Panel Report and Response: EVLA Systems Design Review

Below is the response to the Panel Report on the EVLA Systems PDR. The original text is marked by a ">".

>Part I. Comments from the Panel Review

>   A. Computer

>   Don't repeat the decision made on the VLA design to skimp on computer power in order to buy hardware. Consider a local big computer; Beowulf clusters are ideal.

There is money in the project for acquiring a Beowulf cluster in 2009, in the event such a configuration is clearly indicated at that time.

With respect to monitor and control we will not skimp on computer power. However, the computing power will not take the form of a single, local, big computer. Instead, the Monitor and Control system will use a distributed computing model, with systems for array level monitor & control, antenna monitor & control, correlator monitor and control, and correlator backend processing. A Beowulf cluster is planned for the correlator backend processing. We are now in the process of choosing the processor and RTOS for the module interface boards. These decisions will determine if additional processors are needed to handle the antennas as a whole. We are also examining processors and RTOSes for the correlator board. Each of the 300+ correlator boards will have an M&C processor, and these processors will, in turn, converge upon higher level, more powerful systems for higher level correlator monitor and control. Crates will also be speced for array level control. Initial investigations, now in progress, indicate that a 32 processor Beowulf cluster may suffice to meet the initial requirements for the correlator backend processing.Attention is being paid to the need for redundancy in both the slave nodes and the control node of the Beowulf cluster.

CJ>   Provide time for tests by astronomers during the e2e software development. Periodic targets even without external users would be a good idea.

Agreed.

CJ>   Provide adequate manpower; 65 FTE for e2e may be optimistic.

We continue to track the costing carefully. One point is that the number is comparable to other ground-based projects of similar scope. We expect to have a firmer handle on costs at the end of two iterations, at which point we will hold a meeting of an e2e advisory group meeting. Priorities will then be set to met the 65 FTE-years estimate.

CJ>   The M&C software development seems to be lagging behind hardware design, though the problem may be delay in filling

positions
CJ> rather than having insufficient positions provided for. As a result of the current lag in schedule and staffing problems, the dates for the software CDR and the first test antenna seem optimistic.

This will be very carefully watched. In part, the lag can be attributed to the fact that work on ALMA designs gave the Electronics Division a "headstart" on the EVLA designs. And in part, the lag can be attributed to the fact that the Computing Division was and is understaffed w.r.t. the EVLA effort. The understaffing is not due to insufficient positions, but rather to the time it takes to fill positions. We began with four vacancies. One has been filled, a 2nd was filled, but the candidate later withdrew her acceptance, an offer letter is about to be issued for the 3rd position, and we are corresponding with a possible candidate for the 4th position. Until we are fully staffed, M&C software efforts will continue to lag. We are moving aggressively to fill the vacant positions, while attempting to use the manpower that is available to us to address the most pressing issues.

CJ> M&C system development dates may need adjusting. Design and management for the M&C system under one person may not be realistic. To meet current goals, it may be necessary to divide responsibilities.

Having a single leader of the M&C design effort is extremely desirable. Much of the routine personnel management of the M&C group will be provided by the AOC Computer Division head. We are also offloading some of the management and miscellaneous duties to others, developing leaders in various technical areas to handle some of the design issues, and by choosing candidates to fill current vacancies who have technical lead and management experience. In other words, we are beginning to divide and distribute the responsibilities.

The M&C PDR date has been adjusted. To address the need for earlier progress on the hardware, an M&C Hardware PDR focused mainly on the needs of the antenna monitor and control has been scheduled for March 13, 2002. The M&C software PDR has been rescheduled from February 2002 to May 2002, with formal presentation of the M&C interfaces to e2e software possibly deferred until the e2e PDR of July 2002.

CJ> Definition of the correlator backend computer seemed vague, though apparently the backend is a topic planned for better definition in early 2002.

Correct. A person to handle correlator backend issues has now been hired and reported for work on 1/7/2002. That individual is now developing an initial requirements document and design for the correlator backend.

CJ> A rapid response override mode should be planned for operational software.

This is part of the scientific requirements for the M&C system.

CJ> Although hardware liaisons between the Widar Correlator and VLA site installation have been identified, a counterpart from the software viewpoint should also be appointed.

We have identified two positions which will involve close contact between the M&C software and the WIDAR Correlator development. One is the above mentioned individual who is working on requirements and design issues for the correlator backend. We expect, once efforts are more fully underway, that he will spend several months of each year at Penticton. Eventually, a prototype Beowulf cluster for correlator backend processing will be sited at Penticton. The other position is the individual who will work on the correlator monitor and control. This vacancy is not yet filled. When it is filled, the individual in question will spend the bulk of 3 to 4 years in Penticton working closely with the correlator designer.

CJ> Would more precise RT requirements help in hiring appropriate people?

Our experience with acceptable candidates have not shown this to be a hiring barrier. We feel that we have good job descriptions already in hand. It simply takes time to find high quality people who are a good fit to what is needed.

CJ> Questions on documentation, use, and maintenance of Glish should be resolved.

Glish is well-documented and maintained. Only some minor and new features have lagged in documentation but this is not a major problem. It would help to have another developer be cognizant of glish internals. We expect to address this soon in the AIPS++ project.

> Provide time for tests by astronomers during the e2e software development. Periodic targets even without external users would be a good idea.

The spiral development model calls for tests at the end of each cycle:
so every 9 months.

> B. Hardware
>
> All work necessary to meet project requirements such as the upgrades to servo and encoders, for example, should be included in the Project Book.

Agreed.

> Keep the MIB (Module Interface Board) simple; don't make it >too large and require it to do too many things. The MCB board for the >VLBA had more power than necessary for most modules; a smaller, >simpler unit might have been more widely used. Don't feel constrained >to make the same MIB work for both correlator and antenna sicne the >correlator MIB will require many more features than needed at the >antenna. In any case, parsing of tasks between antenna computer and MIB needs further definition.

