Project Response to Technical Review

Joe McMullin, NRAO
AIPS++ Project Manager

Steven T. Myers, NRAO
NRAO AIPS++ Project Scientist

March 28 2003

1 Introduction 1
2 Major Recommendations 1
3 Conclusions 12

1 Introduction

On 5-6 March 2003, a Technical Review of AIPS++ was held at the Array Operations center in Socorro, NM. On 17 March, an interim report from the panel was issued. This is a response from the Project Manger and Project Scientist to that preliminary report, and reflects the intentions of the AIPS++ Project team in regard to the changes proposed by the panel. In conjunction with this response by the project, the NRAO director and advisory committees will take the panel report as critical input to overall strategic planning in the area of observatory-wide computing and data management.

The eight members of the review panel were: Roger Brissenden (CfA), Hilton Lewis (Keck), Chair, Andrew Lumsdaine (U. Indiana), Dave McConnell (ATNF), Steve Scott (OVRO), David Silva (ESO), Doug Tody (NRAO) and Rick White (STSCI). First and foremost, the AIPS++ Project would like to thank the panel for their participation in this review and for producing the interim report on the extremely short timescale needed to present it at the ALMA Computing PDR. We feel that the review discussions and report contain an impressive set of useful suggestions that will be extremely helpful to the project, and as we state below we fully accept this advice and plan to implement the proposed changes in a timely manner. We hope that the panel members will agree to participate in a follow-up review on the timescale of approximately one year, and look forward to continued interaction as we make progress in the coming years.

2 Major Recommendations

The following are the Major Recommendations of the panel taken from the Interim Report of 17 March 2003, along with our responses (in *italics*):

1. **Change the focus of AIPS++ from a system aiming to be a general-purpose radio astronomy package to a system that is developed to meet the strategic goals of the Consortium.**

   Until now, the aim of the AIPS++ project has been to develop a general-purpose package for the analysis of radio astronomy data. The intent has been to provide tools to allow the analysis and processing of data from current and future radio telescopes, including the VLA, GBT, ALMA and EVLA. After 10 years of development, the functionality of the system is approaching that required to analyze a subset of existing VLA science cases, but with performance, usability and some code defect issues. The system is not yet at an appropriate level of maturity to deliver to users.

   We believe that the major reason for the lack of delivery has been the focus on producing a general-purpose system at the expense of producing software with specific science goals driven by actual radio astronomy analysis use cases and data.

   In order for the Consortium to meet its analysis and data-processing software goals, the approach to developing AIPS++ must be driven explicitly by the needs of the current and planned major radio astronomy projects. This implies a modified development process consisting of

   - science staff identifying use cases and developing requirements (including performance requirements) from the use cases
   - development staff flowing the requirements to the software design and developing a build plan with content and delivery tied directly to the project milestones
   - a test process that involves science testers using test data based on the use cases
   - software acceptance by the relevant science or project lead

   This approach is fundamentally different from the current AIPS++ development paradigm and implies tight coupling between a master schedule and the AIPS++ software deliveries.

   Project Response: We agree that we need to move toward a more project-oriented development methodology. The system of project requirements with a lifecycle including drafting, auditing, planning and acceptance (based upon the ALMA Computing model) that we started last year is a step in this direction. We plan to extend this to all projects that AIPS++ serves, making use of the various scientific staff and user groups available in the Consortium.

   ALMA is a high-priority project and will be a major driver of the future software development. In principle, it is preferable to examine high-priority use cases from

all Consortium projects and derive an integrated set of requirements and build content from these. In practice, however, this diffuse focus, coupled with a small and distributed project team, is unlikely to succeed. We propose instead that the task of completing sufficient functionality to demonstrate a full set of VLA use cases be viewed as a high priority, not in order to satisfy users (who are generally using AIPS), but in order to complete a set of core functionality required for ALMA. A subset of complementary use cases of other Consortium instruments (such as ATCA and BIMA) should also be incorporated in order to match remaining core ALMA requirements. Examples of these requirements are complex spectrometer specification and configuration, mosaicing and linear polarization.

We see the completion of the VLA use cases and a carefully selected subset of complementary use cases from other Consortium instruments as an essential part of the ALMA development.

