- The suggestion was made that perhaps the sub-band lag 0 autocorrelator before requantization could be used as a fast "time-burst" interference detector which then could be used to blank data valid. This is instead of the "FAST interference detector" currently on the Station Board block diagram (in NRC-EVLA Memo# 001). This function could probably only be performed if the AGC (ALC) in the initial quantizer has a long integration time or is not active.

Action: NRC will consider including this function in the design of the Station Board (perhaps integrating this as a programmable function in the FIR filter chip).

- Some comments were made as to the requirements of the sampler boards. The sampler boards should include a quantizer state counter, step attenuator (into the quantizer), an AGC controller, AGC blanking for burst interference rejection, and WALSH function "deswitching". (A post-meeting suggestion by NRC is to also include a variable-gain broadband noise generator whose output gets added to the signal before quantization. Tests indicate that when narrowband interference amplitude is such that resulting quantization noise is higher than the broadband noise level in the signal itself, harmonics and intermodulation products in the quantizer output result. Most of these by-products are removed by the WIDAR frequency shifting technique—but some of them remain because of the frequency shifts of the intermodulation products. Harmonic generation can be prevented by adding some broadband noise to the signal before quantization (a well-known technique), and the resulting SNR degradation due to added noise is no more than with not doing it. A survey of VLA interference monitor data seems to indicate that interference levels are not yet at the level where this is problematic—but nevertheless, the function might prove to be useful at some point in the future.)

- NRC stated that one way of further suppressing quantizer generated harmonics is to perhaps use a 7-level phase rotator/fringe stopper in the correlator chip since it presumably has a lower 3rd harmonic amplitude than a 5-level fringe rotator. This could also eliminate the need to perform 4-bit initial quantization—3-bit initial quantization could save in FOTS (Fiber-Optic Transmission System) costs and, with a 7-level phase rotator, would yield equivalent performance. A 7-level phase rotator would impact the cost and perhaps number of lags in the correlator chip and so this option would have to be carefully considered. (It has subsequently been found that the 3rd harmonic amplitude of a 7-level function is about –20 dB relative to the fundamental. However, it now appears that some intermodulation products of narrowband signals are not attenuated with frequency shifting and so a 7-level phase rotator would not yield beneficial results especially when coupled with 3-bit initial quantization. Intermodulation effects from narrowband quantized signals is the subject of an upcoming memo.)

  Action: NRC will consider using a 7-level phase rotator. It is currently NRC's opinion that the initial quantizer should be 4 bits rather than 3 because of the intermodulation effect mentioned above.

- There was some informal talk about the Australian (M. Kesteven) post-correlation adaptive interference cancellation technique and how it might be included in the EVLA. NRC has since learned more about this technique and its impact on correlator design. The technique looks very effective, has no impact on the current correlator design, and does not get applied or change any data until post-processing (and only then if a user decides to apply it). There is some requirement for additional hardware

for the EVLA and this technique should be a topic of discussion in the next NRAO-NRC meeting.

13 Array Timing Issues

- All array and correlator timing will be based on UTC.

- It was agreed that there would be some sort of timecode that will be generated at the antennas which will then follow the data and appear with it at the Station Board. The exact format and mechanism of this timecode requires further definition—however, to be compatible with the correlator's TIMECODE, it should probably be a 1 Hz time tick. The generation of the antenna timecode should be such that, once calibrated on the sky for delay relative to the correlator's TIMECODE, antenna power failures and/or interruptions in the FOTS should not require recalibration.

  Action: NRC will consider the impact of this antenna timecode and design the system accordingly. Namely, the fact that the phase of the antenna timecode can be arbitrary relative to TIMECODE and that this phase will be calibrated on the sky.

  Comment: Phase-II of the EVLA (New Mexico Array) will most likely use a public-switched FOTS. In this case, fiber delays can be arbitrary and it is impossible to distinguish the fiber delay from the delay of the antenna timecode epoch to the UTC epoch (at the antenna). Thus, it may be necessary to include GPS receivers at these antennas so the antenna timecode epoch is at (or very close to) the UTC epoch.

- It was agreed that array timing is to be based on VLBI standard frequencies. That is, sample clocks are: 1 MHz * 2^n (n is an integer). This ensures that the phasing subsystem can be fully digital—eliminating the need to produce analog and then resample it. Thus, the nominal wideband sample frequency is 4.096 GHz, and the highest clock rate in the correlator is 256 MHz. It was NRAO's opinion that these frequencies would fit within the planned 10 Gbit/sec FOTS.

