CALIFORNIA INSTITUTE OF TECHNOLOGY
VLBA CORRELATOR PROJECT
31 JULY 1985

RESPONSE TO NRAO REVIEW OF CORRELATOR ARCHITECTURE

INTRODUCTION

This document is provides responses to questions raised in the "Review of 'Architectural Design for the VLBA Correlator'" by Clark et al., dated June 28, 1985. The Architectural Design was presented in Correlator Memorandum VC041, dated April, 1985. The format includes relevant parts of the Review offset in the text and preceded by a ">" character. Our comments follow, left-justified.

After this Response is discussed and revised, we may issue a revised Architectural Design report.

SUMMARY

>The committee found that the report represents an
>acceptable architecture for the VLBA correlator. It incorporates
>the expected features and covers the necessary ranges of all
>parameters.

Thank you.

> We have ... reviewed the cost estimates dated June 10.
>These are very difficult to evaluate because the basis of their
>calculation is not given (except to say that programmers produce 10
>lines per day). It is also difficult to relate the cost estimates
>to the design, because (except for purchased computer hardware) the
>costs are not broken down functionally. Finally, the cost estimates
>contain some internal inconsistencies (e.g., the PWG labor estimate
>is given in three separate tables as 9.7, 15.7, and 30.4 work-mos.,
>respectively.)

In general, one must remember that estimates provided with the design report are not based on a detailed design. That was to have been undertaken in Phase A2 of our contract. Obviously, cost estimates would have been much refined in that process.

Nevertheless, we agree that cost estimates can be clarified to some extent. See Ted Seling's memo "Basis of the Revised VLBA Correlator Cost Estimate" of July 23, 1985.

MAJOR COMMENTS

> A. The plan to calculate station phase to 4 bits and then
>difference at the cross multipliers is probably inadequate.
>Systematic effects (resonance between fringe rate and phase update
>rate) will cause persistent roundoff error, and this will be too
>large unless more bits are transmitted. We recognize that further
>study of this has been done since the report was written, but a
>clear demonstration that there is an adequate scheme for phase

>transmission is still lacking.

Refer to memo VC044 in which Dave Fort discusses these issues. We do not regard them to be particularly urgent at this time. It is clear that more bits can easily be sent by using higher speed transmission or more wires. The point for this Report is that substantial economies will be realized by cutting the number of fringe phase generators from about 190 per channel to about 20. The communication scheme must be designed carefully in order to meet the overall phase and delay tracking specifications. (See below.)

> B. There are major omissions in the software for
>post-correlation corrections: very slow fringe rate case; very fast
>fringe rate case; quantization correction.

The quantization corrections for two-level and four-level sampling will be done in the TOP. The corrections will be applied to the data before it is transformed into spectral frequency channels.

Very fast natural fringe rates will require corrections for three effects :

a). acceleration in the phase causes mistracking in the lobe rotator phases (<20 deg at 100 kHz, 4 ms rate updates),
b). high fringe rates and narrow video bandwidths mean different portions of the bandpasses are correlated on different baselines,
c). the lobe rotator phases are updated every 0.5 microsecs. At 100 kHz, there will be a saw-toothed phase error pattern that swings between +/- 10 degrees.

All three of the above produce non-closing errors. The phase errors caused by a) will be slowly varying (a turn of phase in 10's of minutes). The a) phase errors can be removed in the Transform Output Processor. The amplitude and phase errors caused by the misaligned bandpasses depend on the complex responses of the individual bandpasses. Since this is information that the correlator may not have, we may be required to calculate and apply corrections for b) in the post-processing software. This will be reasonably straightforward to do. The sawtoothed phase errors c) will swing centered on 0.0 degrees phase error if the correlator model calculates its phases for the correct time. We will, however, need to apply some small amplitude corrections due to the c) phase errors. These can be applied in the TOP.

In the case of very slow fringe rates, we propose to shut off the lobe rotators on the affected baselines and rotate the fringes in the DSP's. The accumulators with the slow fringe rate data will be dumped more rapidly than the rest of the correlator, and the DSP fringe rotation will be accomplished in software. The question of low fringe rates requires further attention. For example, how is the DSP told what fringe rates to use? Can the correlator chip communicate baseline phase downstream to the DSP?

> C. There are major omissions in the hardware descriptions
>which make it difficult for us to assess the design. Such
>statistics as estimated chip counts and power consumption for each
>major subsystem are needed. Similarly, the number of p.c. or w.w.
>boards and the interconnection scheme are of interest.

Refer to correlator memo VC052, Subsystem Estimates.