To be addressed in MCB PDRs.

> Add 8-bit ADC case to DTS documentation.

We will add information to the DTS documentation as it is developed.

> The analog option for data transmission from the antenna to the CB may not have been adequately explored, according to some reviewers. Analog transmission would obviate the need for FIR filters to reform the IF at the control building during the transition phase, eliminate the need to decimate low frequency data at the antennna, reduce RFI and complexity at the antenna, and save money. The PDR presentation did not show a thorough study of the analog option. Others predict an extreme effort to adequately test the analog transmission possibility sufficiently to make a sound judgement on the analog-transmission feasibility. They argue digitizing early in the signal chain has advantages that far outweigh the disadvantages. However, if the digital fiber-optics solution is chosen, then the self-RFI from the samplers, the digital equipment, and fiber transmitters in the antenna must be properly suppressed. There is some help from the LO offsets and removal in the correlator, in that RFI that gets into the signal path *after* the LO offset is introduced, will wash out. This implies that the offset be introduced as far up the LO chain as possible, and that adequate shielding and filtering be provided to prevent leakage into the analog signal before
the LO offset.

Some analysis was performed (mostly by Ron Beresford and Dan Edmans) early in the ALMA project. It was determined at that time that the gain flatness and phase stability of an analog link of this bandwidth was not of sufficient quality to meet the specifications of the ALMA project. When we began looking closely at the EVLA, it was realized that the EVLA's specifications were sufficiently close to ALMA's that the analog solution was probably not a good choice for EVLA either. In addition, substantial investment in time, money and manpower had already been made in the AMLA design and it did not make economic sense to design an entirely different system for EVLA, especially given the amibitious schedule of the project. We believe that we can deal with the RFI issue at the antenna. We also don't believe it is appropriate to compromise the performance of the final EVLA system to satisfy transition requirements - hence why we have essentially ignored the transisiton in selecting the best IF transmission technique.

>Plan contingencies so that if the ALMA sampler is not available by a certain date, there will be a recovery plan to develop a sampler through some ohter means. Plan RFI suppression based on estimated emissions from the sampler so that RFI can be kept below harmful levels.

We plan to design and build an 8-bit sampler in-house. Production of the ALMA sampler seems safe since the schedule is driven by a much larger project.

> The Ka-band system will work with doublers if a high-side LO is used for signal frequencies from 26.5 GHz to 32 GHz, and a low-side LO for signal frequencies from 30 GHz to 40 GHz. Admittedly this method results in a minor problem after the band 29 GHz - 33 GHz, which then cannot be converted to the IF in one piece, a problem that does not occur with triplers. However, conventional wisdom is that doublers are easier to use than triplers. Thus it may be best to use doublers for Ka-band and accept the small problem in the tuning.
The decision should depend upon performance tests of available doublers and triplers. In any case, the first LO in the Ka-band front end is shown in the block diagram as using a doubler

where the Project Book text describes a tripler. The two documents should be consistent.

We had an error in the Block Diagram that has since been corrected. We indeed planned to use a Tripler in the baseline plan but were prepared to consider a Doubler if it turned out to be cheaper. As it turns out, if we buy them from Spacek, the Ka-Band Tripler/Mixer costs the same as the Doubler/Mixer ($5,300 each). Triplers do simplify the conversion scheme as noted, and are used at 40-50 GHz successfully.

> Isolation of the IF channels carrying signals from cross-polarization outputs of the feeds should ideally be no less than the (bandwidth x averaging time)**1/2 factor. For the EVLA, the >isolation should be ~15 dB greater than for the VLA. The difference should be borne in mind while checking isolation, for example, between the two first mixers which are fed by a common LO power splitter.

We have added isolators on the LO drive to each mixer in the lower frequency receivers. On the higher frequency receivers, which have Doublers or Triplers, we haven't since the the multiplier should provide us with a fair bit of extra isolation. So any signal trying to leak from one mixer to the other will see at least 60 dB of isolation thru the LO path (ie: mixer + isolator/multiplier + splitter). Rick Perley is working on a polarization specification.

> The plan for time synchronization needs further definition.

To be addressed in LO/IF design.

>An overall plan for RIF mitigation/suppression/excision should be developed. RFI mitigation will have a critical impact on the reachable sensitivity of the EVLA.
Agreed. Draft plan prepared.

>If RFI mitigation will be done in "post-correlation" software, how will this be done? Does this work require research? Are resources (FTE) available to study this problem, etc.?
See comment under Part II A.

> The impact of the E-array configuration, especially on the fiber optic cable layout, needs to be understood and clarified.

We have identified an area around the center of the array where the E array might be located. We will route the fiber in such a way that the impact of the E antennas on the fiber will be minimized. This will be done by routing the fiber where E antennas can not be placed. For instance the cable will be run close to the present antennas.
Extra conduits will be installed under the tracks and into the building to allow for the additional E array fiber cables. Space will be left for additional termination panels and patch panels.
> Why are Tsys estimates for the new receivers no better than 10 years ago?

The table below comparing EVLA and VLBA shows decent improvement.

Rx Band	VLBA TRx (K)	EVLA TRx (K)
-----	--------	--------
Q	56	44
Ka	-	38
K	61	28
Ku	34	21
X	35	21
C	24	16
S	13	11
L	17	9

Tsys estimates were (and still are) preliminary. Tsys is made up of many factors, of which LNA noise temperature is only one, and usually not the dominant component. Improvement in LNA performance will have only a small effect on system temperature. Improvement in feed performance, particularly at L-band, will also help, but was not considered in these estimates.

> The pointing specification needs to be qualified for conditions; e.g, windless nights, when the specification will be met, worst case. The specification should be the root sum of the squares, IMHO.
Noted.

> Specify amplitude stability required for receivers.

Rick is preparing this spec now.

> Specify acceptable closure errors.