- The geometric delay memory on the Station Board must also compensate for the delay through the FOTs. It is expected that the FOTS delay will be a factor of two times the air-path delay (except, perhaps for the New Mexico Array antennas which may need external packet assemblers and delay boxes). Thus, the delay memory must compensate for three times the geometric delay. With a maximum baseline of 300 km, this amounts to 3 milliseconds—or 768k words of delay at 256 Ms/s.

  Action: NRC will design the geometric delay memory on a daughter board that plugs into the Station Board. If longer baselines are one day included in the array, then this daughter board can be upgraded.

- There was some discussion as to the impact of requiring Local Oscillator (LO) offsets in the antennas. This does complicate the design of the antenna electronics somewhat

but it was not seen to be a problem. In order to accommodate narrowband radar mode, the LO system will be designed to be capable of introducing offsets with a 100 Hz resolution (i.e. at a 30 kHz bandwidth, and with 40 antennas, a maximum of 4 kHz of edge bandwidth is lost because of the offsets). Additionally, the LO system will be designed so that a power cycle at the antenna will still keep the same phase of this offset frequency (i.e. so that recalibration on the sky to obtain this phase is not required).

14 Fiber-Optic Transmission System (FOTS)

- NRAO will design and install the FOTS for the EVLA.

- The FOTS nominal carrier bit rate is 10 Gbits/sec. Thus, each carrier will accept two bit streams at 4.096 Gbits/sec each for a total of 8.192 Gbits/sec. To get up to the 10 Gbits/sec rate, this 8 Gbits/sec stream will either be 8B/10B encoded (guaranteeing 1's density and zero bias) or it will be external padded with no encoding. 8B/10B is the most desirable since it guarantees best link operation, but it (8.192 Gbit/sec user data) may not fit within the nominal 10 Gbits/sec carrier bit rate. External padding should guarantee link operation and it should be possible to stay within the 10 Gbit/sec rate.

- All data from one antenna is WDM (Wavelength Division Multiplexed) onto one fiber. Thus, one fiber will handle 2 x 10 Gbit/sec x 8 = 16 Gbits/sec if 4-bit sampling is used.

- WDM demultiplexing is done before the Station Boards. Therefore, each fiber going into the Station Boards contains only one carrier frequency and 2 bit streams. Thus, there are 4 fibers going into each Station Board (fiber receiver daughter board)—2 for each baseband.

- NRAO stated that a clear definition of the FOTS receiver daughter board will be available at the detailed design stage.

  Comment: In order to provide compatibility for VLBI correlation, additional signals may be added to the FOTS daughter board interface. Thus, it might be useful to include NRC in the definition of the FOTS daughter board interface.

15 EVLA Implementation/Transition Plan

- ~2.5 years after t=0[1], upgraded antennas start to come "out of the barn" at a rate of about 4 or 5 per year. Thus, if t=0 in 2001, all 27 antennas will not be upgraded until about 2008 or 2009.

- Irrespective of whether the expansion goes ahead or not, the MODCOMP computers at the VLA site must/will be replaced by 2004.

[1] t=0 is when funding starts—for NRAO or NRC depending on the context of the sentence.

- Due to the computing environment, it is not possible to handle any upgraded antennas until mid 2004.

- At least two NRAO software people will work on correlator software design and implementation starting at about 1 year after t=0 (for NRC). NRAO wants to have one of these software people in Penticton for a year or more working on software interfaces, CORBA, and getting involved from the "ground up" in the design of the system. The goal here is to help speed the development of the correlator as well as eventually have an NRAO person in Socorro who is intimately familiar with the correlator system. The second software person would (presumably) work on higher level software or back-end data handling.

  Comment: NRC welcomes the addition of an NRAO software developer to the correlator development team at Penticton.

- A straw-man correlator development schedule was created by NRC during the meeting since no schedule (other than a rough estimate of total project completion time) has yet been defined. This schedule is as follows (<u>all times are after t=0 for NRC</u>):

  1. 6 months to 1 year – all NRC engineering staff for development in place. NRC noted that core engineering staff should be moved from existing/completing projects onto the correlator project as soon as possible after t=0. NRC would have to hire at least one additional hardware engineer.

  2. 2 to 3 years – 1st correlator "alpha" engineering prototypes should be available. It is probably safest to assume that this will take at least 3 years. Some prototypes will be shipped to NRAO so that they can be used for software testing.

  3. 3 to 5 years – full production and shakedown of the system. During this time, boards will be produced, burned-in, tested, and shipped to the VLA site for installation and testing—probably on a rack-by-rack basis. The initial installation, testing, and shakedown will occur in the "basement" of the operations building. During this time, data from a few antennas should be available for testing.