Project Response: We agree that the focus on the fulfillment of core VLA (and ATCA, BIMA) use cases aimed at core functionality applicable to ALMA is an excellent way to integrate the toolkit, demonstrate competence, and to harness our scientific testers. Verification of AIPS++ for select end-to-end VLA processing modes will also test a wide range of critical systems in the package, touching the deepest code levels. In our current requirements process, we had relegated use cases to a supplementary role with respect to the bulk categorized requirements – we will move the drafting of use cases up in priority and build our development and release plan around the fulfillment of end-to-end complete processing for these modes. We propose to deliver the first sets of use cases by June 2003 and to fold these into the following development sub-cycles. Note that the NAUG has produced a draft of an AIPS++ VLA Audit which can form the basis of the use case auditing, and thus we might hope to have some simple modes (e.g. continuum single-field) verified by October 2003.

The support for external non-Consortium users should be de-emphasized until an accepted system can be completed. We recommend taking AIPS++ off-line until sufficient functionality has been demonstrated per use cases. At present, VLA users are generally using AIPS, and the GBT users could be considered as the first target users of the system.

Project Response: De-emphasis of user outreach and support for general users will relieve some of the extreme pressure the AIPS++ developers find themselves under in the current project-driven climate. Of course, we plan to continue user support for defect resolution with priority given to defects in the systems that span the project-critical parts of the package. We appreciate that the panel recognizes the time-

critical nature of GBT support, and we plan to give high priority to user support for this instrument.

We believe that adopting these recommendations can, in time, lead to the AIPS++ goal of providing a general-purpose radio astronomy package, a package that will provide a general-purpose toolkit, cater for a variety of instruments and provide a good environment for algorithm development.

2. Modify the Development Process

In order to deliver software tied to the specific milestones of the Consortium projects, the present software-development approach needs to be modified. Specifically, we suggest considering incorporation of the following elements in future development:

- Develop use cases for the major project drivers
- Derive requirements from use cases (including performance)
- Develop build schedule with explicit release dates with incremental functionality tied to key project milestones
- Develop design and hold design reviews
- Develop tests, test data
- Develop code and documentation
- Peer review code
- Test using data tied to use cases
- Acceptance of release by Project Scientist

Many of these elements are currently present, but are not focused on the delivery milestones derived from the needs of the major Consortium projects.

Tying the content and deliveries to the projects may result in the need to abandon the current 6-month routine delivery cycle, which is more suited to a system in long-term maintenance than one undergoing significant development. The recent adoption of mini-deliveries each on 4-6 week timescales is a move in this direction.

A critical aspect of developing the AIPS++ build schedule is creating a master schedule tied to the needs of the Consortium projects. The resultant AIPS++ master schedule should be consistent with the major project schedules (such as GBT, ALMA, EVLA).

This approach would allow the answer to the question “What’s AIPS++ needed for” to be something like:

- Perform analysis of {BIMA, GBT, ALMA, LOFAR, etc.} data per the use cases and requirements documented in w, x, y and z
- Support pipeline processing of {VLA, ATCA, etc.} data per document x, y

The approach would have an associated set of processes that would be applied to all development projects, e.g., a process for generating a use case, by whom, how to ensure its part of full coverage, etc.

Project Response: We plan to move towards a process where there are clear milestones with deliverables scheduled to correspond to project deadlines and releases. The new system of monthly build plans is a step in the right direction, but we need to integrate all the projects into a unified scheme. We intend to start the proposed changes as soon as possible, starting with the drafting of the VLA use cases (NAUG), creation of a unified build schedule in the new Project Office, and building towards a review of our progress approximately a year from now.

3. Strengthen the Project Management and the Project Team

The project team is clearly highly talented, energetic and dedicated. Although these are essential ingredients for success, by themselves they cannot result (and indeed have not resulted) in achieving the goals of the AIPS++ Project. In particular, it is evident that there are a number of weaknesses in the management of the project that must be addressed.

Project management must be strengthened in a number of areas. The Panel was struck by the lack of standard project management methodology and reports. Suitable formal project management techniques should be adopted, both for managing the project and communicating the state of the project to the Consortium management and partners. Key metrics such as earned value and progress against project milestones should be adopted, in order to show progress in a meaningful way and to allow for adjustment in scope and schedule as warranted.

Project Response: We recognize that substantial improvements need to be made in our management practices and methodology. The move toward project-oriented strategic goals with releases timed toward project milestones will necessitate changes. We welcome suggestions from the panel on the proper techniques to adopt, and we particularly found that discussions during the review on this topic were highly useful!

The current practice of sharing the Project Manager role between several individuals, none of them on a full-time basis, presents a problem. For a project of this magnitude the Project Manager role must be recognized as a full-time position, one that cannot be shared with other duties and responsibilities.