> E. The DPS Control Interface discussion ignores the
>requirement to provide clear separation of those functions which are
>record-technology dependent from those which are not. This
>separation may be logical and not physical, but it must be definite
>and rigorous...

The presentation in the Architecture Report could be improved semantically. We have always intended to separate these functions.

DETAILED COMMENTS

OVERVIEW

> 1.3 - Is the full flexibility of the crossbar switch
>necessary? This is a rather large switch; how much does it cost?

The latest designs do not include this switch as a separate entity. The forthcoming correlator memo "Revised Signal Switching" by David Fort and Martin Ewing will describe the current switching plan. Basically, the ECA input switching that is required to handle various observing modes, coupled with IF channel switching in the Station Electronics, appears to provide all the capability of the crossbar switch in VC041.

SPECIFICATIONS

> 2.1 - Mode names should have some mnemonic value.
>Narrow-band spectroscopic modes are missing from the tables.

We agree that better names will be needed; we would prefer to let the design settle before going further in this matter.

The mode tables are not meant to be exhaustive, but to give an overview. Each mode listed can be multiplied several fold. For example, mode 6d is the highest resolution mode in the tables (~16 kHz). At least 4 options (additional modes) are available to achieve higher resolution. A tape speedup factor of 2 or 4 may be used, corresponding to recording a 4 or 2 MHz IF band. Playback may occur in 2, 4, or more passes to cover more delay space. Combinations of these options would give frequency resolutions of 8, 4, 2, or 1 kHz while keeping correlation time comparable to observation time. Occasional experiments could employ more playback passes to achieve even higher resolution.

> 2.2 - Is any other mode than "8" not possible?

Mode 8 is only prohibited for 2-bit sampling with more than 5 (or possibly 6) antennas. Twenty DPS units can play back 512 million 1-bit samples for each of 10 antennas or 1024 million 2-bit samples for each of 5 antennas. Mode 8 is the only mode that has this particular restriction.

> 2.4 - The meanings of "delay tracking error" and "phase
>tracking error" need clarification. Are these just roundoff errors
>in the calculations or total errors for the whole correlator?
>Calculations to show that the proposed design achieves these specs
>are not given anywhere in the report. Is the maximum rate 128 or
>256 kHz (cf. p 31)? Is this a "station" or "baseline" rate?

Correlator Memos VC049 and VC050 discuss the phase and delay tracking specifications.

> 2.7 - The 4-level samples will be recorded in ones
>complement, not sign/magnitude. The correlator should use
>sign/magnitude only if there is a good reason to do so, and then
>should perform the necessary conversion. The input data should be
>ones complement.

It may well be that sign/magnitude is more useful for the correlator's multipliers. All our designs are based on this, but it is trivial to make the conversion. From our perspective, this is part of the DPS interface specification, many features of which have yet to be determined. (Note also that the "need" for 1's complement is recording technology dependent.)

> 2.10 - The option to support 2x oversampling at full
>frequency resolution is vaguely expressed.

We would have been happy to ignore the possibility of oversampling support at the chip level. It gains only a few per cent sensitivity for a few kinds of observing. It is difficult to assess the extra cost of providing it, although we estimated $150K in memo VC 047, the option list. Our preference would be to proceed further with VLSI design before stating the actual costs that would be incurred by the extra chip functions.

> 2.12 - What imposes the limit of two subarrays? What does
>the sentence in parenthises mean? We think that support of at least
>four subarrays should be required. [sic]

This comment from the review committee was based on the following passage from VC041:

2.12 Simultaneous correlation of unrelated experiments

* 2 experiments. If two independent experiments involve a total of 10 or fewer antennas (including possible duplicated antennas), they may be processed together in ``full'' mode. If the total is 11--14, they may be processed in ``half'' mode; 15--20, in ``quarter'' mode. The two experiments must be of compatible types. High-resolution spectroscopy cannot generally be processed in parallel with wideband continuum, for

example. In some cases, four simultaneous experiments might usefully be correlated; this capability is considered an implementation option.

A misunderstanding has apparently arisen because we failed to state clearly what we mean by "independent experiments" and "sub-arrays". These concepts should be distinguished. We suggest replacing Section 2.12 by the following text:

"2.12 Simultaneous correlation of sub-arrays and independent experiments

"(a) Sub-Arrays. During normal observations, the Array will frequently be divided into sub-arrays observing different sources, possibly at different frequencies. Antennas will be moved from one sub-array to another as sources rise and set. Sub-arrays may also be used to isolate single antennas or small groups for engineering work. The observe-time sub-arraying will be tracked at the correlator, except that single-antenna sub-arrays will be ignored. All the sub-arrays will be correlated in parallel; the only difference between this and processing the complete array as one unit is that data from baselines between antennas in different sub-arrays will be discarded and will not appear in the output archive.