  4. 6 to 7 years – final installation and system debugging. The correlator will be moved to its final location sometime during this time period once it is found to be sufficiently ready. This move will involve a shutdown of the VLA for a week or so while the system is moved and all antennas are connected to the correlator. At the end of this time, the system should be on-line and fully operational. This time frame meshes well with the antenna upgrade schedule.

     Comment: The exact transition plan has not been fully worked out. Depending on correlator delivery timing, it may require plugging old antennas into the new correlator (a plan for which is outlined in NRC-EVLA Memo# 001).

16 Miscellaneous Issues

- There was some thought that the LO offsets eliminated the need for Walsh function switching. It was agreed that Walsh switching should be included in the antenna design just in case there is some subtle effect that requires it. Walsh deswitching will occur on the sampler board (initial quantizer).

  Comment: This has no impact on the correlator design. Mechanisms for Walsh deswitching will NOT be built into the correlator.

- There were some comments about the deficiencies and complexity of the nomenclature for observing modes defined in NRC-EVLA Memo# 001. It was agreed that a better nomenclature should eventually be defined (i.e. for more general use) but for now, the existing nomenclature is useable (mostly because nobody has come up with a better one!).

- NRC wanted to know if NRAO is satisfied with the proposed –48 VDC mains power system distribution. NRAO indicated that it seems like a good plan but Ray Escoffier expressed concern about the longevity and replacement availability of circuit board DC-DC converters. We agreed that this concern can probably be mitigated with proper circuit board design (i.e. install hooks to allow installation of a different footprint power supply if necessary).

- The EVLA analog output will be quasi-baseband from 2 to 4 GHz. Thus, the sampler input must be able to handle signals with frequency components as high as 4 GHz. Sampling will down-convert this to 0 to 2 GHz with inverted frequency sense.

- NRC has confirmed that the correlator can be wired for 27-antenna operation (i.e. before the Phase-II expansion happens), and that doing so will double the spectral resolution on every baseline. 27-antenna and 40-antenna configurations can use the same modules and cables (and numbers of modules and cables).

- NRAO wants NRC to consider providing *at least* two readout controllers on the Baseline Board to meet fast dumping goals for pulsar observations.

  Comment: This should not increase the output memory size but it may require an additional high-speed data connection to the (daughter) board.

- NRAO made it clear that they do not want a reduction in the number of spectral channels, or a reduction in the sub-banding capabilities of the correlator over what is proposed in NRC-EVLA Memo# 001 (with addenda from this memo). As such, NRAO will consider some total EVLA cost shuffling in an effort to achieve maximum scientific benefit from the expansion project.

- NRC has stated that the cost of the samplers is not included in the $10 million cost envelope of the correlator. NRAO and NRC will have to negotiate for the funds to finance the construction of the samplers.

17 Action Items

- NRAO wants NRC to write a memo on cost and descoping options. The memo is to include parameterized cost equations that show cost as number of antennas are reduced, number of sub-band correlators are reduced, and number of basebands are reduced.

- NRAO will write a memo describing the scientific justification for very fast dumping (25 microsecond) for pulsar observing. This memo will be included in the NRAO EVLA memo series.

- NRAO wants NRC to write a memo on how phased-VLA output could be fed back into the correlator and what the design and cost implications are.

18 Conclusions

The July meetings in Socorro have helped to refine the technical specifications of the correlator considerably. There are many minor changes to the correlator design outlined in NRC-EVLA Memo# 001, and this document is the first record of those changes. Thus, the refinements in this document *must* be incorporated into appropriate future design documentation.

19 References

[1] Carlson, B.R., Dewdney, P.E., et. al. The S2 VLBI Correlator: A Correlator for Space VLBI and Geodetic Signal Processing, Publications of the Astronomical Society of the Pacific, 1999, 111, 1025-1047.