Project Response: We plan to reduce (and eventually eliminate) the development duties of the Project Manager as indeed we agree this

must be a full-time position, particularly given the suggested changes. However, we also note that some of the management duties, such as the Consortium site managers, must continue to be distributed among various personnel. We do not see this as a conflict with the panel's suggestions, and this will also relieve the Project Manager from some of the more mundane management burdens and allow them to focus on the more important strategic issues.

Project management must foster a clear customer focus among the project team members, where the customers are the major Consortium programs requiring AIPS++. This is a distinct change from serving an amorphous general user community. The development of an AIPS++ master schedule tied to the Consortium projects will assist in this regard.

Project Response: We fully agree with this suggestion, and plan to go further in having a Subsystem Scientist who is assigned as a direct liaison with the developers for particular parts of the projects. We have found over the past five months that contact between the project team and the users (e.g. the NAUG) helps foster customer focus, with a noticeable increase in morale in the project and the testers.

Every effort must be made to fill current vacancies. The team should be strengthened by the addition of professional software engineers; sufficient astronomical expertise is already present in the existing project staff. One of the project staff positions should have the role of software architect, responsible for the overall integrity and consistency of the core classes. This is especially important, as significant further development will be required in the future.

Project Response: We plan to fill one, if not all, of the currently open positions with software engineers, or with astronomers with professional software engineering expertise.

More effort must be made to utilize full-time FTEs. Staff allocated at a small fraction of their time over an extended period are inefficient, and in the case of marginal involvement (<10%, of which there are several) they often amount to a drain on overall project resources.

Although this program could benefit from additional personnel, the Panel believes that the current team (utilizing all currently open positions, of which there are several) should first demonstrate their capability to use what resources they have more effectively. We also note that several Consortium partners are currently contributing significantly less staff than that required of them. This is a matter for the AIPS++ Executive Committee to resolve.

Project Response: We agree that fractional FTE personnel can be ineffective (though we also recognize that we have benefited from even

small FTE fractions from certain exceptional individuals!). We will ask the AEC to consider this issue carefully. We feel that with the reorganization last year, followed by this Technical Review, we have an unprecedented opportunity to reinvigorate the AIPS++ Project, and the Consortium in particular. The move toward project goals, which should allow more effective use of Consortium resources and a better ability to fulfill Consortium members' project requirements, will hopefully be incentive for some of these improvements to happen. We also acknowledge that we need to demonstrate more effective use of existing (and soon to be available) personnel before asking for increased resource allocation. We note that our resource requirements will become much clearer as we obtain and audit the full set of project requirements, and in the coming years we can make a strong case on this basis for (or against) increased commitment.

4. Further Strengthen the Project Scientist Role

The Panel applauds the appointment of a strong and active AIPS++ Project Scientist. It is clear that this has already had a significant and positive impact on the project. The Panel also wishes to acknowledge the recent activities of the NRAO AIPS++ User Group (NAUG). The interaction of the NAUG and the project in recent months has been clearly beneficial, and demonstrates the need and power of actively involving users in the development process.

Project Response: We appreciate the support that the panel expresses for the work of the NAUG and our other user groups. We also would like to acknowledge the important role that these scientists have played over the past 3 years in bringing user concerns to the attention of the project and paving the way for the substantial improvements made in our user integration process over the past few months. We note that the proposed changes will require even more work from groups such as the NAUG, and hope that the visible successes shown at this review will build enthusiasm among the testers for a more active presence. In particular, the Project Scientist will be relying upon the support and input of the NAUG (and parallel groups in the Consortium) to help in carrying out of the plans suggested below!

The Panel believes the role of the Project Scientist is critical to the success of the project and, therefore, needs to be expanded in a number of ways.

Analogous to the Project Manager, the Project Scientist should be appointed at the Consortium level and be responsible to the AIPS++ Executive Committee. Individual Consortium sites and projects should appoint individual Site Scientists. The activities of these Site Scientists would be coordinated by the Project Scientist. This is similar to the user group structure presented by Steve Myers during the review, but formalized in the following ways.

In coordination with the Site Scientists, the Project Scientist should have the following responsibilities:

- Develop science-based use cases for all AIPS++ supported major Consortium instruments.
- Based on use cases, derive project science requirements. Note that science requirements should include baseline performance requirements.
- Identify and appoint Consortium staff scientists as use case testers. It is expected that these use case testers will be active scientists with research experience with their home instruments.
- Organize and prioritize use case tester feedback for project action.
- Verify that project deliverables fulfill use-case-derived requirements.
- Organize development of high-level end-user User's Manuals ("cookbooks") for major scientific threads.