"The correlator hardware imposes some restrictions: essentially, all the sub-arrays must be configured in the same mode. If this restriction cannot be met, the different sub-arrays must be processed sequentially rather than simultaneously.

"The total number of sub-arrays that can be handled at one time by the correlator is a software parameter: we suggest that 10 is the appropriate number (20 stations can be divided between no more than 10 non-trivial sub-arrays).

"(b) Independent experiments. When the processing of a single experiment does not require the full capacity of the correlator, it will be possible to use the excess capacity to simultaneously process a second, independent experiment (possibly involving the same antennas). The two experiments are independent in that they were (presumably) recorded at different times, and the correlator is not required to synchronize the tapes of one experiment with those of the other. Each of the two experiments may be divided into up to 10 sub-arrays. The two experiments will be controlled by separate command files, and may (if this is thought to be desirable) be monitored by separate operators using separate display terminals.

"The correlator hardware again imposes some restrictions: two independent experiments are compatible if (i) they are processed in the same mode, and (ii) the total number of antennas in the two experiments does not exceed the maximum allowed for that mode (see Section 2.1). [This statement is actually over-restrictive, as two experiments can sometimes be compatible even if they are processed in different modes, so long as they use the same major mode; e.g., one experiment using 10 (or fewer) antennas in mode 1a is probably compatible with another experiment using 10 (or fewer) antennas in mode 2a. Such compatibility questions must be examined on a case-by-case basis until we complete a full analysis of compatibility restrictions.] There is also a logistical restriction:

it is not possible to correlate, say, one piece of an experiment simultaneously with another piece of the same experiment recorded a few hours later, if this would require the same tape to be mounted simultaneously on two different DPS units.

"The specification for the correlator is that it should be able to process at least two compatible independent experiments simultaneously. We do not at present see a need for processing more than two independent experiments simultaneously, and the book-keeping involved would make this unattractive."

> 2.13 - The maximum integration time should not be limited to >10 sec; something like 30 to 100 sec is needed. We understand that >from some viewpoints there appears to be a linkage between this and >a tacit MINIMUM model-switching interval spec. The later should be >the subject of a separate spec, and the linkage should be reduced to >the restriction that a model switch cannot occur during an >integration. Alternatively, multiple accumulators for serveral >models should be considered. [sic]

The maximum time specified applies only to the correlator's output - the archive tape. There is nothing to prevent post-processing doing more averaging. There would be little tape savings if longer averages were performed in the correlator, and providing parallel sets of accumulators would definitely complicate TOP software and require more memory.

Yes, we should probably state the specification differently: The minimum model switch period is 10 seconds and must be equal to or a multiple of the integration time. The integration time must not exceed 10 seconds.

> 2.14 - Along which dimensions can the correlator restriction >to achieve higher dump rates best be implemented in this >architecture?

The dump rate restriction applies to channels, according to the table

Max. Dump Rate	Maximum No. active channels
2 Hz	16
4	8
8	4

(This is equivalent to imposing a maximum aggregate I/O limit on each correlator quadrant. The active channels must lie in separate quadrants for the 8 Hz dump rate, for example.)

Refer also to correlator memo VC051, The Data Concentrator.

> 2.15 - Note that .1% of .1 msec is 100 nsec. Can we really >gate to this resolution? Are these numbers in correlator time or >observing time; i.e., what is the effect of a speedup factor?

The 0.1% spec does not seem reasonable. Instead, we would specify the rate at which on/off decisions are made (e.g., 2 MHz). If, as suggested in Correlator Memo VC048, the gating window is specified in a RAM, pulse period and percent time resolution trade off in a simple way.

> 2.16 - Where will the "Van Vleck" correction be done? There >is no provision later in the report for meeting this spec. Strictly >speaking, "Van Vleck" is the two-level quantization correction; what >about 4-level? Another major omission here is a correction for the >systematic errors of very fast and very slow fringe rates.

See discussion under major point "B", above.

> 2.18 - Since this is not independent of 2.14, the two specs >should be combined.

We may delete 2.18 or simply note that the 0.5 MB/s sustained output rate is compatible with current computer tape technology.

> 2.19 - Make 9 track, 6250 bpi a minimum spec, not an >absolute requirement.

Perhaps NRAO should set this specification, since it is equivalently an input specification for the post-processing system. We put down 9 track 6250 bpi because there is no practical alternative now. The specification may change in the future.