20 Appendix I – email of NRAO's Internal EVLA Correlator Discussions

This appendix contains a record of email documenting NRAO's internal discussions held in May 2000, and NRC's response to them.

```
Subject: Re: WIDAR notes; AOC meeting May 3 and 23; to DRAO
Date: Tue, 22 Aug 2000 15:38:09 -0700
From: B Carlson <Brent.Carlson@hia.nrc.ca>
Organization: National Research Council of Canada
To: Dick Sramek <dsramek@aoc.nrao.edu>
CC: Peter Dewdney <Peter.Dewdney@hia.nrc.ca>,
     Ken Sowinski <ksowinsk@zia.aoc.NRAO.EDU>,
     Bill Sahr <bsahr@zia.aoc.NRAO.EDU>, Jon Romney <jromney@zia.aoc.NRAO.EDU>,
     Jim Jackson <jjackson@zia.aoc.NRAO.EDU>,
     Rick Perley <rperley@zia.aoc.NRAO.EDU>,
     Barry Clark <bclark@zia.aoc.NRAO.EDU>,
     Steve Blachman <sblachman@zia.aoc.NRAO.EDU>,
     Michael Rupen <mrupen@zia.aoc.NRAO.EDU>,
     John Benson <jbenson@zia.aoc.NRAO.EDU>,
     Mike Revnell <mrevnell@zia.aoc.NRAO.EDU>,
     John Weber <jweber@zia.aoc.NRAO.EDU>,
     Ray Escoffier <rescoffier@zia.aoc.NRAO.EDU>,
     "Larry D'Addario" <ldaddario@zia.aoc.NRAO.EDU>,
     gvanmoor <gvanmoor@zia.aoc.NRAO.EDU>, tcornwel <tcornwel@zia.aoc.NRAO.EDU>,
     Boyd Waters <bwaters@zia.aoc.NRAO.EDU>, pnapier <pnapier@zia.aoc.NRAO.EDU>,
     Brent Carlson <Brent.Carlson@nrc.ca>

Hi Dick:

I just got back from Jodrell Bank and I had a look at your email.  I will try to answer
your questions and concerns at least to 1st order in this response.  As agreed at the
Socorro meeting in July, I will also write a memo that summarizes all of the discussions
that I had with NRAO people.  I can add this discussion to the memo as well if you would
like.

--Brent.

PS.  Those who are not keenly interested in this topic may want to hit the delete key
now!

Dick Sramek wrote:

> Brent, Peter,
>
> In May we held two one hour discussions at the AOC to bring out some questions
> from the NRAO staff about the proposed DRAO WIDAR correlator.  My notes from
> these internal NRAO meetings are given below.
>
> These questions were more or less aired with Brent during his July 10 to
> July 14 visit to Socorro to review the WIDAR design.
>
> Therefore my notes are now somewhat redundant, but for the sake of completeness,
> I'd like to add them to our record of the WIDAR discussion.
>
```

> Also attached is a copy of a more recent exchange between Larry D'Addario and me
> that raises additional concerns.
>
> We should try to settle any remaining differences when you get back from travel
> later this month.
>
> ..Dick Sramek..
>
> *****
> May 3, 2000
>
> Discussion of Engineering aspects of the WIDAR Correlator.
>
> The discussion took the form of questions or concerns to pose to the DRAO design
> group.
>
> Q1 Clock rate - Is 250 MHz the best clock rate vis-a-vis power dissipation and
> ease of implementation? Concerned that 250 MHz clock on this scale is new to
> DRAO.
>
> Suggest demonstration with test board and array of FPGA and bus LVDS.
> Demonstrate that the DRAO group can build using this technology. Suggest that
> this be done early in the project.

In some ways this clock rate is new to DRAO although we have built systems that use 1 Gbps for transport within a correlator. The speed capability of the correlator is primarily determined by industry device and connector support, signal design and synchronization choice/philosophy, and expertise in timing analysis and circuit board layout. A survey of the industry indicates that device and connector support is sufficient and improving to support 250 MHz system clock rates. Xilinx will soon anounce their "Virtex-2" FPGA devices that are fabricated in 0.15 um CMOS and contain six layers of **copper** metalization. These devices will be even faster than Virtex-E, and Virtex-E is sufficient for 250 MHz clock rates. Many high-speed, high density connectors are available and other devices such as memory are capable (right now) of about 133 MHz clock rates--requiring a factor of 2 demux in certain areas which has already been taken to account in the design. The signal design and synchronization philosophy, I believe, is sufficient. Namely, final synchronization occurs on-chip where it needs to occur and embedded information in signals allows this to happen. Chip facilities such as DLL's also make this possible (i.e. manufacturers are keenly aware of what chip facilities are required to facilitate high-speed operation). Finally, I believe we have the expertise to design high-speed circuit boards, provided sufficiently powerful design tools are available. This may require purchase of an expensive but very capable Mentor Graphics seat so that as much board-level simulation and testing can be done as possible.

Any additional prototyping and demonstration is going to take additional time and push back the delivery of the system. I assert that the best way to ensure success is with powerful design and simulation tools. Certainly, any bussed LVDS design (and I believe you are referring to the Baseline Board design) is certainly going to have to be looked at carefully.