See Chapter 2 of Project Book.

C. General

> Define key interfaces between tasks for both hardware and software; e.g., between correlator, M&C, and e2e. Formal interface documentation should specify tasks and responsibilities. Specify requirements for interface documentation and set up target dates.
The Systems engineer plans to provide an interface document by 3/15/02.

>Consolidate requirements document in the Project Book Systems Chapter.

A separate document will be developed instead of adding the information to the Project Book.

>Specify EVLA observing modes, or better yet observing scenarios, that are required or most likely to be used. Doing so will help designers know what they are doing, reduce ambiguity, and help identify required software tools.

We will update requirements in this area already reported in EVLA Memo 15 on Scientific Requirements.

>Specify goals for test antenna more specifically.

A Fiber optic system will be installed to the test antenna that is fully functional. This includes a 12 channel Digital Transmission System, a fully functional MCB using 1Gbit/100Mbit Ethernet, an LO distribution system with a Round trip phase measurement system. Test fiber spools will be used to simulate the furthest antennas (22 km).

In the control building, the termination panel will be installed and functional for the test antenna by December 2002.

A temporary IF patch panel will be installed in the Present correlator room. This patch panel will have enough equipment to run two test antennas into the old Correlator by June 2003. Additional capabilities will be added as more EVLA antennas come on line.

The LO patch panel will be located in the present electronics room. Again, this patch panel will have enough equipment to run two test antennas by June 2003. Additional capabilities will be added as more EVLA antennas come on line.

The MCB patch panel will be located in the present correlator screen room. This patch panel will be installed in its final configuration and will provide MCB to the test pad by April 2003.

>Develop more specific plans and dates for the transition tasks. The overriding problem with the transition plan is the length of >the schedule though budget constraints presumably limit the rate of progress.

Each arm of the EVLA will be wired for fiber operation in the following order, East, West, and North. The only exceptions are the two test antenna locations, which will be cabled first. This plan may allow an arm of the waveguide to be decommissioned when there are enough ELVA antennas. (ELVA antennas can not use the waveguide, they can only be operated by the fiber system.)

>A more rapid completion of the project would be a clear advantage to astronomy.

The schedule is dictated by budget considerations.

>Part II. Comments during the open meeting review presentation.

> >A. Science requirements

> >-- Include time resolution requirements as well as space resolution requirements?

These have been added to the science requirements --

Section
2.1.13.1: time resolution requirements better than 100 microsecond, plus absolute time accuracy of 10 nanoseconds.

> >-- What is the specification for frequency range overlap and flatness across band? What is gain/sampler stability at band edge?
>
This detail hasn't yet been considered. I don't understand the connection between sampler stability and the band edge. I'll have to ask a knowledgeable person.

> >-- Do you have performance tests to show the antenna can be operated at the 2.5 deg/min tracking speed to support OTF?
>
Sort of -- Barry and Ken have both stated that this should be achievable, but both are dodgy on the details. A specific test would be useful, and I'll ask Ken about this.

> >-- Is the LO/IF design and VLBI compatible?

To be addressed in LO/IF PDR.
>
> >-- The pointing specification should include limits such as night time, wind conditions.
>
This has been added. It says that the precisions pointing specs are for calm (wind less than 10 m/sec), nighttime, thermal equilibrium conditions.

> >-- Where should absolute time be made available? What is the plan for providing absolute time?
>
As noted above, it's going to be 10 nsec. And the experts tell me that GPS can (or will shortly be able to) do this.

> >-- Is the 50 dB dynamic range specification sufficient to prevent ringing in the presence of strong sources and to accomodate RFI in L-band? Should the range specification be a function of frequency range? Brent Carlson recommends that the range specification should
be a function of the ratio of RFI power to noise power in the band that the initial quantizer sees.
>
Brent has done *extensive* simulations to show that his correlator will provide > 55 dB linearity in the presence of strong RFI.
These refer to the frequency response.
The spatial response to a strong object is a subject which lies in the e2e area -- or more specifically, it's an imaging problem. The specs here are actually higher -- 70 dB. Some day we'll figure out how to do this.

> >-- What is the plan for post RFI excision? Is there an overall coherent plan for dealing with RFI?

A draft plan has been prepared. It calls for suppression of self-generated RFI and providing a linear design as essential priorities. More to the point, a number of groups worldwide are working on this problem. I think there is general optimism that such post-correlation removal can be done, but the accuracy of the removal is yet to be determined. We will remain watchful, and will incorporate such methods as are developed. At this time, what is critical is to ensure our 'pre-correlation' system response is as linear, and as flexible, as we can achieve, without sacrificing our sensitivity goals, or our budget.

> >-- Time limits need to be specified for stabilities.

This has been done.

> >-- How well do the signal gains match from antenna to antenna and band-to-band? What are the relative priorities of stability, gain, flatness?

Done.

> >-- What is specification for clipping noise?
>

I'm not sure what is meant here. Our proposed 3 bit and 8 bit (8 level and 256 level) sampler schemes add negligible noise. Errors in the 'clipping noise' (meaning samplers, I think) must add something, but I don't know what. Whatever it does, it can't add more than a small amount of overall noise -- the overall system specs (which are clearly specified) will set the limits on this term. Somebody else will have to determine what these are.

> >-- More work needs to be done to formalize the scientific and performance requirements. There is no single source listing requirements. Desired, nominal, and worst acceptable parameters for each requirement would guide designers as well.

As commented earlier, working on it.

>B. Systems
>
>-- What is absolute time distribution scheme? Are you taking advantage of COTS, ethernet time distribution schemes/equipment to save money?

We will probably use a combination of COTS ethernet time distribution along with a higher precision timing signal for precise synchronization. The precision timing will be generated using low cost programmable logic technology.

>>-- Digitizer specification should be realistic; if you can't get 8-bit, then don't specify it.