In coordination with the Project Manager, the Project Scientist should have the following authority and responsibilities:

- Provide scientific prioritization for all project activity, based on clearly defined scientific use cases, and subject to the high-level project priorities established by the Consortium Executive Committee.
- Assure that project activity is aligned with the scientific requirements.
- Certify and approve for public release any new version of AIPS++, after the Project Manager is satisfied that a build is technically ready for release and after the satisfactory testing of the scientific use cases.

Technical decisions should continue to lie with the Project Manager. The Project Scientist will generate use-case-based requirements, but it is up to the Project Manager to decide how to achieve those requirements.

Project Response: We plan to implement the suggested methodology and have augmented our User Integration Plan (presented at the review) to include these changes. The re-introduction of the Project Scientist into the AIPS++ Project has triggered a number of changes in our process over the timescale of only a few months, and we are gratified to see that these were seen to be in the right direction. We agree to the list of duties, authority, and responsibilities presented by the panel, and plan to ask the AEC to approve the establishment of an overall Project Scientist and Site Scientists to carry out this plan. We propose that the current NRAO Project Scientist (Steve Myers) act in the role of overall AIPS++ Project Scientist in the interim for the timescale of the coming year, in order to facilitate the extension of the current new process to include the proposed changes.

5. Short-Term Priorities

Development efforts over the next 12 months should focus on the following specific issues:

- Significant improvements in reliability and stability
- Performance improvements approaching the performance of AIPS
- Synthesis imaging capability covering many common use cases of the VLA
- Reduced emphasis on graphical user interfaces

The Panel endorses the current systematic approach to investigation of performance issues in AIPS++ presented at the review. Benchmarking performance of components and feeding back problems to the developers for resolution has been quite productive. It is important that the planned "best practices" manual cautioning developers against performance-robbing practices be produced soon. It seems likely that continuation with this program will increase performance to the point where it approaches that of AIPS.

There are defects in the design of many of the present GUIs, some related to the underlying implementation of Glish and others to usability issues. However, script-driven processing is adequate to demonstrate the core functionality of AIPS++ at this stage. A Viewer GUI is recognized to be essential for this mode of operation. Other graphical interfaces are required in the longer term, but should have a lower priority for now. Redesign of existing GUIs will be required, but should await resolution of the underlying GUI performance issues. The project intends to achieve the needed level of performance through a re-implementation of significant parts of the current technology. However, other less ambitious options are also available (e.g., based on the Qt widget library), and should be considered if the implementation of the new technology takes longer than planned.

Project Response: We agree with these short-term goals and plan to adjust our current development plan accordingly. In particular, we plan to immediately focus on critical defect resolution and reliability improvements, and to continue and ramp-up our benchmarking and profiling efforts. Development of use cases and targeting of end-to-end fulfillment of key modes will follow. We acknowledge that usability improvements such as GUI redesign must come second to alleviation of reliability and performance problems that are often show-stoppers to the use and testing of the package, though we do plan to start the investigation of GUI options in order to make it possible to achieve progress in this area a year or so down the road – in particular we plan to hold a NAUG focus group to storyboard possible integrated GUI look-and-feel. Note that these efforts will not detract from the suggested high-priority improvements.

Future releases must include user-level "cookbooks" for processing data with validated use cases. Given the toolbox nature of AIPS++, such cookbooks are the

most effective way for non-expert users to use the system successfully and productively.

In the case of the NRAO, the VLA provides an ideal test bed for AIPS++, not only because of the many common observing modes, but also for organizational and sociological reasons. The Array Operations Center co-locates many of the AIPS++ developers and Project Scientist as well as a large pool of scientific users of the VLA. The recent increase in involvement of the scientific staff is very encouraging as is the quick turnaround on bug fixes. The effort to identify the VLA scientific use cases needs to be completed along with systematic testing of these use cases. Successful achievement of these efforts opens the door for an AIPS++ VLA pipeline, which would complement the proposed archive. Routine operation of such a pipeline would give much needed credibility to AIPS++ and confidence that it could fulfill its Consortium obligations.

Project Response: We would again like to acknowledge the contributions of the NAUG over the past years, in particular the preparation of user cookbooks and the auditing of the VLA requirements (led by Debra Shepherd). The recent successes in this area are due in no small part to scientific staff involvement and we plan to continue these activities in the coming years. We are sure that the NAUG will applaud the focus on VLA testing and end-to-end processing capability, and agree that fulfillment of these goals should build confidence in the ability of the package to carry out its mission for ALMA and EVLA and beyond.