Nevertheless, our fall back position--should 250 MHz design be impossible--is 125 MHz.

>
>
> Q2 Maintenance - Combination of 81 chips on a board using ball grid array may
> be a maintenance problem.
>
> Need a demonstration that chips can be successfully removed and replaced on the
> board. We should gain some experience with a rework station.

Ball grid arrays are new territory but it looks like that is the direction the industry is moving for reliability, manufacturability, package size, thermal performance, and speed. Rework stations (e.g. PACE TF2000) are commercially available. It should not be difficult to purchase a rework station early in the project to get some experience with BGA/FGA rework. (N.B. The Baseline Board design has been scaled back so there are 64 correlator chips on it rather than 81.)

>
>
> Q3 Maintenance – Need on-line diagnostics that can detect a failure and localize
> fault down to the board level in all subsystems. Need off-line diagnostics,
> test fixtures, hardware and software, that will localize fault to the chip
> level; this should be supplied with the correlator.
>
> Diagnostic software should be supplied by DRAO.
>
> Or diagnostic hardware should be provided on each PC board that will localize
> faults to the chip level. Can be run with board removed from system.

This issue was discussed somewhat in Socorro in July. The basic plan is to have testvector/testing capability within the context of normal operation of the system that is able to pinpoint hardware faults to at least the board level (and probably better). That is, a test observation is just another observation--except for a few key hardware settings and facilities--so that all of the S/W facilities that are developed for normal operation are available for testing. Thus, the NRAO high-level software design must take this into account and provide sufficient facilities for on-line (and perhaps) off-line testing. Additional high-level test software can be developed for off-line testing if needed, but I see the off-line test system as just another (but much smaller) correlator where you plug-in your faulty boards to test them. Of course, the size and complexity of any additional test software can grow depending on desired ease-of-use, automation etc. and its impossible to commit to anything that nebulous at this point.

>
>
> Q4 Size and power /cooling requirement
>
> Minimize power usage and floor space. Correlator should fit within existing
> correlator room at VLA. Power/cooling requirements should be given to NRAO as
> soon as possible.

I visited the VLA site and discussed with Rick Perley where the correlator would go and additional space that could be had if absolutely necessary. An estimate of correlator system power requirements has been provided to NRAO. The installation is assumed to be a benign office environment at 23 C ambient.

>
>
> Q5 Should we adopt for the EVLA system design that fringe rotation will be done
> in the correlator?
>
> Tentative agreement that this is a reasonable approach and that the idea should
> be developed further.

This is entirely up to NRAO. The WIDAR correlator contains fringe rotators and so all fringe rotation could be done in the correlator.

>
>
> Q6 A narrow deep absorption line (20 kHz in 2 GHz, 35 dB line depth) will be too
> time consuming to simulate with adequate SNR, but building special purpose
> hardware to evaluate this situation is not practical.
>

> We probably need to accept that this situation will not be tested.

This would be time-consuming to simulate--mostly the operations to generate the narrow absorption line. I can't imagine that the correlator would just not work properly in this case. In the worst case, the absorption line shows up on a sub-band boundary and there is an SNR loss. In the nominal case, the line is in the middle of a sub-band and there is no aliasing or reduced SNR effects.

>
>
> Q7 - Is it acceptable that spectral points at joints will have higher SNR?
>
> Extent of problem depends on the synthesizer step size? Refer this question to
> the science requirements group.
>
> Q8 - Simulations don't use a long enough bit stream to achieve realistic SNR.
> Should a hardware prototype be built to verify that the stitching works well
> enough?
>
> Estimating 6 months for two people to build a prototype for this demonstration,
> plus software development, the concern probably does not warrant the effort.

Again, additional demonstration hardware would push back the schedule. In the July Socorro meeting I explained to NRAO all of the nuances in the design that I had come across. Perhaps NRAO has an interested person who would like to build a simulator and think about the concepts in sufficient detail to see if there are other nuances. Many simulations I ran yielded very high SNR outputs (albeit with high SNR inputs)--refer to NRC-EVLA Memo #001 and April/00 Penticton meeting copies of simulation results.

>
>
> Q9 - Are there strange phase effects at the joints?
>
> Phase effects should be less of a problem than amplitude. Request comment from
> Brent Carlson on this.

I have tested the phase response of the symmetric FIR filters and have found zero phase response well past -15 dB in the transition band. Since the sub-band boundary would normally be chosen to be at a much higher cuttoff than this, then it would seem that there is plenty of margin. The tests I ran looked at the phase of the cross-power spectrum where data from one source went through the FIR, and data from the other source did not. Thus, it tests the absolute phase response rather than the differential phase response which should be even better (if that's possible). Fundamentally it is known that symmetric FIR filters are guaranteed to be linear phase (which becomes zero phase once the fixed delay through the shift register in the filter is removed).

