EVLA Advisory Committee Report

Meeting: 14/15 Dec. 2004; Report: 04 Feb. 2005

Sean Dougherty, NRC Herzberg Institute of Astrophysics

John Dreher, SETI Institute

Gianni Raffi, ESO

Mark Reid (Chair), Harvard-Smithsonian CfA

Luis Rodríguez, CRyA UNAM

Alan Rogers, Haystack Observatory

Stephen Scott, Caltech, OVRO

Marco de Vos, NFRA

Sander Weinreb, JPL

Contents

1 Executive Summary

Most aspects of the EVLA project are proceeding very well. The antenna re-construction is proceeding well and no critical problems have been encountered. The very large L-band feed, clearly the most difficult of the new feeds, is performing well with excellent spillover reduction compared to the old feed. The monitor interface board, a critical component in many systems, appears to be working well. Progress in the electronics lab is accelerating, after a worrisome, slow start. Software efforts are progressing on several systems. Requirements documents are mostly in place, the correlator software is improving, and recent success in testing task in AIPS++ is encouraging. The phase-II proposal has been submitted, and (after a bewildering delay) the NSF has finally sent it out for refereeing.

Some hardware tasks are behind schedule, a problem that is clearly recognized by both management and staff. A plan to recover the schedule slippage is in place and, in large part because of extra hours worked by many staff, it appears that the original schedule can be recovered within a year.

While the feed tests have all been positive, there is some worry about purchasing all feeds without complete system testing. The committee recognizes that interim observations, which mix old and new feed designs, may have large polarization differences; these problems may be calibrated-out, but in any event this will be a temporary problem and should not be cause to divert effort to address it.

The electronics laboratory needs to have two test systems operating simultaneously in order to simulate interferometric observations and test various component subsystems.

The system integration and general debugging activities are currently in the hands of a few senior staff. The committee feels that it is very important to bring some younger NRAO staff into the debugging/commissioning activity. This has two main benefits: 1) it increases man-power on this important activity, and 2) it will give new staff a broader understanding of the entire system, which will be important for future operations, maintenance, and enhancements.

The software efforts, including monitor and control, e2e, and post-processing, at some levels lack full system designs. This may result in extra work, and possible delays, down the line during the commissioning phase.

The monitor and control is a complicated and challenging system which needs attention. Neither the scope of the system, the scientific requirements, nor the interfaces appear to be clearly defined. This may result in significant problems during the commissioning phase. The committee recommends that the specification of the maximum number of antennas that can be missing during interim operations be relaxed to minimize inefficiencies which may slow the entire project.

The e2e effort is important to the EVLA project, yet there is a large mismatch between the requirements of the system and the available resources. A realistic assessment of the required resources is only possible when an adequate system design is in place, complete with a chosen framework, and definitions of subsystems and their interfaces. The lack of a such a design impacts the ability to take advantage efficiently of software being developed for other projects, particularly ALMA. Given the lack of resources, a minimum subset of priority 1 requirements needs to be defined.

The AIPS++ software has made good progress recently, largely through the need to meet the ALMA software schedule, and may emerge as a usable data-processing system. As with e2e, the available resources is a concern. With much of the development driven by ALMA, the EVLA-specific requirements need to be identified and prioritized, so that the required resources can be identified. We note, however, there is still a large risk as the user interface is untested, there is no complete science requirements document, and no clear and specific design plan for the entire system.

The current projection is that contingency may be inadequate to complete the entire project, given post-proposal growth in the desired software. We recommend a careful scientific trade-off study that considers where best to descope. This would include postponing desirable software development and not completing one (or more) frequency bands.

2 Committee Charge and Meeting Information

The charges to the 2004 EVLA Advisory Committee are to evaluate and advise the NRAO Director regarding: (a) the EVLA Phase I project's technical progress and issues; (b) the project's software development plan and resource requirements; (c) the EVLA Phase I management plan, schedule, and cost, including strategies for the effective use of project contingency; (d) the scope and maturity of the project's operations plan; and (e) the relative priority and scientific impact of options for change in project scope.