INTERFACES

> 3.2 - The second sentence indicates that author is not >familiar with VLBA MCB nor with EIA standards. The MCB conforms to >RS-422, and also (more relavently) to RS-485. It does not support >global commands, but this would be undesirable because it would >make the programming of subarrays awkward.

The distinction is between the MCB protocol and other, more suitable protocols that might be used. (The wording should be adjusted accordingly.)

We now feel that "global addressing" of DPS units is undesirable, increasing complexity with little or no benefit. We would delete it from the interface discussion.

> 3.2.3 - The critical concept of maintaining strict >separation of record-technology-dependent functions from others has >been ignored. The "align" command is particularly offensive in this >regard. Among the commands dealing with tape-dependent matters, the >tape serial number should be read from the DPS, not sent to it; and >the problem of handling tape reversals needs to be considered (it >probably cannot be handled in the DPS in a way transparent to the >correlator).

We would rewrite this section using more neutral language, such as MOUNT VOLUME xxx on DEVICE yyy, following common computer usage. Some degree of technology dependence may be necessary, however. Tapes do have to be aligned, and tape is a sequential access medium.

MAJOR BLOCKS

> 4.1 - Except for the calibration task, a MicroVAX II might
>do, rather than a VAX-11/750.

This is hardly a criticism, since the MicroVAX II was not announced at the time the Report was issued!

If the MicroVAX development continues toward more powerful, less expensive systems, however, they might be good candidates to replace some or all of the 68,000 processors.

> 4.2 - The so-called "station electronics" is not organized
>by station at all, but rather by channel. Why? For many purposes,
>the other organization would be better. For example, it would be
>easier for an SMP to handle all channels for a few stations rather
>than all stations for a few channels (with most of the calculations
>being duplicated in other SMPs).

We tend to use the term "station electronics" to mean everything that does not scale with the number of baselines. The "channel-by-channel" approach has a major benefit in simplifying interconnections in the correlator. We feel this outweighs the cost of having the SMPs computing unrelated models. Under some conditions, several independent models per DPS would have to be computed anyway. Probably all station parameters would be read into all SMPs, adding a minor cost in memory.

> 4.3.3 - Fig 4-6 uses "CH" when apparently "STATION" was
>intended...

We will change "CH" to "STN" in Fig. 4-6.

> 4.4.2 - Why isn't the phase response of the DSP perfect? Do
>we really buy enough with this to justify its cost? Depending on
>the accuracy criteria used, the improvement in fringe rate bandwidth
>may be much less than a factor of 4, compared to boxcar averaging.
>We would like to have a more detailed description of what the DSP
>would do, along with its separate cost. There is also a concern
>that these filters may introduce map artifacts that are not well
>understood, especially due to interaction with gridding filters.

The basic reason for including a digital filter (the DSP) is to reduce the data rate through the Transform Output Processor and Archive Writer, and reduce the size of the VLBA archive data library. Spectral line observations of H2O masers sources set the maximum data rates that we feel we must support with the full correlator. The spatial structure of the largest H2O maser sources extend nearly continuously over 10 arcseconds; we need to support 22 GHz single-field observation of 10 arcseconds. That means preserving residual fringe rates of up to 0.6 Hz. Since we want to make very high dynamic range maps with the VLBA (at

least 1:1000), the fringe amplitudes should uncorrupted to a level of 1% or better.

The digital filter used in this architecture report is a two-stage half-band FIR filter. We accumulate to 0.125 ms and normalize the visibilities before the filter stages. The accumulation is box-car filtration, and the response at 0.6 Hz is 0.991. The FIR filter decimates the sample rate by a factor of four (2 Hz output). The FIR filter response at 0.6 Hz is 0.995, so the total boxcar x FIR response is 0.986.

The phase response is flat to less than 0.001 degree. The limit of 0.001 degree comes from the test programs that were used. In principle, this filter should have zero phase error.

The concern that the filters may interact with the uv gridding convolutions is certainly valid. We will investigate this by filtering and mapping model data. We expect to have this test completed by the September Design Review Meeting. Benson will also write a VLBA memo that thoroughly describes the filter.

> 4.4.3 - The accumulator/filter/COP tree is not related in
>the report to the EC/ECA/input tree. Why is the COP a 68000; isn't
>it just a switch? In Fig 4-10, it seems tight to push 4 COP outputs
>into one Unibus.

The trees are independent. The COP has a number of functions, controlling correlator logic, performing diagnostics, controlling DMA datapaths, buffering and synchronizing data transfers. More recent designs have a 68000 "data concentrator" between the Aptec and the COPs. It acts as a front end to the Aptec.