>
>
> May 23, 2000
>
> Today's WIDAR meeting took the form of a rather loose discussion about
> alternative approaches to the EVLA correlator
>
> 1) Has DRAO done a comparative cost estimate to arrive at this particular /16
> multiplex hybrid design. Are there trade-offs of hybrid designs which have
> greater hardware complexity but less software complexity.

I did cost a /32 WIDAR design for 125 MHz clock rate and it is in NRC-EVLA Memo #001 (page 65). It was more expensive mostly due to the higher number of circuit boards (Baseline Boards) required to hold all of the correlator chips. However, for the same spectral resolution, the same number of correlator lags as the /16 design are

required so the correlator chip cost is not that much different. The second statement would be generally difficult to answer except on a case-by-case basis. For example, if we reduced the flexibility, eliminated recirculation, and only used /16 sub-band bandwidths then the software is much simpler but the capabilities of the system are reduced considerably. Generally, flexibility==complexity for the implementors of a system--our challenge is to not make this the case for the users of the system.

>
>
> 2) Introducing new modes will be harder in WIDAR than with a simple XF approach.

Not sure what is meant by this and what exactly is meant by "a simple XF approach". I assert that the presence of FIR filters greatly improves the flexibility of the system and thus the introduction of new modes.

>
>
> 3) What are the cost trade-offs using fewer FIR filters and/or fewer spectral
> channels?

Fewer FIR filters probably does not impact the cost that much unless it is carried to the extreme. Fewer spectral channels will impact the cost and, in the extreme, a WIDAR correlator is not advantageous (i.e. for a pure continuum correlator). My understanding from the July Socorro meeting is that 16384 spectral channels per baseline in wideband modes is now considered to be the *minimum* requirement.

>
>
> 4) What RFI spec is realistic for the VLA? Are the simulations calculated for
> WIDAR realistic, i.e., a few isolated carriers? It was suggested to try
> modulated signals.

I just ran an extensive simulation of 2 isolated carriers comparing 4-bit WIDAR and a simple 4-bit correlator. The reduction in the number, and level of intermodulation products is dramatic. I will include the results in an upcoming memo.

>
>
> An alternative is to record real interference and present this as the input
> stream to the simulation. Or we could synthesize rfi and present analogue
> signal to samplers and record a bit stream.

This could be a useful exercise. It may turn out that the interference is more extreme than I've tested and so a low-frequency, (relatively) narrow-band, 8-bit sampling with FIR notch filtering (sub-)system may be *required*. If this test is done, it would be best from my perspective to do the fringe stopping at the antennas, and then record some data with offset LOs and some without offset LOs so that a comparison can be done. A 4-bit system, even with WIDAR, does have rfi harmonic/intermodulation product suppression limitations.

>
>
> 5) Should the simulations include cases where the total power in RFI exceeds the
> system power level?

This would have effects on the analog electronics that I could simulate but may be difficult to match to what the system will actually do. In this case, if simulation is to be done, acquisition of real data would produce the best results. This test would probably not help with correlator fine tuning but it may give people an idea of what to expect in the correlated output.

>
>
> 6) What assumptions are made in the simulations about digitizer levels? Is the
> presence of an ALC assumed.

In my simulations, I have put errors in the initial quantizer threshold levels but, perhaps, not some errors that could be present (such as different quantizer thresholds for positive and negative going signals). Nevertheless, every FIR filter sees all of the data so any spectral/power effect this has is going to be seen by all filters. The requantizers operate as well as they could with real hardware (i.e. the simulation is exact since the requantizers are digital). I used an ALC--but with a one-time setting which blindly set the thresholds based on the RMS level into the initial quantizer--not optimum in the case where there is one powerful narrowband interferer in the band. As far as I know, WIDAR can operate with or without an ALC in both initial and requantization stages--output data can be normalized provided the statistics are measured and applied correctly.

If there is any doubt, perhaps an interested NRAO person could perform the simulations, think about this in sufficient depth, and verify the conclusions and equations I have developed.