The committee met in Socorro on 14 & 15 December 2004. Presentations on all aspects of the project outlined below were made by NRAO staff and Peter Dewdney (HIA, Canada). The committee toured the laboratory facilities in the AOC and the telescope construction, fiber-optic layout, and new correlator facility at the VLA site.

3 Project Management

The Committee was pleased to see that a recovery plan to get the project back on schedule has been put in place. The milestones completed are currently apace with the recovery plan, and all involved in the project are commended for this effort.

The attainment of first EVLA-EVLA fringes is exciting news and a clear indicator that the project is progressing well. The submission of the Phase II proposal was another positive issue achieved during 2004.

The cost of the effort contributed to the EVLA project from VLA Operations is of concern. Currently, this contributed effort is 10 FTE/yr higher than the original baseline plan and cannot be sustained by Operations at this level, especially with the foreseen tight operations budget. Returning the cost of this effort over the next few years to the construction budget, and then recovering the cost from Operations in the closing years of the project (2010-2012) might mitigate this problem. However, there is the risk that Operations may not have the necessary funds at that time, and work that can be delayed until late in the project needs to be identified. The committee suggests this strategy should be investigated further, including exploring potential methods of mitigating the risk of Operations having difficulty contributing the effort in 2010-2012.

Another issue that needs to be given serious consideration is software resources. The staff available for e2e development are far short of those needed to meet the requirements of the current e2e design, which is of concern given the ambitious goals for e2e and the maturity of the EVLA project. The cost of software development, both e2e and AIPS++, also needs consideration. We note that $500K was contributed to AIPS++ development in 2004/05, yet there was no mention whether or not contributions of this size will be required in future years, and whether or not this has been budgeted in forthcoming years.

The goal of keeping as few antennas out of the VLA for retrofitting with EVLA systems is commended. However, this may be placing undue stress on the project staff, who are trying to get functional antennas back into the VLA as soon as possible. The committee will be interested in reviewing this issue in the future.

Finally, the possibility of having to descope some aspects of Phase I was of concern to the committee, since the two options mentioned, dropping some bands or reducing the bandwidth, are serious and affect in important ways several scientific possibilities of the project. We urge NRAO to search with NSF for a solution to this possible problem.

4 Antennas & Feeds

4.1 Antennas

VLA antenna 14 has now been outfitted with several EVLA feeds and amplifiers so that testing on an EVLA-EVLA baseline between the previously-outfitted antenna 13 has started.

4.2 Feeds

The EVLA L-band sensitivity from 1.2 to 2 GHz looks excellent. The spillover of the L-band feed on the EVLA antenna is much less than the VLA, so that the Gain/Tsys performance is a considerable improvement over the VLA. The C-band performance of the EVLA feed is also an improvement over the VLA. There is some concern over the polarization performance of the L-, S- and C-band feeds, which are similar in design. These octave bandwidth feeds use an OMT followed by a quadrature hybrid to form dual circular polarization. The return loss of the OMT is being improved, which in turn is expected to improve the efficiency, remove the trapped mode resonances, and improve the polarization orthogonality. However there may be a significant lack of orthogonality between the EVLA and VLA feeds, which will make the "D" correction terms on the EVLA-VLA baselines uncomfortably large during the changeover period. Polarization observations will be more complex during the interim period before all VLA antennas have been upgraded with EVLA feeds. It is not clear that anything can be done about this, other than performing the necessary calibration.

There was some discussion about the possibility of going to linear polarization by removing the quadrature hybrids in the EVLA feeds and the quarter wave dielectric wedge polarizers in the VLA feeds, but there was not much support for this option and so this option is not suggested by the committee. The progress on the antennas and feeds is very good, but there is still a need for complete testing with the EVLA amplifiers before the feeds are procured in quantity. However it should be safe to go ahead with procuring the large sections of the large L-band feeds, as their design should not be affected by the refinements in the OMT design. The committee recognizes the urgent need to start the procurement of L-band feed.

The work on the refinement of the OMT design should be given the highest priority. A high priority should also be given to building and measuring the system noise and polarization of a prototype wideband front-end over the full octave frequency band. For this purpose C-band would probably be the best choice.