I'm not sure exactly what is meant by this. If it menas can we get 8-bit devices, yes we can get 8-bits now. As shown in the PDR presentation, suitable devices exist now (Maxim MAX104/108) that would allow us to build a 2:1 inteleaved 2gsps 8 bit device. Simulations have shown that the 2:1 interleaving is not a problem. If it means do the 8-bit devices produce an effective 8-bit resolution, probably not - the actual number is estimated to be around 7 bits. That doesn't change the fact that there are 8-bits coming from the device and that we need to transmit and process them all of the way through the correlator.

>-- Pay attention to board design techniques and new technologies such as distributed capacitance on Vcc plane.

Agreed. Several engineers have been to a high speed digital pcb design course and I hope to send more this year. The first step in combating RFI is at the circuit/PCB design level. Sinusoidal phase rotation is explained in detail in EVLA memo 11 and the actual digital phase rotation function is further analysed in EVLA memo 12. This type of phase rotation is used in VLBI correlators, and in particular the lag-based phase rotation is used in the Canadian S2 space VLBI correlator in Penticton (described in Carlson et. al., PASP, 1999, 111, 1025-1047).

>C. LO/IF

>-- What is the source of specification for 0.7 ps and 0.07 ps stabilities?

To be addressed in LO/IF PDR.

>-- What will be impact on phase of flexing LO fiber cable? Is there room for a cable wrap-up if necessary? How would this be done? Are there provisions for keeping fiber out of the sun and otherwise keeping the LO cable temperature stable? Should you be using tube housing for fiber on telescope?

To be addressed in LO/IF and Fiber PDRs.

>-- Are isolators necessary?

To be addressed in LO/IF PDR.

>-- The use of sinusoidal phase rotation instead of phase switching should be explained. Brent Carlson explains that sinusoidal phase rotation is a requirement of the correlator and, as Barry Clark says, "is the moral equivalent of phase switching". Phase rotation may also be more flexible and have better artifact decorrelation properties than 180 deg phase switching.

Sinusoidal phase rotation is explained in detail in EVLA memo 11 and the actual digital phase rotation function is further analysed in EVLA memo 12. This type of phase rotation is used in VLBI correlators, and in particular the lag-based phase rotation is used in the Canadian S2 space VLBI correlator in Penticton (described in Carlson et. al., PASP, 1999, 111, 1025-1047).

A comparison between sinusoidal phase rotation and 180 deg phase switching has not been carried out in detail, but it is known that some artifacts of phase switching have no effect or are easily handled by phase rotation. Spectral artifacts introduced from using interleaved samplers (due to unmatched gain) wash out with phase rotation, but do not with phase switching. I suspect (but have not yet proven) that RFI introduced into the signal after phase switching (but before "de-switching") is not removed like it is for phase rotation.

For phase switching to be effective requires that integration times are integral numbers of entire Walsh switching cycles. This restricts flexibility particularly when wanting to dump rapidly and differently on different baselines and when synchronizing integration times to pulsar timing. Phase rotation has no such restriction. The only requirement is that within an incoherent integration time, enough phase cycles be present so that digital mixer edge effects are not apparent. Exactly how many cycles are required is not precisely known, but generally a minimum of 10 cycles (providing about 20 dB of aliasing rejection and "after the LO offset" RFI rejection) is sufficient. I have recently asked Gareth Hunt to ask Fred Schwab to look into this in more detail, and establish a true minimum requirement. Nevertheless, the number of cycles in an integration is a free parameter depending on the LO offset and required integration time.

Finally, at one time I suggested that phase switching be included in the sampler board design (EVLA memo 23) "just in case" there is some lurking artifact that can only be handled by phase switching and can't be handled by phase rotation. Barry Clark has since indicated that he finds this unnecessary.

>-- Is the cleanup loop necessary or could this function be performed on fiber?

To be addressed in LO/IF PDR.

>D. Receivers, feeds
>
>-- Will the L-band OMT require a new dewar?

The L-band dewar will be re-used by replacing the extension which holds the OMT with a longer piece.

>-- The OMT size estimated for L, S, C, and X bands seems small.

It probably is. The actual length is not yet determined.

>-- Is the scaling method to project Trx correct or are the results on the chart optimistic?

The scaling method was used as a way to extrapolate from existing receiver performance. It was not as accurate as it could have been, but the numbers are probably not far off.

>-- Should the quad hybrid be located ahead of the LNA for phase stability?

Yes, that is the plan.

>-- Gain blocks will balance the amplifiers and provide 3 dB additional head room.

We hope to use the L-Band balanced amplifier gain blocks which Rich Bradley is developing for the GBT. Above L-band would require new designs. They also require twice as many stages for the same gain.
>-- Will the Iridium filters be used on EVLA? Why?

Currently, no. Iridium is obnoxious not so much for its strength as that it is pulsed at a rate that beats with the 9.6 Hz VLA cal cycle. EVLA can use another cycle rate and excise the actual carriers downstream. Should it prove necessary, the filters can be re-installed.

>-- Have you planned adequately for production issues such as LNAs from CDL, high frequency component electroforming, acid etching?

Noted. This is an area that will require careful planning, attention, and cooperation. A request has been forwarded to CDL for 488 amplifiers.

>-- What are plans for pulse cal and why?

No plans for pulse cal at this time.

>-- Should you be using a split, single source Tcal?

This seems to be a suggestion for a single broadband noise diode which could be switched between several receivers. This might work in theory for the the lower frequency receivers. I'd be really worried about achieving a 0.01 dB Tcal stability over 10's of minutes. Plus calibrating the Tcal signal would be difficult if the noise diode wasn't built into the receiver (it's currently done while we do hot/cold load tests in the lab). But it bears some study.
> E. Correlator
>
> -- Does chassis construction provide adequately for insertion force required for high pin density connectors?