>
>
> %%%
> 2 AUGUST EXCHANGE BETWEEN D'ADDARIO AND SRAMEK
>
> Dick Sramek writes:
> ...
>
> > Q5 Should we adopt for the EVLA system design that fringe rotation
> > will be done in the correlator?
> >
> > Tentative agreement that this is a reasonable approach and that the
> > idea should be developed further.
>
> You do realize, of course, that this implies that the correlator must
> be twice as big for the same performance (bandwidth, resolution), or
> that the performance is halved for the same size, compared with having
> the fringe rotation done in an LO. For this to be considered
> acceptable, you must say what gain is achieved that offsets this
> loss. I do not believe that you can claim a simplification of the LO
> system, which needs nearly the same hardware even if fringe rotation
> is not done there.

Complex correlation is required with WIDAR for anti-aliasing. Also, the complex correlator allows digital sub-sample delay tracking and has the additional benefit of a substantial increase in spectral dynamic range in the presence of powerful interfering signals--equivalent to at least two extra bits (and maybe more) in the quantizer. One could use a simple correlator and take the aliasing hit at the sub-band boundaries but I think that would be a huge mistake for a factor of two increase in spectral resolution--given the number of spectral channels that will already be available and the additional benefits mentioned above.

>
>
> ...
> > May 23, 2000
> >
> > Today's WIDAR meeting took the form of a rather loose discussion about
> > alternative approaches to the EVLA correlator
> ...
>
> The discussion as reported did not include anything about the

> fundamental considerations in choice of architecture. On such
> grounds, the WIDAR architecture does not seem to make sense. It would
> have to be shown that there are practical considerations which
> outweigh the fundamental ones in order to justify the proposed
> approach. Perhaps this can be shown, but so far no attempt has been
> made to do so.
>
> The fundamental considerations to which I refer are these: The
> well-known extremes of architecture are the FX and XF correlators,
> depending on whether the frequency analysis is done before or after
> cross-correlation. But more generally you can have FXF, where part of
> the frequency analysis is done before and part after. The WIDAR
> proposal is in this intermediate category, as is the old "hybrid
> correlator," once implemented as a spectrometer for the 12 M
> Telescope. The WIDAR design improves on the HS by doing the first "F"
> digitally rather than with analog filters, but structurally it's the
> same. The argument for the FX extreme is that it can be much smaller
> (less hardware) than XF for the same performance because the total
> rate of multiplications and additions is less (at least for the case
> of many spectral channels and more than a few antennas). The
> counter-argument in favor of XF is that it has a far more regular
> structure, enabling the (admitedly larger) number of operations to be
> implemented in many identical and simple circuits, with clean
> signal-distribution paths. When the number of antennas is very large,
> the hardware may be dominated by signal transmission ("copper") rather
> than signal processing ("silicon"), favoring the more regular
> architecture. (A secondary argument for XF is that the individual
> multiplies and adds are on few-bit numbers, whereas most of those in
> FX must be many-bit.) For the WIDAR design, it has not yet been shown:
>
> a) that the particular compromise between before and after frequency
> analysis is somehow near optimum. If pre-correlation analysis is
> better, why not go all the way to FX? If not, why not use XF? If
> there is an optimum in between, why is it not 2x or 4x or more
> different from that proposed?
>
> nor
>
> b) that the pre-correlation analysis using FIR filters is better than
> the drastically more efficient FFT. The FIR filters produce better
> control of the frequency response, but is this really needed? In both
> FX and XF architectures, the spectral channel response is sinc(f/B)
> due to the FFT; why does WIDAR need to do much better? If FIR is
> used, how do you make the tradeoff between complexity (number of lags
> in FIR) and filter shape?
>
> In discussions at Penticton, it was suggested that these choices have
> something to do with the practical internal clock rate.

>From a pragmatic point of view, I assert that the practical internal clock rate is what drives everything. Based on fundamentals, FX correlators win hands down, but if one considers the data expansion problem after FFT, the n^2 fanout problem to the correlator boards, the wide bands that are being correlated, the complexity of the FX correlator multiplications, and the number of spectral channels that astronomers want, the choice of correlator is clear. An XF correlator requires minimum bandwidth from the station electronics to the baseline electronics, and with WIDAR, results in an efficient lag-correlator implementation (plus other benefits mentioned above).

There is certainly a finite probability that a more efficient wideband FX correlator architecture exists but I would leave that to FX correlator design experts to prove that this is the case given current technology limitations.

> That's not
> correct. In either FX or XF, any bandwidth and any number of channels
> can be implemented with any clock rate by using appropriate
> multiplexing and/or demultiplexing.

No doubt...the question is, which scheme does it more efficiently? The choice could mean one room of equipment with a 100 kW power requirement or 10 rooms of equipment with a 1 MW power requirement for the same output--not to mention cost.