5 Electronics

5.1 Schedule

While there has been excellent progress on the development of new modules to accommodate the very wide bandwidth and frequency ranges of the EVLA, the development has been slower than expected due to challenging requirements, the need for design robustness before making large quantities, and the complexity of testing. The design engineers appear to understand the remaining problems and are making appropriate corrections. We recognize that some problems must be fixed before proceeding with manufacture, while others can be accepted as causing limited effects upon a small range of observations. We believe the EVLA management has the knowledge and experience to make these decisions.

5.2 Staffing

In general, we note the increased knowledge of the engineering staff over the past few years of working on the EVLA. Many of these engineers had previous RF or communication experience, but were not familiar with the particular problems of radio astronomy arrays. This transition is noticeable and commendable.

We see a vital need in radio astronomy for engineers or scientists with a broad knowledge of all aspects of the observations in the mode of the retiring B. Clark. This is needed for debugging of the increasingly complex and sensitive systems and for planning of future instrumentation. This knowledge includes: a) the types of sources that are observed, and their spatial, temporal, spectral, and polarization characteristics, b) atmospheric and RFI effects upon phase and noise, c) global understanding of hardware functions such as polarizers, lobe rotation, dynamic range, sampling, and frequency synthesis, and c) familiarity with software both for system debugging and image processing. Some suggestions for developing this expertise include the following: 1) encouraging engineering staff to attend internal (such as the recent New Initiatives Workshop) and external meetings (such as URSI) where instrumentation development is discussed; 2) short courses within NRAO to broaden knowledge, perhaps in the form of 1 hour lectures by experts in different topics; and 3) a scientific staff and student hiring policy which looks for applicants with broad areas of interest.

5.3 Testing

We recommend that more complete system testing be implemented in the AOC laboratory. Last year's committee recommended that equipment for two antennas be set up for laboratory testing. We note that there is presently a test rack for only one antenna. Perhaps this was due to the pressure of getting "first light" at the site, but we believe the laboratory tests are important for assessing phase stability contributions and measuring spurious responses and emissions from individual components of the system. Laboratory tests should include the effect of changing temperature internal to a closed module and due to orientation with respect to gravity.

6 Monitor & Control

6.1 Module Interface Board

The monitor and control system has adopted a bottom-up approach to the system, concentrating on the lowest layers near the hardware first. Supporting development of the hardware is a top priority, and the current state of the monitor interface board (MIB) software is up to the task. The status of the general MIB software appears to be very good. The design documents appear adequate and the demo at the VLA site looked good. The framework and service port documents appear particularly thorough and usable.

The use of UDP as a communications layer can result in dropped packets and a thorough examination of the impact of dropped packets should be given a high priority.

6.2 System Design

The overall system design lacks detail and is inadequate to use for staffing and schedules. The design is also insufficient to verify whether or not the required overall performance and throughput will be reached. The design does provides an overall scheme but has little detail on monitor and control. In particular, in the system documentation there are no interface definitions, the standard foundation for system design. Early development should be driven by the monitor and control point definitions for the individual MIBs, with some level of administrative control to give them stability. At this stage of the project, full estimates of monitor and archive bandwidth and volume are expected. Sub-arrays are an important construct, which should be addressed in the early phase of the design, since they have a significant impact on the monitoring and control architecture at all levels.

The high-level design mentions a multi-level hierarchy for the control implementation, but there was no evidence of further development of this concept. Top-level specifications mention a quality control feedback mechanism, which is indeed essential for complex instruments like the EVLA. This concept has to be better integrated in the design of the monitor and control system. The monitoring and control system is currently somewhat operator oriented, and it should be considered to what extent science users will benefit from a closer interaction with its higher levels.

6.3 Software Engineering and Testing

There was little information available on software or systems engineering. There are no apparent design and coding standards, or other guidelines for building a system or module. The high level design mentions the programming languages that have been used to date, but not what will be used for or why. Automated unit and integration tests are essential for maintaining a stable code base and should be included in the software development plan.