6. Proceed With the Proposed Technology Changes

The current AIPS++ architecture is built around Glish, which serves as the software bus for task communications, the task control system, the scripting language for high-level applications, the command-line interface (CLI) for user interaction and the GUI system. Since the AIPS++ system has its own custom approach to each of these needs, it is effectively a closed system that cannot easily benefit from the many developments in the rest of the software world. In fact, many of these functions might be better performed by standard components developed in the software community over the last 10 years.

The proposed new technology appears to be a sound basis for changing AIPS++ to an open system that uses industry standards where appropriate (e.g., CORBA, Java) and allows the use of robust open-source software in place of AIPS++ software (e.g., Python instead of Glish). By moving Glish out of the core of the system, a major bottleneck for improving the performance of GUIs is also eliminated. Using the ALMA Common Software as the basis for the bus is also appealing, since that software is already in use within the project.

The case made for a new component-based distributed architecture is strong. A preliminary analysis of implementation technologies has been performed. This analysis should continue and should be the major focus of any proof-of-concept. In particular, the Alma Common Software (ACS) appears to be well aligned in terms of architecture and technology. However, the current ACS implementation was designed for a different purpose and may not be suitable for use within a data analysis framework without significant rework and extension. Collaboration with ESO on a common "ACS-light" core could be beneficial to all concerned, but there may be difficulties in reconciling the development and release needs of the ACS in support of ALMA with the needs of the AIPS++ project. The project should prepare for the possibility that they will be responsible for the work to integrate ACS (or other technology identified by the feasibility study) into the AIPS++ system. A common software framework is desirable, but it may prove to be impractical to share this core software with other systems.

The first step in considering the technology change should be a feasibility study that evaluates alternative technologies relative to the long-term project requirements. It is essential that this evaluation not interfere with the ongoing improvements in functionality, robustness and performance required for the AIPS++ tools and applications. Parallel efforts focusing separately on completing the current AIPS++ science data-processing capabilities and prototyping a new system framework are therefore called for. However, the staffing for the current AIPS++ applications group must not be reduced; moreover the effort of evaluating and implementing the new technology will require staff with a strong software engineering background, rather than that of astronomical data processing. These requirements point to the need for additional resources, probably carried out by a group that is independent of the existing project team.

Project Response: We are gratified to see that our proposed technology upgrade, and the proof of concept in particular, is seen by the panel to be a reasonable approach. We acknowledge that the current proposal is not a real design, and plan to build from the proof of concept tests to a full design followed by a design review (possibly on the timescale of one year, coinciding with the revisiting of this review). We agree that pursuit of this technology change should not interfere with the other critical improvements (e.g. in reliability, performance, end-to-end capability, and management) but we also point out that the personnel involved, for example in the proof of concept, are not the same as those currently involved in the core project due to the different expertise (particularly in software engineering) required. Our plans for investigation of such a design will respect the needs of the rest of the project, and the changes suggested by the panel.

3 Conclusions

The AIPS++ Panel proposes the adoption of an ambitious set of changes to the process and methodology of the project. However, it is gratifying to see that the vision that we presented at the review was seen to be generally in the right direction, in spite of having been developed and implemented in many cases in the short span of four months! We look upon the proposed action items as a logical extension of our current plans, rather than a complete change in direction. Furthermore, we note that the panel has found no fundamental flaws in the architecture of the AIPS++ package, rather that the toolkit still has problems in some areas with poor reliability, performance, and integration between components. Thus, we conclude that the core classes and tools can form a sound basis for fulfilling our project requirements provided we augment our management and development processes in the manner advised. Finally, we are encouraged to see that the panel feels that our proposed technology and framework upgrade is a reasonable approach to further investigate, and we will continue with the plan to develop a proof-of-concept and eventual design based on this.

As stated above on a case-by-case basis, we agree with the panel's assessment and advice and plan to implement these on as early a timescale as is feasible. Although we note that the reduced staffing (e.g. lost to increased management responsibilities, the loss of recent personnel and the delayed availability of the new positions for work in these areas) will mean that our progress in the actual implementation of code improvement will be slower than might be desired otherwise, we do think that if we correctly implement the proposed plan we will show significant improvement on the timescale of a year. Therefore, we agree that a re-review around Q2 of 2004 is in order. In particular, the Project Manager and Project Scientist appreciate the confidence in our efforts shown by the review panel. However, we also acknowledge that we should be judged not just on the plans we have presented at the review, but upon the progress made over the coming year, and that the proof of the effectiveness of the new project leadership and our ability to carry out the necessary changes will lie in our performance during this time, and we therefore welcome the prospect of coming back and demonstrating real improvements.