SKA Signal Processing

Concept Design Review

(SP CoDR)

14th to 15th April 2011

Report of the Review Panel

Robin Sharpe (External advisor UK) (chair)

Peter Dewdney (SPDO UK)

Andre Jongeling (JPL USA)

Phillip Stanley-Marbell (IBM research CH)

Lou Morales (iSine USA)

Tim Hankins (New Mexico Tech, NRAO USA)

Summary

The key findings and recommendations of the Review Panel are summarized below:

1. The preparation for the review was of a high standard including both the documentation, which was distributed before the meeting, and the presentations given during the review itself.
2. It was clear that the required signal processing concepts and algorithms were mature and in general very well understood.
3. Several alternative architectures are available with realisations of these utilising technologies ranging from software on a general purpose computer or graphics processing unit to FPGA and ASIC implementations.
4. All of the solutions presented appeared to be feasible, although estimates of cost and power consumption showed a large degree of variation.
5. In moving into the next phase it will be urgent to adopt some baselines to allow proper comparisons to be made for the alternative architectures that are available.
6. The requirements for pulsar signal processing appeared to be less well developed and there was much less implementation experience evident. This area will need stronger attention in the next phase to bring it to the same level of maturity as the other areas.
7. Taking into account the preceding comments, the Panel believes the Signal Processing Element is ready to move into the Definition Phase.

1 Introduction

The SKA Signal Processing Concept Design Review (CoDR) was held on April 14th and 15th, 2011, at the University of Manchester, Manchester, UK.

This meeting reviewed the progress within the Signal Processing Element of the SKA project versus the Project's CoDR milestone and assessed whether the Project has achieved sufficient maturity to move ahead into the next phase.

The CoDR Panel consisted of four members external to the Project with experience across the fields of Radio Astronomy, Signal Processing, Consumer Electronics and Semiconductors. In addition to these the Panel was also joined by the Project Engineer from the SPDO. The Panel membership is given in Appendix 1, and the Panel's charter in Appendix 2.

The Panel's initial observations were fed back to the SKA Director, the SPDO, presenters and observers on April 15th. This report outlines and further details the observations and recommendations made by the Panel.

The Panel was in full agreement in its assessment of the Project together with the comments listed in this report.

2 Preparatory Documentation

The preparatory documentation for the review, which is listed in Appendix 3, was of a high standard and consisted of 27 documents totalling over 850 pages. These were distributed to the review Panel about ten days before the meeting. Many of the documents were technically detailed and the review Panel acknowledges and thanks the staff in the SPDO and the contributing organizations for this standard of preparation.

There was an opportunity for the Panel members to submit written questions prior to the meeting and the Panel would also like to acknowledge the fast turn-around of the answers to these questions.

3 Overall Progress

It was clear to the Panel that the signal processing concepts and algorithms were generally mature and very well understood. Many of the presentations described specific implementations together with their architectural and implementation technology choices as well as detailing performance results for these. The Panel made the following general observations:

- Of the solutions presented, there was much more emphasis on correlators, and relatively less information presented on "non-imaging" pulsar processing.
- The Panel found that the performance results were not easy to compare directly due to not having a consistent technology baseline or framework in place.

- Never-the-less, the Panel's consensus was that feasible solutions were available within a cost and power consumption level that would be acceptable to the overall SKA1 project.
- Also the later scaling up to address SKA2 was also seen as feasible.

All the information presented gives a high level of confidence in the maturity of the concepts of the signal processing Element (as defined in the System Hierarchy).

- **The Panel judged the overall progress of this Element of the System to be in good shape and ready to move into the Definition Phase, with recommendations for the next phase.**
- The Panel observes that the Project will be moving from PrepSKA governance to governance under the Project Execution Plan (PEP). This plan envisages a more rigorous approach, in which work will be carried out by work-package contractors. It will important that the Signal Processing participants and the Project be capable of providing sufficient definition (documentation, requirements and other aspects) to effectively contract out subsequent work. This may require a more structured approach than is currently under way, although there is insufficient detail available to provide specific advice. In order to be successful, preparation will have to begin much earlier than the Signal Processing SRR. The panel recognises that there may be inconsistencies in timing and funding issues associated with the transition to a structured approach, but these aspects are beyond the scope of this report.
    - Preparation for the work-package contractor stage may encourage current participants to collaborate and concentrate their efforts, so as to corral the range of qualifications and technical coverage to be able to carry out work-packages in this area.

