VLA Test Memorandum #179

Report on VLA Technical Issues Seminar #4 held on October 28, 1993

Juan M. Uson

A new VLA correlator

The discussion was led by B. Clark and what follows is my summary of it.

The main reason why we should develop a new correlator is that we would like to observe with more bandwidth. The present bandwidth is about as high as can be handled in a single bitstream. In order to increase it it is necessary to process it in parallel. The cost would be proportional to the bandwidth used. It would seem reasonable to shoot for 2×500 MHz with full polarization capability which would require 2 Gigasamples per second. Full polarization would result in another factor of 2 in the cross-multipliers which would make this correlator ~ 32 times the size of the VLBA correlator.

So what should an ideal correlator look like? Its block diagram follows

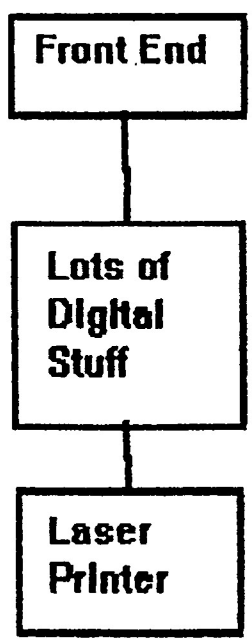

At some point, FX correlators become cheaper than lag correlators. However, although FX correlators involve fewer operations per second, the operations are simpler with lag correlators. The cost factor might actually end up being a toss-up. The FX correlator might be preferable because of its conceptual simplicity.

So what should it actually look like? It would seem desirable to simplify the vertex room which now contains mixers, synthesizers, fringe and timing generators, modem,..., because we presently need to choose where to extract the 50 MHz bandwidths that are sent to the correlator. The new vertex room would implement the following block diagram:

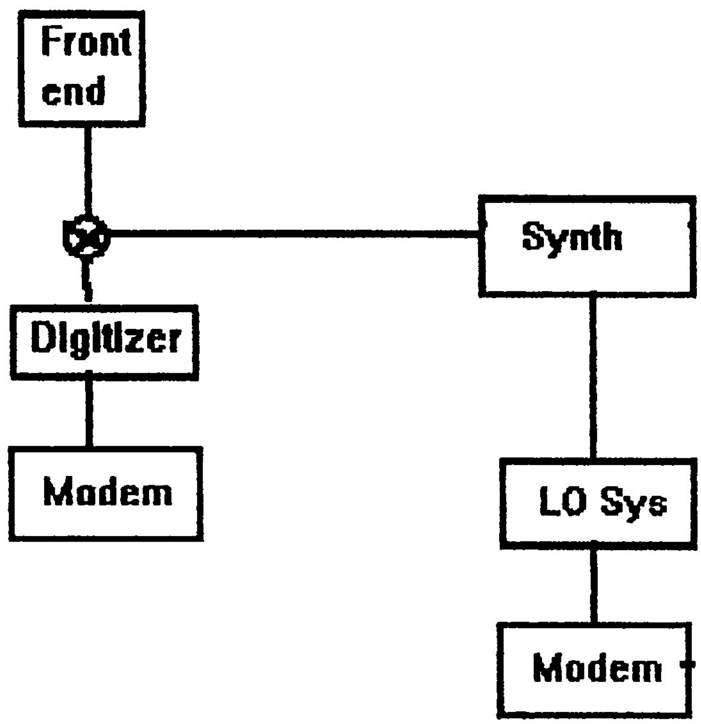

and would contain one synthesizer, one digitizer, one fiber modem, an LO system and a second fiber modem. Multiple fibers would be used (2 GHz sampling at 4 bits might require 4 fibers). One more fiber would be used for analog transmission of the LO. Control information would have to be sent as well bringing the total to perhaps 8 fibers per antenna station. The cost of the fibers and modems might be of order $2M.

The digitizer is an interesting device. Indeed, if it used four bits it might not need an ALC (Automatic level control) device, just an occasional 3dB attenuator (for example for solar observations).

The correlator itself could perhaps use a hybrid design. For example an FX block per antenna might be used to get 32 MHz filtered data streams that could then be fed into a lag correlator. There are algorithms other than the FFT which still involve a number of operations proportional to N logN which should be explored (most might be easily discarded).

There are advantages to the FX correlator. For example, it does not require an explicit fringe rotator, as this operation would be done by the FFT engine.

J. Uson asked if this would not be necessary anyway in order to accomodate the nearby VLBA antennas (Pie Town as well as planned antennas such as those to be located at Dusty and Bernardo).

B. Clark responded that was not necessary as those antennas could be easily retrofitted with fringe rotators with speeds up to $\sim$ 1 MHz.