> 4.6.1 - It is good that at least two sets of hardware were
>investigated. But no justification was given, nor were the
>considerations mentioned. Was the cost a factor? How reliable is
>the Aptec device? How well established is the vendor?

The criteria used were technical suitability (throughput), availability "off the shelf", software support, and integration with a VMS host. Cost is obviously a factor - general purpose computers were ruled out for their high cost. Aptec gives a list of happy customers, of course, but we have not made a detailed investigation since we are at a rather early stage in the project. "The technology of such systems is evolving rapidly, ... and we shall continue to evaluate alternative systems." We seem to have a few years to let the market shake itself out.

> 4.7 - The plans for clock calibration are much too
>pessimistic. If masers are performing properly, then extrapolation
>of clock drift will be within 5 ns for at least a week; using
>observations of only this accuracy made several times a day would
>probably make the "corrected" data worse. The decision about how
>much computing power is needed for calibration tasks interacts with
>operational plans, and must be made on a project-wide basis.

We would be delighted not to worry about calibration! Geodetic/ astrometric experience, however, seems to indicate that masers can not be trusted to this level. The point of calibration is not to "correct" the data, but to set the correlator windows correctly.

A more interesting question would have been, how is a new antenna brought on-line? How does one find fringes and clock offsets initially? This is an operational problem we have not yet studied.

> 4.8 - It appears that considerable effort and expense will
>be devoted to development of a VLSI chip; we therefore think that
>some effort should be made to design a general-purpose correlator
>chip that would have applications beyond the VLBA. NRAO expects to
>develop many other correlators in the future.

We are considering the possibilities of a generalized design. It is clear, however, that the optimum chip for the VLBA would have input switching, phase rotation, etc. that are specialized to the VLBA. It is likely that the chip would have a "vanilla" mode for other applications. NRAO should state its requirements, if any, beyond those of the VLBA proper.

> The choice of on-chip prescalar and accumulator lengths
>needs to be justified. We suspect that the prescalar is too long,
>espeicially for the pulsar case. [sic]

A number of tradeoffs must be made in selecting counter lengths. Prescaler bits are significantly less "expensive" than accumulator bits in terms of gates, output busing, and I/O pins. We lengthened the prescaler until the SNR penalty was just acceptable (~1%) for the shortest anticipated integration, 125 ms. Shorter integrations (fewer samples correlated) will cause a greater penalty. For a pulsar gated with a 5% duty cycle, the minimum integration would be 20 x 0.125 = 2.5 s to suffer the same penalty. Shorter integrations are allowed with lower sensitivity. This seems an acceptable tradeoff.

The accumulator length, 8 bits, was as long as we felt safe to plan for on a 6,000 gate array. The 6,000 gate size seems to be the most cost-effective for a complex function like ours. (Cost increases substantially for larger arrays.) If the accumulators could be lengthened to ~28 bits, we could eliminated the RAM/adder hardware. Although this would simplify our system, it is unlikely to be cost effective in the present gate array technology. More modest lengthening of the accumulators would reduce the required scanning rate of the RAM/adder, and allow us to build fewer of the RAM/adders. The cost of the RAM/adders is under $100 on a $3,100 correlator board, however.

>The decision to provide an 8 bit
>buffer register is new; why change from earlier plans for serial
>readout?

There was never a serial readout (except in the AT chip). We had multiplexed 3-state counter outputs onto an output bus without holding latches; this would have required correlator blanking during readout. We discovered an economical method of providing registered 3-state

outputs in the LSI Logic gate array technology. These registers avoid the need for blanking at very little cost on the chip.

> The discussion of normalization is quite obscure. How many
>counters are there and exactly what are they counting?

In order that very slow fringe rate conditions be handled properly, we have determined that two normalization counters per 16 complex lags may be required. One counts the number of valid sine-rotated samples, and the second counts the number of valid cosine-rotated samples. These counters would be accumulating the AND of the two validity data streams, multiplied by the sine and cosine fringe values, respectively. If the rotation-after-multiplication architecture is adopted, this would entail having two extra fringe rotators on the chip. The total number of samples in the integration period, also needed for normalization, can be determined a priori from the length of the period and the sample rate.

The optimum partitioning for the correlator VLSI is, according to our analysis, 8 complex lags per chip. In order that all the correlator chips might be identical in design, we would provide two counters on each which could be configured to function as either a sine or cosine counter, or as the sample counter.

ARCHIVE

> B.3 - A significant omission is the sign of the net LO
>frequency (upper/lower sideband flag).

Agreed.

NOTE

The following people contributed to this Review: J. Benson, M. Ewing, S. Kator, T. Pearson, and T. Seling.