4 Technical Adequacy and Maturity to Move into the Next Phase

To assist in structuring its comments and feedback the Panel chose to consider the Signal Processing function at three different levels: The Concept level, the Architectural Level, and the Implementation Technology choice. This section contains two recommendations, and spawns a few more recommendations, which are described in the next section.

4.1 Concept Level

This level describes the basic functions and algorithms to be carried out on the received streams of data from the antennas. These are Array Beamforming, Channelisation, Correlation, Pulsar De-dispersion, Pulsar Timing and so on. It was clear to the Panel that at this level the mathematics and algorithms required for these functions have been developed over many years and are clearly well understood and mature, although in the case of pulsar processing, the efficiency of the search algorithms appears to require further investigation.

The required experience to move forwards is certainly available within the Project and this in itself gave the Panel a high level of confidence that the Concept Definition milestone has been achieved.

In addition to the basic concepts much of the information presented in the review was about specific architectural choices coupled with the selection of a specific implementation technology. The descriptions at the concept and the architecture levels could have been more clearly delineated. It would be useful to clearly indicate where concepts and architectures must be coupled, and where such coupling is not necessary.

4.2 Architectural Level

At this level the Panel noted different approaches. For example, memory based versus pipelined structures, and different ordering of beam-forming and correlation were presented. It was also noted that different architectural partitioning among other System Elements (e.g. Stations) could affect, for instance, the optimisation of power consumption.

In many cases architectural approaches being considered had chosen differing implementation technologies and thus it is difficult to form an objective assessment of the relative benefits of the different architectures (see recommendations 1 and 4).

Figure 1, which was presented, indicates feedback from the Element architecture level to the System architecture. The Panel was left unsure of how this was going to be accomplished. This would be easier if a clearer definition of each of the available architectures could be established.

Recommendation 1: The Project should clearly define a number of suitable architectures to be investigated further in the next phase. These will probably align with options to be carried forward into the next phase. Each of these architectures may have different impact at the system level, and these impacts should be carefully examined at both levels.

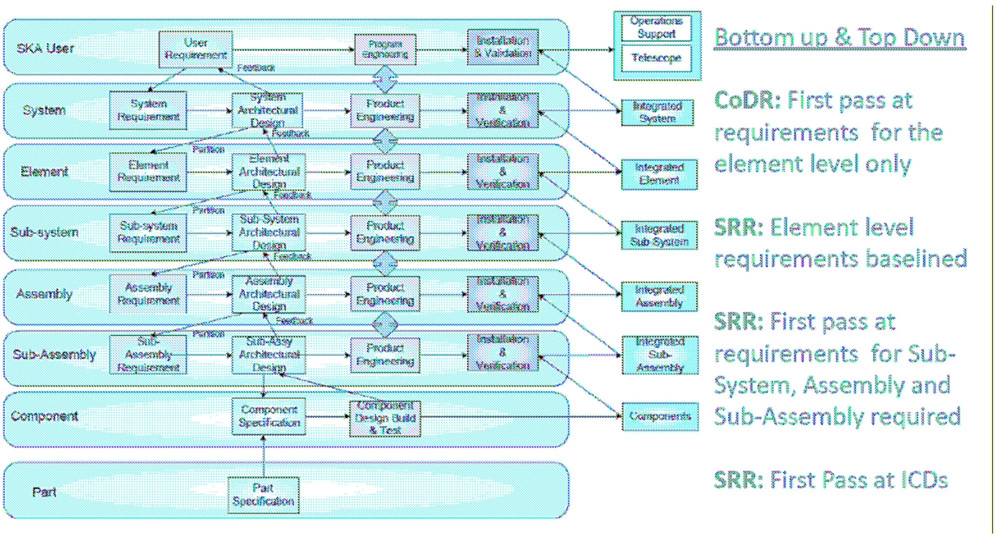

Figure 1

4.3 Implementation Technology

The implementation of the various architectures was achieved in technologies ranging from “pure” software targeting a general purpose compute engine, more specialised software running on a graphics processing unit, hardware implementation on programmable logic arrays, to full ASIC implementation. In addition to these degrees of implementation “hardness” there were also differences in terms of the basic underlying semiconductor process technology choices such as semiconductor feature sizes of 65nm versus 22nm.

These implementation technologies can be seen to have differing benefits (e.g. high flexibility for software versus high computational efficiency for hardware). Furthermore, the combination of architecture and technology can be expected to perform differently in respect of key factors such as reliability (perhaps related to the number of backplane joints, the number of fans or other failure prone items).