The insertion force has been calculated and is estimated to be about 100 lbs (EVLA memo 31). It is planned to provide adequate board stiffening and chassis design to support this force. As a comparison, a 12U VME board utilizing 4, standard 96-pin DIN connectors has a calculated insertion force of about 85 lbs.

>
> -- Has adequate cooling been provided for 128 W circuit boards?

We believe so. Each correlator board will have the equivalent of 3 complete VME slots to itself. Thus, there will be 2 empty slots (1.6") between each board. This should be adequate for any additional chip heat sinks (and mezzanine cards) that might be necessary.

> -- What is failure rate for 11,000 chips in correlator and 5,000 chips in FIR? Frequency of failure will impact software. What is the impact of temperature on failure rate? Are diagnostic

and maintenance procedures adequately provided for?

Failure rate of complex, highly density devices is difficult to analyze. Some indication of this difficulty is in the article by Parry, Rantala, and Lasance "temperature and reliability in electronics systems-the missing link", Electronics Cooling magazine, volume 7, no. 4, Nov. 2001. It is generally not possible to apply reliability prediction techniques of the past such as those of MIL-HDBK-217 to modern VLSI devices. Nevertheless, the techniques that are being used to ensure high reliability are as follows:

a) Use BGA packages. BGA packages have a higher production yield than QFP packages (i.e. fewer solder faults). This should ensure that the number of questionable solder connections that can lead to future failures should be low. Also, BGA packages have better thermal performance than QFP packages in that they are able to conduct more heat to the PCB. This results in lower temperature and increase reliability.

b) We plan to incorporate a rigorous program of accelerated life testing on production boards. This should find most failures such as chip die problems, marginal timing, and marginal solder faults before the boards are installed in the final system.
c) Provision for very good cooling...lots of space between boards and keeping the total power within a rack to within a reasonable limit. This should eliminate hi-temperature related failures.
d) All boards are hot-swappable. If desired (but at extra cost), extra Baseline and Phasing Boards can be added for N+1 redundancy.
There is some, but not complete, redundancy of Station Boards. Full redundancy would require a large and very costly cross-bar switch in front of the Station Boards.
e) On-line synchronization codes ensure data transport integrity. Nevertheless, there are some failures (such as a FIR or correlator chip gate failure) that cannot be detected on-line. We plan to include test features that can be run when the correlator is not in use (e.g. antennas slewing between sources), that will easily be able to find these sorts of failures.

Finally, past experience (of this engineer anyway) has demonstrated that modern digital chips and boards do not fail on a regular basis (after accelerated life testing) provided they are from reputable manufacturers, and that device temperatures are kept within reasonable bounds. There is no expectation that correlator boards will fail on a consistent and regular basis, given our current plans and estimates. But, all of the above mitigation plans are the backup plan in case boards do fail on a regular basis. Also, we plan to do more research into the area of reliability and reliability prediction.

> -- Do we want to specify the correlator for 80,000 km baseline required for space VLBI?

The current plan is to provide a 0.25 sec delay buffer that is capable of supporting about 25,000 km baselines. Going to longer baselines is a matter of cost. It is expected that once we commit to production, that larger memory chips within the projected cost envelope will be available and will enable longer baselines. However, there is no guarantee. Changing the spec. to 80,000 km baselines at this point may have a significant cost impact.

> -- Correlator software support division between Penticton and NRAO was reported to be NRAO: 2 software engineers, Penticton: 1 software engineer.

Correct.

>F. e2e

>-- Will data from auxiliary functions such as API (Atmospheric Phase Interferometer), weather station, solar cal be provided for in real time equipment? Weather data will be part of the ancillary data products. The requirements for the speed of measurement of the solar calibration values are still under development.

>-- Are you taking adequate advantage of ALMA development?

Yes.

> -- Are you going to run into developmental delays and "design- related" problems like AIPS++?

We have changed our development model explicitly to avoid some of the problems encountered in AIPS++. Also e2e is internal to NRAO and will thus avoid some of the problems associated with multiple organization projects like AIPS++.

> -- How are you providing for a meaningful first examination of data quality? With the data rates planned, won't the monitoring have to be performed by machine rather than operator?

Yes. This is a key part of the design of the pipeline.

> -- Very smart algorithms are going to be needed to bring the data examination/editing state to automation.

We disagree. For the most part, simple editing algorithms can be used once the system is operating well. Experience with the VLA and VLBA supports this view.

> -- Is 1 man-year for proposal handling too low? Are all manpower estimates too low?

The questioner is incorrect: the estimate for proposal handling is 10 FTE-years.

> -- What are software acceptance procedures? Are they acceptable? Isn't the 6 months acceptance period that you have scheduled too long?

For the M&C, these are presently undefined. They will have to be developed once all of the requirements have been agreed upon.

> -- What software tool kits are provided?

To be addressed in future PDRs.

> -- Is Glish the language to use? Is there a documentation problem with Glish? How will you maintain costs?

Glish is well-suited to the needs of e2e. There is no documentation problem (see the AIPS++ web site). I'm not sure what the last question means.

> -- Is it a problem that the AIPS++ Group is the sole maintainer of Glish? Is having only one Glish expert in the group a single point of failure?

No, it is no more of a problem that that AIPS++ is the sole maintainer of the C++ libraries.

> -- Do you know what are the current VLA obsering modes and that you will be able to support all, even during transition?

Yes, and we intend too; although, it is probable that some modes will not be fully supported during transition. A plan for the concurrent operation with existing and enhanced antennas is under development. It will be reviewed at the M&C PDR.

> -- Should the e2e design target a wider audience of scientists, even those not trained in radio astronomy?

It does.

>-- Tim Cornwell commented that attaining the target of 60 dB dynamic range could require correction of the time variable primary beam caused by the alt-az mounts of the VLA. This may be a very compute-intensive problem, depending on how often one updates the correction. A large computer system may be required to compute such corrections in a reasonable time.