Possible channel widths should be as low as 8 Hz in order to accomodate planetary radar experiments. This would be sufficiently narrow channels that no delay system would be necessary for an earth-based interferometer. This could be done by using multiple stages (in order to get down from 1 GHz to 8 Hz). This would involve 27 binary butterflies. Two additional butterflies would be used for a fringe rotator and a delay compensator.

The IFs could be processed with parallel FFT engines. The VLBA FFT chip can process 32 megasamples per second. Therefore, 64 engines would be necessary to process two 500 MHz bandwidths for each station. The VLBA chips do not have enough on-board memory so external memory would be needed. This adds unwelcome complexity but not a significant cost as it would only need $\sim$ 100 Mbyte of memory (the memory required in the system is about twice the ratio of total to synthesized bandwidth).

An additional change would be needed. The VLBA processor is a floating point device but the exponent can only accomodate 12 bits while 13 bits would be needed here to manipulate the complex numbers.

The processor would then produce too many channels which could not be dumped at the required rates, so they would have to be averaged in groups for all but the narrowest-band applications. This would be done in the multiplier/accumulator. It would be straightforward to step the accumulation to the next output channel in order to get any nice set of non-overlapping square bandpasses which could be easily convolved in order to avoid ringing (of order 100000 channels could be averaged). Wider channels would also have to be provided for special purposes (like pulsar gating), which could be done by bypassing FFT stages. Ultimately, the number of output channels would be limited by how fast they could be dumped onto tape, perhaps up to a total of 100000 channels per second.

J. Uson remarked that the on-line system could be fed a larger number of channels in order to allow for some on-line processing such as needed for RFI excision, with subsequent averaging prior to writing to tape.

B. Clark indicated that it would perhaps be practical for the on-line system to handle an extra factor of 20 above the data-rate onto tape.

Requirements for the multiplier/acumulator (MAC)

A full multiplier requires four multipliers per FFT engine for polarization. In order to process a full bandwidth of 500 MHz this results in 4 $\times$ (32engines) $\times$ (351baselines) or about 45000 multipliers. They are essentially quad FFT butterflies. This number is comparable to (15stages) $\times$ (2IFs) $\times$ (32engines) $\times$ (27antennas) or about 26000 chips.

This are quite a lot of chips to package. The 70000 chips would require more than 2000 printed circuit boards. The 64 FFT engines for each antenna should still fir in one rack, so that part of the electronics room would not change much. The MAC would likely require 15 racks, somewhat larger than the present VLA correlator. Since it would all be fast digital electronics, shielding it would be a problem.

What would it likely cost? Some wild guesses are:

Fiber-optic connections to all stations:	~\$1Million, installed.
27 modems for 8 Gbits/s:	~\$1.6Million.
Vertex-room electronics (mainly the new LO systems and samplers):	~\$1.6Million.

The FFT/MAC would require:

70000 arithmetic chips @ \$150 or about	\$10Million.
"Glue" estimated at 10%:	\$1Million.
Power supply, racks, other hardware:	\$0.5Million.

And the design, assembly and testing would require about 20 person-years.

Discussion:

C. Walker: Will chips get better and/or cheaper?

B. Clark: Yes, perhaps by a factor of 2.

F. Owen: What about not doing polarization at full bandwidth?

B. Clark: That could save about one-third of the cost although it would make the design more complex.

R. Sramek: It would be necessary to send timing stamps.

B. Clark: It would be necessary to send an occasional bit in order to measure the travel time.

R. Sramek to M. Kesteven: Are you happy with the implementation at the Australia Telescope?

M. Kesteven: Yes, but there remains the complexity of having the fringe rotator at the antenna. The question of aligning the data streams is non-trivial. The AT solution is robust, but had a few teething troubles. (Due to the fact that we solve the "fractional-bit" problem by tying the samplers to the fringe rotators. Barry's stripped-down racing model would avoid these).
Nonetheless, you need to align the data streams to within one sample. The problem arises if you interrupt the stream, then restart - you don't want to have to do an astronomical determination of the delays; the AT in effect has a counter at each antenna driven from the LO signal - as long as the LO is uninterrupted, the data stream can be restarted without trouble.

P. Napier to M. Kesteven: Are the Australia Telescope still planning to allow setting the bandwidths flexibly?

M. Kesteven: We still intend to do so although it is not done at present. The original AT design had the bandwidth selection (below 64 MHz) done at the central site. However, that implementation proved too difficult/impractical. The present (interim) solution has a modest set of filters at the antenna (with bandwidths of 256/128/64/32/16/8/4 MHz), narrower bandwidths will await a new "central site" solution.

P. Napier to M. Kesteven: Is there RFI from the digital devices picked-up by the low-noise receivers?

M. Kesteven: Yes, especially so from the sampler.

B. Clark: It should be possible to do no operations below 100 MHz so that the generated RFI pattern will have the spikes separated by 100 MHz, leaving enough room to observe in between.