These disparities of implementation technology should be reduced as much as possible, and a technology baseline established. A specific recommendation is made in Section 5.3.

5 Measures Identified to Address Gaps or Shortcomings

This section describes in more detail recommendations, some related to the issues described above, and some additional ones.

The Panel felt that these steps are urgent for the Project in order to give the option providers the maximum time to optimise their architectures and also give the earliest possible notice of the basis for the sub-selection at the Signal Processing Requirements Review (SP-SRR) in Q3 2012.

5.1 Define a Fixed Element Boundary

The boundaries of the Signal Processing Element may already be well defined at the system level, but there seemed to be a number of variants among participants presenting at the review. For example, station beamforming does not appear to be included in Signal Processing, but this was treated in some presentations as though it were. The Panel recognises that this could lead to tricky situations when optimising at the system level. This a good reason for clarity among all the participants and the participants in adjacent Elements.

Recommendation 2: SPDO should define a fixed boundary for the Signal Processing functions again to ensure that comparisons are meaningful, and to ensure that the participants are aware of the boundaries. If ideas for optimisation are put forward that cross the boundaries, this should be taken up at the system level.

5.2 Optimise at the Overall System Level

The previous section discussed defining the boundary of the Signal Processing Element. The Panel also noted from some of the information presented, that there could be the possibility to improve some aspects of the signal processing performance by adjusting the requirements of closely related elements in the signal flow path to achieve an overall optimisation. An example of this is the trade-off between the number of antennas per sub-array, or station, where local beamforming is carried out, versus the number of stations in the overall System. This balance could be optimised, for example, power consumption whilst keeping to the overall SKA target specification.

While there could be an argument for defining element boundaries by technology, the Panel is not in general agreement with such an approach. The boundaries should make sense from the system perspective, regardless of technology. However, the foregoing paragraph indicates that optimisation can be made across element boundaries. Clearly these must be done at the system level.

But the requirements management approach being followed by the Project will eventually bring the Element level inputs together at the overall System level, but the Panel felt that this process would be too slow to allow the opportunities for optimisation to be identified and pursued.

Recommendation 3: The Panel recommends establishing a regularly scheduled (monthly) overall System requirements review between key Elements, where significant cross-Element optimisation could take place. This could take the form of a System Architecture Board (SAB) at the signal processing level, with a charter to be responsible for this overall optimisation. The SAB should examine or commission studies on the requirements trade-offs between various aspects of the overall SKA, where changes at that level would have a significant impact on the signal processing.

5.3 Define a Common Technology Footing

It is essential that all of the participants make the same assumptions for technology roadmaps. In particular, the use of Moore's Law should be precise. Appendix I contains a brief summary of the impact of Moore's Law projections in the relevant time frame.

Recommendation 4: Define a baseline implementation technology for the purposes of normalising performance results to allow a meaningful comparison to be made of the benefits of the alternative architectures. As an example, the Panel suggested adopting a currently available semiconductor technology generation in which CPUs, GPUs, FPGAs and ASICs are all available. Given that ASIC implementation would be one generation behind due to development delay, an approach could be to use 28nm for COTS FPGA & GPU and 40nm for ASICs. The semiconductor technology roadmap can then be used to scale all the potential solutions forwards in time on a uniform basis.

5.4 Clarify and Reduce the Number of Options

The Panel observed that the process by which the narrowing of options will take place is not currently defined. This was identified as an important gap.

As shown in the Figure 2, the Project plans to take forwards a number of potential solutions to the Signal Processing Sub-System Requirements Review (SP-SRR) milestone and make a decision at that point on a reduced set of Candidate Options. The date scheduled for the SRR is Q3 2012, which provides about 15 months for the further development and optimisation. Also, the System Requirements Review (SRR) will be held in early 2012. There is likely to be significant feedback to the system level from the Signal Processing Element level, which could influence the SRR. It will be important that this feedback be based on clearly defined options.

The Panel was concerned that carrying forward too many options could be wasteful of resources and hence counter productive. The Panel felt that it would be desirable to focus the number of options as early as possible and that this could perhaps be achieved through encouraging collaborations that would allow merging some of the options.

