ATNF Scientific Computing Group

Response to ATNF/ATUC Review of 3 May 2005

Lister Staveley-Smith, Vince McIntyre
17 August 2005, updated 19 January 2006

Background

A review of the activities of the ATNF Scientific Computing Group (SCG) was held on 3 May 2005. A number of recommendations were made. This document responds to these recommendations, and summarises a possible future direction for SCG activities in the near term. A separate document itemises the proposed work breakdown for the coming year.
The current SCG structure is listed at:

http://www.atnf.csiro.au/computing/computing_marsfield.html

and its current activities are listed at:

http://www.atnf.csiro.au/computing/activities.html

Overview

We generally support the review panel's recommendations, namely:

- That the SCG is a valuable part of the ATNF and should be maintained with critical mass.
- The core business of the SCG should be:
    - Maintaining and developing existing data reduction packages.
    - Acting as a coordinating body for ATNF software standards and data archiving.
    - Strategic algorithmic (for xNTD, SKA, etc) and standards (e.g. WCS) development.
- That the development activities of the SCG should be better defined and, where possible, better aligned with the core activities of the ATNF, via the now established ATNF project management structures.
- That the level of strategic development of software for large projects such as the xNTD be increased significantly over the next 12 months.

Response

The detailed responses to the panel's recommendations (*panel recommendations in italics*) are as follows:

1. *From the presentations made during the review, it was apparent to the panel that the work of the members of the SCG is of a very high standard, is well*

regarded within the ATNF user community, and is acknowledged internationally. The group members are to be commended for their application to their assigned tasks.

Thank you!

2. *The panel concluded that the SCG is a valuable entity within the ATNF and that this group should be maintained with critical mass, composed of people that have skills and experience in both the astronomy and computing domains. This will be particularly true over the next 5 years as projects with large software components such as the CABB and xNTD are undertaken. The advantage of the SCG is that it forms a competent technical and practical resource that can be drawn upon by ATNF projects, as well as a self-contained strategic resource for the ATNF. How this group should sit within the overall ATNF organisational structure is best addressed by ATNF management.*

ATNF management has decided to maintain the SCG under Astrophysics line management for the year 2005/6. This is in order to reflect the scientific role of the group, rather than its operational role. For logistical reasons it has not been possible to separate out all operational support from the new SCG (e.g. some group members have small but not transferable support duties in maintaining the Telescope Control System), although this may be able to happen in the future. Similarly, certain people with useful skills in the scientific computing domain are not under SCG line-management – e.g. Tim Cornwell, Albert Teoh, Christopher Owen, Mark Wieringa. It is likely that this will continue – i.e. people with substantial responsibility for a particular project may be line-managed by the project leader. In response to the latter situation, the SCG is now convening fortnightly scientific computing meetings to improve communication between Marsfield computing scientists.

3. *The panel was concerned that, while the output of the group was very good, the scope of the activities undertaken was in some ways poorly defined and could be better aligned with the focus of other ATNF activities. The panel discussed the purpose of the SCG within the ATNF, in particular looking forward to future large software projects. The panel concluded that the purpose of the SCG was to focus on activities that contain a significant research and development component and that support core ATNF objectives. The panel suggest that the core activities of the SCG should include:*

    a. *Maintaining and developing existing data reduction packages, in particular extending algorithms for existing and future ATNF observing facilities.*
    b. *Playing a coordinating role in implementing software standards and policy across the ATNF, for example ensuring that ATNF data archiving policies are implemented uniformly across all data types (ATCA, single-dish, continuum, spectral line, pulsar etc).*
    c. *Strategic algorithm development aimed at near term (e.g. xNTD) and far term (e.g. SKA) instruments.*
    d. *Strategic standards development (e.g. WCS).*

Our proposal (partly summarised at the ATUC meeting on Dec 1, 2005) is to maintain existing packages at the following levels. A more detailed description of support level is not yet available, as this depends on available resources:

- AIPS: Standard NRAO AIPS to be maintained on a limited number of linux and solaris systems at Marsfield, Narrabri and Parkes, mainly for VLBI purposes. No support for Mac. Limited support for peripherals (e.g. remote tape decks).
- MIRIAD: regular maintenance at ATNF sites and support for remote ATNF users (linux, solaris OS). Increase coordination with BIMA to make use of useful developments arising from CARMA. Reduce the numbers of problems associated with out-of-date version of MIRIAD[1].
- AIPS++: stable version to be minimally supported for general use on a single operating system (debian linux, currently). Mainly to be used as a development environment. No support for Solaris or Mac, except for code trialling.
- LiveData/Gridzilla/ASAP: : regular maintenance at ATNF sites and support for remote ATNF users, including binary builds. As with AIPS++ support only available for linux.

Future packages and strategic algorithmic research will be developed as follows:

- LiveData/Gridzilla: algorithmic support to be provided for GASS and methanol multibeam surveys.
- ASAP: version 2 has been approved by the ATNF Project Review Board and will provide greater functionality and be more generally useful than the existing version 1.2.
- xNTD: Algorithmic research will proceed using funds from the CSIRO Emerging Science Program and funding that will hopefully be allocated to the xNTD theme in 2006/7. Maxim Voronkov will commence a 2-yr xNTD software scientist position in March 2006.
- WCS: Extend WCS to handle coordinate system distortions.
- Source finding and characterization: we intend to work with the User community on software which may be useful for the GASS, methanol, xNTD and other surveys.