See comments about Beowulf clusters in Part I.

CJ> G. M&C Software

CJ> -- Are you providing an adequate firewall to block hackers?

Legitimate access by remote users is a requirement. The GBT is presently investigating this while providing a block to unwanted intrusion. At present, it does not seem as if firewalls will be needed. However, if they are, then they will be specified and operated by the NRAO security team and procured from regular observatory funds.

CJ> -- Explain clearly where the correlator ends and the correlator back-end begins.

The correlator ends at the baseline boards. The baseline boards perform the actual correlations and output the lags. The correlator backend begins at the output of the baseline boards.

It will be the job of the correlator backend to assemble the lags and perform the FFTs. Some RFI excision may also be done in the correlator backend. The output of the correlator backend will be FFTed subbands.

CJ> -- The transition plan of using SLC for control of new and old M&C designs during the transition phase is not clear. Will the Modcomps ever be providing control via the new M&C?

This will depend upon the outcome of the concurrent operation study.

CJ> -- Will phasing the array during the transition be a problem?

Support for phased array during transition will have to be studied in more detail. The main problem is the feedback of individual antenna/IF phase changes. How best to do this will not be clear until a transition plan has been adopted. Successful phasing of the array may require information exchange between the nascent EVLA M&C system and the VLA M&C system.

CJ> -- Will all existing observing modes be supported during transition or will you have to give up some like sky surveys, mosaicing modes?

It is probable that a few modes will not be available during transition. However, the surveying modes will probably be fully available.

>H. M&C Hardware

>-- Have you explored adequately COTS substitutes for MIB design? There is a "de-facto" industry standard format called PC/104 or PC/104+ that is very inexpensive and seems well suited to the application.

Use of a commercially-generated chip and in-house board design seems best-suited for the antenna MIB.
>-- The plan for module and location ID needs firming up.

Agreed.

>-- Have you provided for packet collisions in M&C design?

To be part of M&C hardware PDR.

>-- Shouldn't "Backup M&C" be absorbed seamlessly into M&C operation?
>Get rid of extra resets such as current need to go to antenna to reset power after a lightning strike.

To be part of M&C hardware PDR.

>-- Antenna should stow if M&C communication to antenna fails.

Agreed.

>-- Who will maintain phone equipment?

Telephone equipment will be given to NM Tech ISD for repair. Transmission line problems will be addressed by Electronics Division.

>-- XDR is a pain. (XDR or eXternal Data Representation, the standared for a machine-independent data representation.)
> One reviewer counter argues that XDR is easy to use as long as one properly uses the "rpcgen" compiler, .x files, etc., although XDR should not be used where high-performance data transfers are required. XDR can be used for data transport across the network (e.g. if using RPCs), or it can be used to easily encode data for saving into a file or decode XDR-ecoded data from a file. Without XDR, saving data into a file and retrieving it is time-consuming and code-intensive. With XDR, the save and retrieve operations are, quite literally, "one liners".

To be addressed in M&C PDRs.

CJ> I. Operations software.

CJ> -- Has target of opportunity observing and dynamic scheduling been adequately provided for? Short notice changes should be supported on the order of seconds.

This is a requirement. It will be part of the M&C design. Dynamic scheduling has always been viewed as an integral component of the EVLA. I believe it has been and will continue to be adequately supported in our considerations. It is, most properly, a subject that falls in the area of the interface between the e2e and M&C areas.

Target of opportunity has been considered in the context of minutes, but not really on a timescale of seconds. A goal of short notice changes on the order of seconds must be investigated w.r.t feasibility. Again, it is a topic which will require cooperation between the e2e and M&C efforts.

CJ> -- During the design phase, what procedures will be used to identify alarms, priority of response, responses and provide that information to operations programmer?

Alarms and the related issues mentioned above have been considered in the first drafts of both the Operations and Antenna Monitor and Control Requirements documents. However, the considerations are still rather vague and insufficiently developed. The procedures so far used to identify these items include observation of the VLA operators by the author of the Operations Requirements document, review of this document by the head of Operations, and discussions with the hardware designers. Over time, we plan to draw individuals involved in Operations more and more deeply into consideration of this aspect of the system. Other issues, such as progress on the choice of MIB processor and RTOS are more pressing at this time and are receiving the bulk of the time available from our currently limited manpower resources.

In short, the implementation of the alarm system will not be started until after all requirements are gathered and the design of the M&C is complete.

CJ> -- Will failures be integrated with the maintenance management system?

Yes. The ability to query the maintenance database has already been mentioned in the 1st draft of the Operations requirement document. The ability to submit problems to the database or otherwise connect failures to the maintenance database has been discussed, but I do not know if some version of these discussions has appeared in written form.

CJ> -- Ergonomics of operator messaging needs careful study.

Agreed. We are fortunate that the individual who will be a primary developer for the operator GUIs already has a long history of development work with the VLBA.

CJ> J. Engineering requirements for software

CJ> -- MIB specification seems ambiguous. Keep the MIB simple but agree on requirements.

The MIB specifications are somewhat less ambiguous now (1/30/2002) than they were at the time of the EVLA system PDR, but still require additional work. We are continuing to develop and refine this specification. The groups developing the various EVLA hardware systems have been queried (via a written document) re their needs, characteristics considered essential to minimizing MIB-generated RFI are being discussed, and a list of candidate chips and RTOSes is being developed.

CJ> -- Settling times for antenna, synthesizer need specification.

Noted. Agreed.

CJ> -- Use of tech laptop may need clarification.

It does need clarification, and we are now in the process of doing so. One point that has become clear is that the tech laptop will have the ability to communicate directly with MIBs in addition to whatever software entity (or entities) are used to encapsulate the antenna as a whole. Screens for communication with the MIBs will be among the first software items developed. The ability to communicate directly with each MIB is essential in many respects, both to the design and development effort, and to maintenance issues.