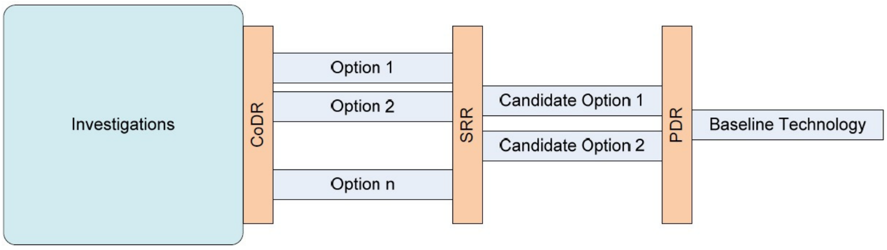

Figure 2

The means by which decisions on options will be made is unclear at this point. It is reasonable to expect such decisions to be made on the basis of performance requirements, as well as on cost.

Recommendation 5: There are about 200 requirements currently defined for the Signal Processing functions. The Panel felt that this was too many to handle for the selection process and recommended that the Project should define a small subset of the key requirements and evaluation criteria, perhaps 10 to 12 or so, against which the different architectural approaches can be compared. This subset should probably contain factors like cost (NRE, variable and running), power consumption, technical performance (S/N, pulsars/min.), flexibility, reliability, fault tolerance, and future proofness. The ability of an approach to exploit factors such as parallelism and locality should also be considered (see appendix I).

5.5 More focus on Pulsar Processing

The signal processing required to meet the DRM objectives for tests of General Relativity by discovery and precision timing of relativistic binary pulsars poses a significant challenge for the SKA. Although the techniques are well-developed, they have not yet been implemented for real-time processing and it is not clear if the DRM goals for pulsar science can be fully met in real time. Much depends on the rate at which the large searches must be carried out. The need to search in "acceleration space" for pulsars in relativistic binary orbits where the pulsar rotation period will be continuously Doppler shifted will be particularly challenging.

The Panel's view is that there was much less implementation experience evident for real-time pulsar processing and that this area will need stronger attention in the next phase to further refine requirements and bring it to the same level of maturity as the other areas.

Appendix II contains more details on this topic.

Recommendation 6: Once the pulsar search requirements, including operational requirements) are clearly understood, the signal processing group should undertake an investigation of pulsar processing efficiency, using various levels of technology available ("soft" to "hard"). Assuming that acceleration searching remains a requirement, the investigation should particularly emphasise acceleration searching.

6 Comments related to specific questions

The Panel was asked to evaluate the overall progress in the project against eight specific questions. These are listed below together with the Panel's observations.

1 Are the requirements complete, and sufficiently defined for this stage of the project?

Yes, provided momentum is continued in refining the requirements and defining verification procedures.

The Panel views the current requirements document, which contains around 350 overall requirements, as a first draft and the assumption is made that the system level requirements will continue to be refined and increase in number.

Where requirements have been initially less well defined "Memo 130" has been used as a baseline. However, it was stated that Memo 130 does not detail requirements for the non-imaging signal processing and as such the Panel has identified this as a gap.

The current document does not make a clear delineation between SKA1, SKA2 and extensibility requirements (see Recommendation 7).

A regularly scheduled process for rolling functional requirements changes back up to the system level and considering the overall system effect has been described and the Panel agrees that this is important. (See Section 5.2 for a complete discussion).

Also with the large number of requirements being tracked, adopting specific requirements management engineering tools may be necessary.

Recommendation 7: Requirements should be grouped so that there is a clear delineation SKA1, SKA2 and extensibility requirements.

Recommendation 8: The project should acquire professional requirements management tools and use them for tracking requirements in the future. The participating organisations may also require the same tools.

2 At the concept level, is the element/subsystem presented capable of meeting the requirements?

Based on the information presented at the review, the signal processing as currently envisaged for both the imaging and non-imaging processing does appear feasible. Requirements will have to be developed in much more detail in the next phase.

In the case of non-imaging processing, cost (capital & operating) may turn out to be an issue, if the system requirements are very ambitious. De-scoping at the system level (e.g. lengthening survey times) would be an obvious way to reduce cost, if necessary (see Recommendation 6).

Almost all of the concepts presented are likely to be capable of meeting system requirements. Cost and performance will be the deciding factors.

3 Have interfaces to other aspects of the system been adequately identified and defined at this stage of the program?

The External interfaces are described in the high-level description document and the location of internal interfaces is also described to first order in functional breakdown form.

The Panel noted that signal processing functions also exist in other system elements such as the Aperture Arrays and that it may be helpful to find a way of combining and/or cross referencing these to allow a more complete view of the processing functions to be taken to ensure optimum partitioning is achieved (see Recommendations 2 and 3).