> This has no direct effect on the
> basic choice of architecture. If there is an indirect effect for some
> practical reason, it has not yet been discussed.

The amount of demultiplexing required and how that can be efficiently processed may have a very direct impact on architecture choice.

>
> These questions should be answered convincingly before proceeding.

I don't think there is any way to answer these questions convincingly to everyone's satisfaction. To do so would require an extensive and exhaustive study with more ideas and input than from just one person or one group. Perhaps interested people from NRAO could undertake such a study. Perhaps a careful look at the Japanese proposal for an FX future correlator for ALMA is the best way to do this. In fact, at first glance as far as I can tell, the Japanese design works like this:

On one chunk of time-contiguous data (samples 0,1,2...n), perform an FFT to produce (0,1,2...n) frequency points. To do this on a stream of high-speed data, it *looks like* "M" FFTs are required (each one handling a different chunk of time-contiguous data) since the output is M*(n+1) words into a "short-term spectrum buffer" in their write up. M is the demux factor according to their terminology. Then, for *each* spectral output point, M frequency points (each one from a different chunk of time) from both stations are multiplied and (using a single register accumulator technique they've developed) accumulated. The outputs of these short term spectral accumulators are then rearranged and concatenated with an output cross-bar switch and go to the LTA to yield the wideband spectrum.

So, it *looks like* M x FFT chips are required for each baseband of each station (analagous to M x FIRs with WIDAR) and M multipliers are required *(for each spectral point—B.C. notes this is WRONG!)* for each baseline (ignoring a factor of 2 that may be there because of
required time-overlapping). Thus, it would seem that for wideband operation with a demux factor of M, essentially M, FX correlators all operating at sample rate/M are built. It doesn't look like they are using an equivalent "polyphase decimating FFT" equivalent to what the WIDAR EVLA correlator is using. So, the WIDAR EVLA design results in an efficient XF correlator independent of bandwidth, whereas, the "X" part of the Japanese FX design is inefficient compared to what a full speed (4 GHz clock rate) FX correlator could do. Of course, one could build an FX correlator with FIR filters and FFTs in the front end and an efficient "X" in the back-end...but then there is still the problem of wideband data distribution due to the increase in word complexity after FFT (as Ray Escoffier has pointed out on several occasions) and fringe stopping somewhere in the correlator would be required. (Indeed, the paper we published on WIDAR does not specifically state that an XF correlator is required.)

>
>
> --Larry
>
> %%
> 3 AUGUST EXCHANGE BETWEEN d'ADDARIO AND SRAMEK
>
> Dick Sramek writes:

> > When I send my notes to Peter Dewdney and Brent Carlson, I'll attach your
> > comments as well.
>
> OK.
>
> > Regarding fringe rotation, they need a frequency shift in the correlator for
> > their scheme of suppressing aliasing between sub-bands. I think this is the
> > main driver for choosing a complex correlator. Using it also for fringe
> > stopping just falls out at no additional cost.
>
> Well, then we can say that the doubling of the correlator size is
> caused by the alias suppression scheme rather than by the fringe
> rotation. But we should be sure to recognize that the architecture
> has an inherent 100% overhead and ask what we have gained in exchange
> for this cost.
>
> ...
> > Many people like the multiple sub-bands for avoiding RFI and the great
> > flexibility of placing the observing channels where they're needed.
>
> First, a pure FX or XF correlator can achieve the same degree of RFI
> suppression for the same initial quantization. If digitizer
> non-linearity causes an RFI line in subband 3 to have a harmonic in
> subband 6, then the harmonic will not be supressed any better by WIDAR
> than by the other architectures.
>
> Second, the observers should specify the tuning flexibility desired
> and we should design to that. It doesn't come for free. In WIDAR,
> each subband is of fixed bandwidth and there are exactly enough to
> cover the digitized bandwidth. While it is true that their center
> frequencies can be tuned (but only within discrete slots) and some can
> be ignored, this does not seem to create any more flexibility than can
> be obtained with a sufficient number of filters in an FX or XF scheme.
> The WIDAR scheme has 8 BB channels, each with 16 subchannels, or 64
> subchannels altogether. I suggest that this is far more subchannels
> than anyone has a use for. The "flexibility" is obtained only when
> less than the full bandwidth is processed; the same is true for an XF
> correlator when only part of each of the 8 BB channels is used, or
> when some of the BB channels are unused and their correlators are
> re-assigned to produce more lags on those that are in use. I don't
> see any huge advantage here.
>
> Cheers,
> Larry

Index

J

L

M

N

O

P

Q

R

S

T

U

V

W