There is no integration and test plan. Integration and testing is very time consuming, and in the case of the transition plan it will be quite complex. Who does the testing, and exactly what tests are to be done, should be part of the development plan. Care should be taken to define a systematic test approach with prioritized testing requirements targeted at verification of the critical system specifications.

6.4 Staffing

The transition plan is very complex and it follows that it will be very difficult to estimate the staffing and schedule required for its implementation. The complexity derives from the underlying requirements, and the committee recommends that NRAO examine ways to relax these requirements in order to simplify this plan.

7 Correlator

The EVLA Correlator project is well organized and coming along well. The Canadian partners have done an impressive engineering job–the bandwidth and flexibility are outstanding. Like all complex signal processing projects, there are sure to be some problems along the way, but the planning and the expertise of the correlator team give us confidence that no really major problems are likely. We do have a few modest concerns as detailed below:

7.1 Hardware

The design team indicated that liquid cooling may be required, but the civil engineering for the correlator cooling plant is already done and the detailed HVAC for the correlator enclosure is in progress, so some coordination is needed on an urgent basis. Furthermore, the correlator room racking plan shows that locally the racks will dissipate about 500 W/square foot. This high a power density is generally found to be a problem in industrial data centers, unless cooling with local chilled-water heat-exchangers is used. (It is unclear to us just what is intended in this respect; if the "rack as a duct" scheme used for the old VLA correlator is planned to be used for the new correlator, we would expect no problem.) If individual rack chilled-water heat-exchangers need to be added to the baseline HVAC design, an increase in cost of about $50K might be expected. It is essential for system reliability that no local hot spots be allowed to develop within or among the racks.

The Correlator Board will require detailed "signal integrity" analysis because of the high density of traces. Also, as the design team already well knows, the clock and DC power distribution to the many FPGAs on the Station Board will need meticulous attention. We applaud the change to a "point to point" connection scheme on the Correlator Board and the decision to implement the FIRs using COTS FPGAs rather than ASICs.

Adding features such as flexible placement of bands and VSI interface is nice, but beware of creeping specifications.

7.2 Software

The software effort looks much improved since the last meeting. The new "memo 18" (not publicly available on the web) seems likely to go a long way to reducing software risk. Since software is being developed at Socorro and Penticton, the project management will need to be sure that an "us versus them" mentality does not develop, most especially given the software resource concerns mentioned elsewhere. Mode-free, rules-based design of the configuration software sounds like a good idea.

7.3 Schedule

There has been some schedule slippage, but it probably is not too serious at present.

The Correlator Chip CDR is scheduled for January 2005; it is important that this review does not slip if prototypes are to be available on schedule for testing. Vendor selection also should have occurred by the time of writing of this report.

Since Xilinx is shipping Virtex IV on schedule, the availability of these parts for the Station Board should not be a risk. However, will the required modeling tools be available soon enough so that the FIR Chip CDR can be held in February 2005 as scheduled?

The time scheduled for testing prototypes is 9 months; that might be enough (but it is notoriously hard to predict debugging time). The hardware team will certainly need to have good (i.e. expensive) test equipment in place when the testing of the prototypes begins in mid-2005 (lead times can be large).

Deferring work on the Phasing Board is a good idea in order to concentrate resources where they are most needed. But, deferral carries a high risk that the schedule will slip for the phased-up array output capability.

Concern was evidenced by the Canadian team over possible delays due to the procurement process that HIA must follow. The EVLA project management should therefore keep an eye on this aspect of the correlator development.

7.4 Cost

As noted by the Correlator Project management, variable exchange rates pose some budget risks, as does the assumption made in the plan that there will be no inflation in the cost of electronics over the next five years. We have no practical suggestions to what, if anything, can be done about this risk.

The Delay Module was reported as needing revision to reduce costs, but no information was given as to the impact if this cost reduction cannot be achieved. Similarly, it was reported that the Correlator chips vendor quotations were expected to come in "considerably" over the planned budget. We hope the budget consequences become clearer at the two CDRs schedule in early 2005.