The Panel felt that interface control documents are not needed at this stage, but should be started soon. This is already noted in the High Level Description.

4 Are the options proposed to be carried forward credible and are the presented data and information in support of each option credible?

The Panel observed that some of what was presented looked mainly like results from existing projects, so it was not really clear to the Panel which of the options were to be carried forward. However during the feedback it became clear that all the presented options were potential candidates and would be carried forward. This raises the concern of the efficiency of the process going forward and the Panel felt there should be some opportunity to combine approaches through collaboration. The SPDO should look for these opportunities and provide leadership (see discussion in Section 3 – "Progress" and Recommendations 5 and 9).

The Panel found that, whilst all the options appeared credible, judging them in detail was not really possible due to the lack of a firm basis of comparison. The options would be more credible if comparison was on an equal footing. As mentioned earlier in this report, the Panel recommends that the Project should immediately define a common footing and parameters to enable meaningful comparisons to be made (see Recommendation 4).

Concerning the level of depth of the analysis presented, the Panel felt that the JPL approach of comparing architectural options looks to be an appropriate level for this stage of the project, although not all the alternatives or approaches were explored in this approach.

Recommendation 9: The SPDO should lead a process of consolidation of technical effort, where options are sufficiently similar. This should be done in conjunction with Recommendation 5 and in the light of the longer term aspect of forming work-package contractors (see the last paragraph in Section 3).

5 Have all the necessary aspects of the specific element/subsystem been considered and addressed during the review or are there gaps and/or shortcomings?

Apart from points mentioned in the other answers the Panel observed the following gaps:

a) There was clearly less implementation experience with pulsar processing and the predicted performance of search algorithms would appear to limit the survey performance of the instrument. As this area is a key part of the Design Reference Mission for the SKA1 the Panel recommends that this area should be given an increased level of attention into the next phase (see Recommendation 6).

b) Linked to the above point, consideration should be given to the re-construction of a high bandwidth signal for pulsar detection and pulsar astrometry. Also note the later additional comments in this report on Pulsar Processing.

Recommendation 9: There needs to be an analysis of spectral channel widths, derived from science and system-level requirements. This should lead to optimization/coordination of the spectral channelization for beam-forming with the varied time resolution requirements (and hence the bandwidth requirements) for pulsar searching and timing (for a discussion, see Appendix II).

c) *The Panel had a further concern that a fully automatic, hands-off, fixed-hardware pulsar search facility may miss serendipitous pulsar discoveries.*

d) *RFI Mitigation processing was touched upon in only one presentation and needs to become more detailed as soon as possible. Particular attention should be paid to the potential for self-induced RFI, to which the beam-formed sum will be quite sensitive.*

Recommendation 10: The signal processing group in conjunction with the system group should carry out an analysis of RFI mitigation techniques at the signal processing level, concurrently with a full inventory of RFI mitigation requirements at the system level.

e) *Due to the way in which the system elements of SKA have been partitioned there does not appear to be a single block diagram or description of requirements showing the complete signal path from ADC to the input of the image processing. The Panel is concerned this could lead to sub-optimisation of the overall system.*

Recommendation 11: The signal processing group should prepare a signal path diagram beginning at the ADCs, which emphasises a signal-processing perspective in the system.

6 Does the risk profile appear reasonably detailed and assessed for this stage of the program?

The Panel felt that the risk profile is acceptable for this stage of maturity of the Project.

The Panel welcomes the development of the technology roadmap, which, of course, should be agreed across the project and used uniformly (see Recommendation 4).

However, from the experience of the Panel members, Moore's Law is now progressing more slowly than represented. The risk of an over reliance on Moore's Law as "saviour" could perhaps be mitigated by rigidly defining a technology generation, based on the project's timetable, which all candidate solutions are required to use for performance benchmarking. Also see the comments in Appendix I.

Further the Panel were concerned that very advanced semiconductor technology is likely to be less reliable (e.g. electromigration at 20 nm scale) and as mentioned earlier reliability should be a parameter against which the alternative solutions are compared (see Section 4.3).

7 Do the stated risk controls and proposed mitigations appear reasonable and executable?

Mitigations in the risk document seem reasonable enough at this stage.

The process by which risks get retired was not completely demonstrated, but the Panel assumes that this will be dealt with in the future.