CJ> -- Subarcsecond pointing precision is not necessary, though it may be necessary to read an additional bit or two of resolution from the absolute position encoders to achieve the 2 arc second reference pointing goal. Use of the additional bits must be coordinated with the group performing encoder electronics modifications, and tests should support what additional bits, if any, are useful.

New VLA electronics being installed provide additional bits should they prove necessary or desireable.

> K. Backend correlator computer
>
> -- Is the plan to use switching of packetized, serial data streams to combine correlator data frames affordable?

Yes, it looks like it is available, based on quotes of Gbit switches, and "press-release quotes" of the Cypress chip on the Baseline Board that enables it. We feel that the correlator backend using either FPDP parallel data streams or packetized serial data streams is not adequately supported by the WIDAR budget. I have spoken with the EVLA project manager concerning this issue and he is willing to make contingency funds available to supplement the WIDAR budget in this area.

> -- Can you hot-swap Ethernet? Brent Carlson responds that COTS Ethernet switches quickly learn what their I/O IP addresses are and so hot swapping back-end computers or Baseline Boards should not be a problem.

I am not aware of hot-swappable ethernet products. I do not know how ethernet NICs would behave in a system designed for hot swap capability such as CPCI. NICs however are only one component. As to hot swapping switches, Brent Carlson's statement is correct. I do not feel that I have enough information, at this time, concerning the baseline boards to confirm or contradict Brent's statement on that point.

> -- Is packet protocol the way to go?

Yes, yes, and yes. The original design for the correlator backend used FPDP (front panel data port). This approach suffered from at least three serious deficiencies. First, FPDP is length-limited. It was unclear that it would be possible to position the input end of the ports close enough to the output ends to satisfy the FPDP spec. Second, the design called for input of the data directly into computers, with 4 ports per system X 64 systems to receive the data. With 4 ports per system, the backplane speeds of each system receiving the data was a bottleneck. We also felt that 64 systems was an unduly high number which might threaten reliability. Third, the FPDP design necessitated fixed ports at both the output and input ends of each interface, i.e. the output of a given baseline board would always be received by the same port on the same system. This arrangement would have necessitated data exchange among the nodes in the Beowulf cluster, another potentially severe bottleneck. The switchable serial data stream design we are now using addresses all of these shortcomings, and does so in a reasonably cost effective manner. Moving the output of the baseline boards from a length-limited non-switchable parallel protocol into the arena of switchable serial datastreams was a major design win.

> L. General

>-- Add a paragraph detailing justification for 8 GHz IF bandwidth over 4 GHz. Doubling the bandwidth at Ku thru Q-Band will not come cheap.

Agreed.

> -- Specify documentation required for design. Documentation should included interface requirements, theory of operation, schematic, maintenance instructions, BOM. Documentation should be uniform between task groups. Technical Reports for VLA can serve as guide.

Agreed.

> -- Add a Table of Contents to all Project Book chapters.

Noted.

> III. Conclusions:
>
> The top level performance requirements for EVLA system design are complete and adequate with the following important exceptions:
> Timing Accuracy
> Rapid response target of opportunity observing
> RFI mitigation, what, how, requirements?
> E array in so far as it effects Phase I
> Need a requirements document that lists priorities.
> It is important that the various requirements documents be integrated into a single chapter in the Project Book, and at the appropriate time put under change control.

The document "Requirements for Interface Specifications (or Interface Control Documents)" is currently in the EVLA schedule. Draft ICD's are required by 1 July and Final ICD's by 1 September -
four months prior to the bench integration.
RFI and E array issues are being addressed.
Timing is being addressed as an LO/IF issue.
Rapid response to targets of opportunity is being addressed as an e2e issue.

> The EVLA M&C system is lagging behind in development compared to the rest of the project. This is absolutely no criticism of the M&C staff; it reflects the fact that the EVLA M&C group is only just now being staffed, while other areas of the project have apparently had significant effort for some time. It will be extremely difficult for the M&C group to catch up; EVLA project management should take all possible steps to bolster the resources available to them. If any of the existing EVLA electronics staff have software engineering expertise, one approach might be to divert these to study some aspect of the M&C system, even if this causes delays elsewhere.

This is currently the area of biggest concern. We have scheduled an M&C hardware PDR for early March. The choice of MIB architecture(s) and processor(s) should be made by that point. We do not really want a system that is already 15-20 years old (VLBA/GBT) nor do the EVLA engineers feel comfortable with the choices made for ALMA. The technique we have chosen is based on Ethernet and is much more "COTS" than either the VLBA/GBT or ALMA systems.

We are working to define
the MIB architecture. While it is clear that a COTS solution is probably best for the "Correlator MIB", it is not clear that a COTS solution is best for the "Antenna MIB" - mainly due to RFI concerns.
We may have to temporarily control prototype hardware with software like LabView. We do not plan to delay hardware designs to wait for the software.

> It's very desirable that end-users (telescope operators, schedulers, astronomers) be continually involved with the specification, design and testing of the e2e system as it is developed. To ensure this involvment, NRAO should consider assigning people from these groups to be e2e project members, and make it part of their job requirements (job description, performance evaluation ?) that they work with the e2e system developers at least part-time on an on-going basis.

I think this is an excellent suggestion.

> The procurement plan for the EVLA project involves many distinct groups (EVLA M&C group, electronics group, e2e project, Canadian correlator project group, etc.) The interfaces between the various sub-systems have not been defined in detail. Without these, there is a strong possibility that the system as a whole will not come together as hoped (or that integration will take far longer than planned). The appropriate interface specifications therefore need to be identified, written and formally reviewed. Until this has happened, we cannot state that an adaquate system design is in place for the EVLA as a whole.

A recently released memo (NRC-EVLA Memo# 015) on correlator M&C and backend software requirements and interfaces is a very good start at addressing this issue (at least the interface to the Canadian group, anyway).