Also as noted by the Correlator management, contingency fractions for the hardware are relatively small for such a high-tech project, and, in general, the entire budget is "slimly allocated." On the plus side, the FPGA solution for the FIRs may save money, as will the use of less expensive interface cables (assuming they work). Still it would be prudent for the EVLA management to plan for some modest cost over-runs just in case.

8 End-to-End (E2E) Computing

E2e software is meant to "wrap around" the Monitoring and Control (M&C) software and provide what is needed to deal consistently with the whole observing process from proposal preparation, scheduling, monitoring, archiving, on-line pipelining, and full off-line data reduction. The need for e2e software in the EVLA is important, and adequate resources need to be allocated to achieving the goals. However, the goals need to be realistic and compatible with affordable resources. At the moment, there is a big mismatch between e2e requirements/goals and available resources.

8.1 Synergy of Software with Other NRAO Projects

NRAO high-level policy is encouraging re-use of software on different instruments. This can be beneficial both for the optimized use of NRAO resources and for the resulting uniformity of software products. However, it does present complications and requires good management.

ALMA software comes into consideration here because of compatible requirements, general design, and timing of the two projects. The EVLA e2e software is a good case for synergy with ALMA software, provided that some basic concepts like Scheduling Blocks (SBs), Observing Modes, Data Models, and probably also technologies like communications and scripting languages are adopted by EVLA (after applying the necessary instrument specific additions to the design).

Re-use of software cannot happen just by expressing it as a goal for individual developers. It requires a precise commitment by the EVLA software management to establish this as a team policy. Software sharing (or re-use) may require compromises, and in this case we anticipate that the EVLA would have to do most, or all, of the compromising, given the more advanced state of ALMA development.

Limited resources on one project (e.g., EVLA) is an encouragement to collaborate. Still even this cannot work when the level of resources is too low, as it can result in "waiting to see" what gets produced by the other project (e.g., ALMA), without the necessary minimum interaction that would enable later re-use. Statements about the wish to re-use ALMA software, without having done the corresponding (same) design choices in well defined terms will not lead to convergence.

In summary, a general synergy policy exists at NRAO for software, but its implementation for the EVLA is immature.

8.2 System Design

Currently, there is not enough detailed design of the e2e software system. The design is developed at the top level, but subsystems are defined only to the level of general concepts. The next step is to define the subsystems (e.g., Proposal Preparation, Scheduling, Telescope Calibration) and their boundaries or interfaces. Note that this is not a detailed design of the subsystems, which can only be done afterward, but the first definition of their scope and interfaces, in order to assign responsibility to groups or individuals for the subsystems. This is urgent because the M&C team is currently designing from bottom-up, without knowing precisely the scope and overall system complexity.

In summary, a general design exists, but it is not sufficiently detailed to define either the subsystems or their interfaces.

8.3 Implementation Plans

The fundamental technology choices (frameworks) for the EVLA e2e efforts have not yet been made. These include the communications technology (eg, CORBA) and utilities (error, logging, alarms, libraries, tools) to be used by all the e2e software. This might be partly due to the still undefined design, although technology and frameworks are specifically made to be general enough to accommodate the needs of different projects. Frameworks often enable collaborations across very different domains (such as between astronomy and high energy physics). It is clear that while technology decisions will have to be compliant with requirements, there are several ways to fulfill them. A timeline should be defined to make a decision on the framework for e2e (which for ALMA Common Software is much more than purely a CORBA encapsulation). The decision should then be taken with technology advice, with resource or other constraints in mind, and should be explicit, e.g. where exactly and how collaboration with other projects is sought. In general the more one relies on other projects without contributing to them, the more one later pays in terms of separate development, to the point where re-use will not exist and duplication of effort will be necessary.

In summary, the e2e basic technology/framework has not been chosen, which hinders progress in the detailed design of subsystems and restricts collaboration.

8.4 Prioritized Requirements

EVLA staff should be praised for having produced a set of very detailed requirements covering all the range of e2e software, namely science, engineering, operations, real-time and data reduction software requirements. These provide a solid foundation for the e2e development work. At the moment the requirements are very ambitious, and are almost surely beyond the scope of the project. The committee agrees with the analysis presented by NRAO staff that it is important to analyze the current "Priority 1" (top) requirements and to come up with a minimum reasonable subset, taking possible re-use of ALMA software into account. Even a subset of the current priority 1 requirements might require a substantial increase of effort in the e2e (see next point).