It is envisaged that the SCG will increase its activity as a coordinator of software standards and policy. This will occur partly through the fortnightly SCG meetings, the regular computerfests and through the revamped Observatory Computing Committee meetings. Internationally, SCG members will use the ADASS meetings to discuss software standards and future directions

> *Of the activities listed in the summary document, the panel felt that some activities did not fit the core purpose of the SCG and should be transferred to ATNF Operations. These activities are: telescope control and monitor (ATOMS, TCS); compute facility.*

[1] Currently this is not resourced.

The TCS activity cannot be transferred out of the SCG currently (it occupies only 0.1 of an FTE, but requires specialist knowledge). Similarly the compute facility (although it's possible that the fraction of the compute facility associated with visitor support could be completely moved to the Operations group at a later stage).

> *In addition to the core activities that the SCG should be involved in, the panel discussed the interaction between the SCG and the ATNF project management process. The panel felt that, in addition to the core activities of the SCG, that SCG members should be available to participate in project-based activities, as required. The panel felt that a clear demarcation between SCG core activities and ATNF project activities would make resource management more transparent. In addition, since the aim of the SCG is research and development orientated, the SCG members, where appropriate, should be given a fraction of their time for self-directed research. This should be accounted for in any resource management model applied to the SCG.*

The ATNF is committed to the clear identification of project-based activities and the tracking of these activities through the project management system. SCG members will devote a significant fraction of their time to future project-based activities, as opposed to operational support and maintenance activities. Members may be seconded to projects if a sizeable fraction of their time is spent on a project.

Other Issues

The Review panel responded to a set of specific questions relating to the performance of the SCG and the ATNF Computing staff in the areas of telescope operations, strategic developments and management. The following specific issues were raised:

> *Users raised a number of difficulties with SPC and have called for a replacement package over the last few years. Development of the new ASAP package was undertaken to address this User request and the software is now almost ready for general release. The uptake of ASAP over the next 12 months should be monitored to determine if the objectives of the SCG and the ASAP project have been met.*

The first phase of ASAP appears to have met most of its objectives. Version 1.2 is downloadable from http://www.atnf.csiro.au/computing/software/asap. A major new version is being prepared which will have further functionality, e.g. polarisation handling. The ALMA consortium have been evaluating ASAP as the basis of its single-dish reduction software (ALMA will be equipped with four 12-m dishes for total-power calibration). See http://almasw.hq.eso.org/almasw/bin/view/OFFLINE/SDEval.

> *The panel felt that maintenance and development of off-line data reduction software was something that the SCG did very well, in general. To ensure that this high level support of existing facilities continues, the panel felt that development of off-line data reduction software should be a core activity of the SCG.*

It is planned that the SCG and other ATNF staff will continue to develop off-line data reduction software.

> *The danger is that the evolution (or replacement) of these packages for the support of instruments like the xNTD are not receiving the attention that they warrant. The panel therefore felt that the strategic development of data reduction algorithms should be a core activity of the SCG. This development may need to be directed from within the SCG since existing project plans may not be able to fully identify the software requirements for these projects.*

ATNF plans for xNTD algorithm development have received a significant boost with funds from the CSIRO Emerging Science Program having been awarded. Significant xNTD funding is expected to flow in 2006/7. Initially at least 1 extra FTE has been hired in this area (Maxim Voronkov commences in March 2006). It remains true that project leaders and scientists often underestimate software requirements. One of the aims of the Project Review Board is to ensure adequate project resources are being requested.

> *A further issue going forward is the fact that MIRIAD does not support 64-bit operating systems and significant development work may need to take place if MIRIAD is to be used on new systems.*

There are several issues connected with MIRIAD development which the ATNF is still addressing. These include 64-bit code, file size limitations, coordinate systems and image displays.

> *The panel recommends that RVS be in effect frozen in a state that will allow it to be used as part of the ATCA online archive/pipeline system. Further, the panel recommends that software developments for data visualisation be highly integrated into projects that support the core activities of the ATNF.*

RVS development has been frozen. Whilst lifetime planning may have been lacking, we nevertheless have a sophisticated visualisation package which, currently, is easily maintainable and is currently being utilised as the visual interface to ATNF data products. We are exploring whether further development of the package can be done by other institutes.

> *The panel felt that in the interests of uniformity and efficiency (and with the benefit of hindsight) that a better outcome could have been achieved if the high level policy on ATNF data archiving (all data types) had been passed to the SCG, with the responsibility of developing and enforcing a uniform method for archiving and archive interfaces. While data archives should be made available for all ATNF data types, and uniform interfaces are an issue for usability, the deeper issue here is the procedure for the implementation and delivery of products relevant to ATNF policy (that fall under the purpose of the SCG).*

The panel thought it reasonable that this type of coordination and responsibility for implementation of high level policy (where appropriate in the software arena) be considered a core activity of the SCG.

Despite the disparate types of data available, and the different access regimes required, the issue of uniformity of ATNF data archives is important, and needs better coordination in future. We are attempting to ensure the newer interfaces such as to the pulsar archive have a similar look-and-feel to ATOA. The immaturity of VO standards is not helping the uniformity of data products.

The current level of strategic or blue-sky software research will therefore likely be unacceptable 12 months from now, far more will be required. As stated above, the panel felt that this would be a natural core activity for the SCG.

In the last 6 months ATNF has embarked on two new blue-sky developments. One is research into wide-field imaging, directly related to xNTD. The second is research into source detection algorithms, particularly in three (or higher) dimensional datacubes. This should have applications to existing surveys, and to xNTD. Each of these areas has one full-time postdoc allocated.