F. Owen: How much can the LO stability be improved?

B. Clark: The new system will not have much effect on the LO stability but will improve the bandpass stability significantly.

J. Romney: Taking the memory off the FFT chips is a major perturbation, on the whole a beneficial one. The limitation of the on-chip memory constrained both the spectral resolution (FFT length) and the spectral dynamic range of the VLBA correlator. On the other hand, I understand that there are engineering considerations, beyond just the number of chips and interconnections, that favor the on-chip RAM.

Cost comparisons between FX and lag correlators are affected more by how the specifications are drawn than by design tuning. Instead of treating the specs too rigidly, one should use them as guidelines in developing the best design each architecture is capable of. These can then be compared on the basis of both cost and performance. The extra effort invested at an early stage will pay off in obtaining a system best matched to the application.

An FX correlator is being built for the VSOP mission. They plan to do a long (32K) FFT on a single wideband input data stream, then select the de- sired bands for cross-multiplication and further processing.

P. Napier: It seems that the drivers are the wide continuum bandwidth as well as the need for narrow channels. What is driving the narrow channels?

J. Uson: It is the Radar experiments but also spectral stability and RFI rejection.

B. Clark: Besides, the FFTs are only one-third of the cost.

J. Uson: How soon would it be desirable and practical to start the development of such a correlator?

B. Clark: That depends on the money!

F. Owen: It would seem best to start with the parts that are not likely to evolve.

B. Clark: Yes, maybe it would be best to start with the R&D on the fiber-optics systems.

J. Uson: We should incorporate some smarts into the on-line processing and some development on the next on-line system should also be started soon.

The following comments were sent to me by D. Bagri after the meeting:
It seems worth to consider whether to use analog transmission of the signals from the antennas to the correlator (I am just trying to be a devil's advocate) because of:

1. RFI over large bandwidths and intermodulation considerations. Is is possible to envision a linear system for analog transmission of signals. For example, a 10 GHz system using a 0.5 GHz bandwidth will produce

inter-modulation products which are at least 26 dB down (probably more).

2. Timing jitters, hysteresis, skews, and sampler threshold variations are all going to contribute to bandpass-stability problems. We should compare these with the equivalent ones for an analog system.

3. RFI due to the fast sampling at the antennas which are totally porous to RFI. Also we have possibility of digitizing at the central location at a very early stage, rather than doing a lot of signal processing before digitizing.

4. If we want the same design for both the VLA and MMA correlators, then the following may be worth considering: For a given overall bandwidth it will be necessary to break it up into several "channels" in order to observe multiple spectral lines simultaneously, especially so for the MMA.

5. I am not sure whether this is a valid point: Generally, optical systems have very good isolators which give more than 60 dB (or better) isolation, compared with RF isolation which is typically of 20 dB. This should imply less severe problems due to reflections at such points where isolators are used.

6. The selection of the wanted portion of the spectrum requires sharp and stable filters in order to reject the unwanted parts and avoid aliasing effects. This might be easier to achieve at a central location whose conditions can be made more stable than those at the individual antennas.

In addition, the following issues deserve some thought:

1. We need to understand the advantages and limitations of the FX and the lag architectures. For example, if we want to use a pulse-cal system, then the FX approach accomodates it more easily than the lag architecture because the fringe rotation can be easily applied after sampling (at the FFT stage). On the other hand the FX approach is more difficult for handling quantization corrections properly.

2. We should also think about phase switching: we will not have the present convenient waveguide dropouts which do provide an easy way of aligning the signals.

3. What about the effect of noise-cal switching on the quantization function? This may be more important for the FX approach; but even in the lag approach we will have signal samples during the noise-on and noise-off phases mixed together, because there are no signal drop outs separating these phases. This might turn out to be an unimportant concern, but it needs to be addressed quantitatively.

B. Clark added the following remarks on D. Bagri's comments:

These remarks, although labeled as remarks on analog transmission of the signals, are actually arguments for a more complex channelization, and have no bearing on how the signals are transmitted, except, perhaps, that complex channelization is a little more cheaply done at a central location, to avoid duplication of synthesizers.

The point about phase switching is interesting. I had been presuming that phase switching would be replaced by sampler state counting and post-hoc correction (wasn't something like that tried at Westerbork?), but it has not been demonstrated to work to the required accuracy.

One of the advantages of the wideband sampling is that quantization corrections become easier. For instance, they go away for line observations. (Even the megaJansky water masers will not raise the correlation coefficient on a 500 MHz bandwidth enough to trigger them.) And they are small enough for continuum sources that the simple approximate formulæ that L. Kogan has derived should do an adequate job. (Possibly with small modifications for slow fringe rates.)

It sounds to me like we need a development project to investigate wide-band samplers.