> Correct design solutions have been selected except for the following:
> Consideration of analog data transmission.

Addressed in Part I.

> Improved designs for L/S/C band front ends to lower Tsys (if > possible),

Addressed in Part II. D.

> Emergency communication/M&C backup subsystem is not well defined,

The Emergency communication/M&C backup subsystem and the new phone system will be incorporated onto the 1 Gbs Ethernet MCB. The bandwidth requirements for the existing system have been estimated and it has been determined that 1 Gbs Ethernet can easily handle all three systems. If additional bandwidth is required in the future, then additional Ethernet systems will be added. Will be addressed further in the M&C hardware PDR.

> Response times required for M&C RTOS (Real Time Operating System) need to be analyzed -- if sufficiently stringent, they will tend to drive the design,

Will be addressed in M&C hardware PDR.

> Use of 1310 nm F/O links for distributing frequency reference to reduce dispersion in fiber (assuming SMF28 is selected as
fiber).

It is recognized that the LO round trip measurement system will require a fiber transmission system with minimum dispersion. Although SMF 28 fiber has zero dispersion at 1310 nm, the insertion lose at 1310 nm also needs to be considered.

> Provide for adequate RFI suppression and provide the LO offset early in the design, if the digital transmission option is chosen.

Addressed in LO/IF PDR and RFI plan.

> In the case of the M&C system, only the most high-level general decisions ("it will be a distributed architecture using Ethernet") have so far been made. This high-level approach is appropriate. However, cannot assess in detail whether the correct solutions have been selected for the M&C system until the design is further advanced. The EVLA project should give serious consideration to adopting the M&C infrastructure from an existing project, or collaborating with, for example, ALMA; even if this means compromising in some areas on the system requirements. Use COTS for MIB.

Will be addressed in M&C PDRs.

> The procurement plan for the computers was the only one mentioned in the Systems PDR. The computer procurement plan calls for "just in time delivery" to reduce obsolescence issues, a good idea. Other plans for procurement should be elaborated during the > subsystems PDRs.
Software is the area where extra vigilence is required to ensure that it is available when hardware is available so that adequate time is provided for system integration and testing. This may require carefully defining (or prioritizing) what the core software requirements are and ensuring that development is focussed on meeting those needs first to keep the project on schedule. Subsequent software releases can add additional functionality or "bells and whistles", as desired.

Noted. Some comment on this issue in Part I.

The following addendum is accepted in its entirety.

>Addendum from Dick Thompson:

>It is noted in the report that there are no quantitative details to support the decision to use digital rather than analog transmission of the IF signals. However, while a quantitative comparison of the costs of the two options is can be made, it seems to me to be very difficult to compare the relative benefits accurately until it is possible to make quantitative tests of the performance. The arguments in favor of the digital choice are:

>(1) It allows full use to be made of present and ongoing developments in the relevant areas of communications technology.
>(2) The frequency response of the transmission system is eliminated. Of these two, I am inclined to think the first is the more important, since the analog response of the fiber should be good enough. Digital transmission will surely be used for the long links to more distant antennas (Pie Town and beyond) in phase 2, and it is also the choice of ALMA, so it makes sense to concentrate the overall engineering effort on the one technique.

>The arguments in favor of analog transmission are:

>(1) It reduces the amount of hardware at the antennas and thereby reduces the potential for self-generated interference.
>(2) Having more of the electronics at the central location is more convenient for repair and maintenance.
>We know from the existing VLA system that the self-generated
>interference is a serious problem, but it it is hardly possible to extrapolate from the existing performance to what we would expect to get on an entirely new system. So any quantitative estimate of the self-interference problem will have to wait until there is new hardware to measure. I think that this is the only serious point in favor of the analog transmission, other than cost. With regard to convenience in maintenance, the digitization hardware at the antennas has to be developed to a level at which reliability is not a severe problem, otherwise there will be problems for the distant antennas in phase 2, and for ALMA. Thus I think that digital transmission is a good choice, so long as the additional self-interference that may result from it does not turn out to be a very severe problem. It is difficult to say much more to justify the digital choice until tests are made on the hardware.
>
>A point that I find somewhat worrying is that we have to wait until the first part of the Widar correlator is available before being able to test a pair of antennas with the full 2 GHz IF bandwidth. In order to forestall any upleasant surprises at that point, it would be good to make tests using the existing correlator to investigate the full band response as soon as a pair of antennas is outfitted with the new system. The existing 50 MHz bandwidths could be stepped across the full 2 GHz band, with several hours of integration for each step, to get something approaching the sensitivity of the final system. Test observations will also be required with the receiving band set to include each potential self-interference frequency, i.e. the fundamental and harmonics of each of the the LO reference frequencies at the antennas as well as the sampler-frequency harmonics, etc. Such tests could take a lot of time, so at least two of the outfitted antennas should be available on a priority basis for tests rather than astronomy. I am glad to see the the "five-antenna rule" in section 11.2.2.1 of the Project Book, which would allow three for testing even if two more were out of action.

>Another point in the report concerns justification of the 8 Ghz bandwidth per polarization, rather than 4 GHz. This is really a matter of the scientific specification rather than the system design. It may well be difficult to find 8 GHz of bandwidth free from interference in the frequency range of the EVLA. However, as I understand it, the main strategy for interference excision is to make use of the spectral-line capability of the correlator to identify channels with interference and delete the corresponding data. This should work well for narrow-band interference, for which the interference-to-noise ratio gets better as the bandwidth is reduced (down as far as the bandwidth of the interfering signal). Broadband interference is likely to be more difficult to handle. However, experience with EVLA observing should gradually build up information on which parts of the various bands are free from interference. So long as, say, half the spectrum is usable for astronomy, the bandwidth provided by the 8 GHz-wide response will be valuable. As with the original VLA system, it may be advantageous to bring the system into operation with half the final bandwidth, and add the extra modules as funds and manpower permit.