In summary, determine a realistic minimum subset of Priority 1 requirements for EVLA.

8.5 Resources Required

While any project's resources are limited, the division of resources for the EVLA software appears unbalanced. Resources seem adequate for M&C software, with an estimated 80 FTE-yrs, but the e2e software will likely require more than the planned 30 FTE-yrs (which may not even be available). Additionally the staffing profile proposed for e2e is very low in the early years, and additional staff might become available from the M&C team too late (2008-2009) and possibly include staff with backgrounds unsuitable for e2e work.

As a reference, the NRAO part of the ALMA Computing effort is almost exactly the same as the EVLA software effort. However, ALMA has resources distributed quite differently, with more FTE-yrs in the e2e-like applications than in the monitor and control. The e2e software of ALMA can be estimated at more than 100 FTE-yrs, assuming the EVLA software metrics, with an additional 30 FTE-yrs invested in the common framework ACS (used both for M&C and e2e in the case of ALMA). Even assuming heavy re-using from ALMA there is still the need to have a minimum threshold of staff to contribute, so that they can start considering in detail and early enough what tasks are EVLA (instrument) specific. We recommend that resources are planned for at least one-third the level of corresponding subsystems for ALMA. In practice this might mean that for every EVLA subsystem (e.g. Observing Preparation, Scheduling, Monitoring/Exec, Pipeline, Archive, Telescope Calibration), where collaboration with ALMA is desired, there should be from 0.5 to 1 FTE available. More precise estimates can only be made when adequate design/technology choice are completed. The baseline planned resources might allow this to be achieved, provided staff are available early in 2005 at a level of about 4.5 staff/year for about 7 years. Below this threshold little software re-use/borrowing will be possible. Still it should be noted that sharing/re-use can exist only if choices in §8.2 and §8.3 are compatible with ALMA software.

In summary, we urge staffing the EVLA e2e effort at a level of at least 4.5 FTE/yrs. With less sharing of software from other projects, more resources will be required.

8.6 Software Staging

Given the current staffing, monitor and control software will be developed before e2e software. While this has advantages of making the telescope itself work sooner, it will probably require extra work to adapt the M&C layer to the rest of the project software. There is a risk of open-ended development in M&C, as the high-level interfaces are not defined and the technology to use in e2e has not been chosen. The M&C team at the moment works with internal standards, but without a general context for communications/ messages, errors, logs, alarms, data models, etc, for the entire e2e system. (We note that the EVLA M&C framework is different from the ALMA's, but we believe there are good reasons for this departure from the re-use paradigm.) The M&C group will have to adapt their output to the e2e needs, but may also have to change formats (models) and existing interfaces. While early choices made by M&C might make interfaces with e2e software more problematic, this is surely possible with some extra work.

In summary, the current EVLA software development may require a conversion layer between the M&C and the e2e software.

9 Postprocessing

It is clear that the AIPS++ post-processing software has made substantial strides on many fronts since the last advisory meeting. There is no doubt that this is largely due to the recommendations made by the AIPS++ technical review (March 2003), the challenge of the ALMA Computing CDR (July 2004), and the focus of the AIPS++ group to meet the ALMA software delivery schedule.

9.1 EVLA features

Management has opted for a more "project office" style of management, with the setting of project milestones and schedules for both development and delivery closely tied to the ALMA software schedule. This has been an important step in terms of transparency and accountability of the project. The close similarity between the core processing requirements being advanced for ALMA and those for the EVLA make planning relatively straightforward. However, it was not clear from the material presented what planning is in place for development of code targeted specifically for EVLA and not for ALMA. At some point in the not too distant future this needs to be considered.

9.2 Testing effort

Improving functionality, reliability and stability has been the recent goal of the project. This has been attained (seemingly in large part) through dropping the graphical user interface and establishing an integrated development and testing plan. In addition to an internal testing cycle, that allows for feedback from NRAO science staff, it also includes a number of external "testers," through which an attempt has been made to obtain an impartial opinion of whether or not the goals are being attained. Judging by reports from the external testers it would seem that reliability and stability are now much improved, and that the AIPS++ cookbooks and a number of the AIPS++ tools e.g. editing and imaging tools, are superior in many aspects to those available in AIPS and other packages. This is very encouraging to the committee and a significant step in the right direction, most particularly toward the eventual development of a much wider AIPS++ user base.

Given the test reports, it is unfortunate that the push to enlarge the user base for AIPS++ has been given a low priority due to the shortage of staff. We strongly encourage that the size of the pool of external testers is enlarged if at all possible. This serves the dual purpose of increased feedback to the developers, at the same time expanding the user base.

9.3 User Interface

Of moderate concern is that the testers note the current command-line interface is far from intuitive and the cookbooks are vital to being able to make progress. It is understood that the CLI is a short-term measure so that functionality and stability of AIPS++ can be pursued, and that development of the user interface is currently a low priority, awaiting further development of the "Framework." Given both the importance of the user interface and that prototyping of the "Framework" is planned for the coming year, it would be reassuring to hear more about development of the interface at next year's meeting.

9.4 Resources for EVLA AIPS++

An enduring concern of the committee is that the bulk of the current AIPS++ effort is driven by the demands of the ALMA software development, testing and delivery schedule. It is clear the EVLA benefits from this effort, with many aspects of the post-processing software requirements common to both ALMA and the EVLA. However, there are some significant differences in the requirements for the two arrays (e.g. RFI rejection, wide-field imaging), and the plan to carry out a post-processing "requirement-by-requirement" comparison for the two arrays is sound. We encourage the completion of this study so that the scope and scale of the EVLA-specific requirements are identified, and the impact on resources and project planning can be fully understood.

9.5 Priority 1 requirements

With the strong focus on post-processing for ALMA, there is some concern that the resources required to address the EVLA-specific requirements will not be available. A careful audit of the EVLA software requirements to ensure the priorities of various capabilities are set appropriately - akin to the ALMA Requirements Audit. We caution that such an audit should not become a mechanism to water down the requirements to such an extent that AIPS++ functionality is so severely limited that little is to be gained from its use. Once these true "Priority 1" requirements are established, a realistic assessment of the resources is then possible.

10 Descope Options

The growth of the project, particularly in the areas of e2e and post-processing software (AIPS++) since the phase I proposal, might require some descoping of other aspects of the project. However, before deciding on specific descope options, a careful balancing between the added capabilities of modern software and the loss of scientific capabilities through descoping other systems is needed.

At this point in time, remaining contingency funds would allow the projected completion of the full project. However, as it is early in the project timeline, NRAO is acting prudently to start to assess descope options, should they be needed later.

The simplest, most modular, descope options involve not constructing some receiver bands. The cases for three such bands were presented to the committee: S-, X-, and U-band. Each could provide savings of roughly \$1.5M, while minimizing the scientific impact of the resulting gap in continuous frequency coverage (which was a highlight of the EVLA concept) While dropping one band would produce perhaps a tolerable gap in frequency coverage, dropping two or more bands could open up huge holes in this coverage. This would likely have a strong negative impact on some high value projects, such as observations of molecules in high redshift galaxies and proto-galaxies.

Clearly a careful scientific trade-off study needs to be done to balance desirable software development, that benefits large numbers of users and NRAO in subtle ways, against the direct losses that would occur by dropping one or more frequency bands.

11 Phase II

The committee congratulates NRAO for submitting the phase-II proposal. The Decadal Survey highly recommended the entire project stating that

> "...The addition of eight new antennas will provide an order-of-magnitude increase in angular resolution. With resolution comparable to that of ALMA and NGST, but operating at much longer wavelengths, the EVLA will be a powerful complement to these instruments for studying the formation of protoplanetary disks and the earliest stages of galaxy formation."

The cost savings that would accrue from constructing both phases at the same time are significant, and we strongly endorse the completion of the entire EVLA project.