Recommendation 12: The Panel agrees with the last slide in the presentation on risks, regarding actions to be taken, namely:

- *Ensure that the risks currently listed are owned, managed and mitigated.*
- *Ensure that new risks are identified and captured.*
- *Continually track and monitor progress on risks.*
- *Review risks at the sub-system level.*
- *Roll up risks and inform the system risk register.*

8 Is the overall plan (including the identification of the tasks, effort, resources, costs, schedule and risk mitigation needed) to complete the subsequent project phases credible?

The Panel is impressed by the amount of work and the thoroughness of the investigations. If the process for comparison of options can be specified quickly, then the forward plan looks in good shape.

The general plan for the next year is less clear than the plan for the pre-construction phase, when work-package contractors will be identified, although preparation of the contributing groups for the work-package contract stage is unclear (see comments in Section 3).

The Panel is not clear on what is actually being carried forward to the next stage (see Recommendation 1).

Recommendation 13: The project should ensure that the output documentation for the next phase results in a work breakdown structure that can be used subsequently in the pre-construction phase.

Appendices

Appendix I: Specific Advice on the Use of Moore's Law

The Panel's opinion is that Moore's Law is slowing down. The experts on the Panel made the following observations:-

- each process node (a technology generation of a specific line width) change occurs every 2 years, up to 3 years.
- for each new node, expect 1.5X improved speed-power performance and 2X improvement in density
- hard data:
    - 40nm-28nm 1.55X speed-power improvement, 1.95X density improvement ;
    - 28nm-20nm 1.5x speed-power improvement, 1.9 density improvement
- most popular nodes will have best IP support: 0.13um, 65nm, 28nm
- leakage will be getting worse, but not dominant for this application

Given the slowing of Moore's Law most future performance gains will likely come from the exploitation of both parallelism (at different granularities, and whether harnessed using custom hardware concurrency in ASICs or FPGAs, or with programmable cores in CPUs or GPUs) and locality (temporal and spatial). The requirements for the compute backends might thus be expected to include information about the potential available parallelism and locality, either inherent to the various computation sub problems, or inherent to various choices for algorithms for solving those computational problems. There are many well-known approaches to performing such parallelism and locality characterizations (so-called "limit studies") which could be applied.

Appendix II: Pulsar Processing

The Panel notes that the scientific objectives of General Relativity tests by the precision timing of pulsars have been deeply considered. The accurate timing of known pulsars is a straightforward application of current techniques. In cases in which more than one pulsar appears within the beam of a dish the timing can be sped up by forming several array beams pointing at specific pulsars, subsequent coherent de-dispersion, period folding, time-tagging and archiving.

To extend the array of pulsars used for timing and to study other gravitational effects, new pulsars are to be searched for and detected in real time. The Panel acknowledges that the huge numbers of new isolated pulsar candidates can be sorted and false positive candidates rejected by neural network algorithms running on general purpose computers. However, the detection of the most interesting pulsars for tests of General Relativity, those in tight binary orbits around massive compact objects, imposes an enormous and complex computational challenge for the SKA, particularly if it is to be done in real time. The computational cost in facilities and power for real-time detection and analysis should be evaluated to determine if it would be more practical to record the data streams from each telescope beam for off-line processing. The Panel suggests cost analysis of extending the bandwidth and storage of the transient detection and recording system to archive continuously sampled data for off-line pulsar acceleration searches as part of SKA1. This will slow the pulsar survey search speed, perhaps considerably, but the speed of the search may not be critical. On the contrary, we suppose that completeness of a search in parameter space (position, dispersion measure, period, period derivative, orbital speed, etc.) to be of paramount importance for satisfying the DRM. A subset of a continuous recording capability also enables detailed off-line follow-up studies of both newly discovered pulsars and well-known pulsars for their emission and polarization properties.

The Panel notes that there is no mention of pulsar-synchronous operation of the imaging correlator/spectrometer in the documentation provided. This capability is important for at least two classes of science: pulsar astrometry and Galactic HI studies. For newly discovered pulsars an accurate timing model and pulsar ephemeris can be determined much more quickly if an interferometrically determined position is fed into the model. Increasing the signal to noise ratio by pulsar synchronous imaging allows faster position refinement. Differencing on- and off-pulse HI spectra plus a Galactic rotation model enables the mapping of HI between the Sun and the pulsar.

There needs to be optimization/coordination of the spectral channelization for beam-forming with the varied time resolution requirements (and hence the bandwidth requirements) for pulsar searching and timing. This may require reconstruction of wider bandwidths by combining beam-formed spectral channels appropriately before the de-dispersion operations. The reconstructed bandwidths are likely to be different for searching and timing, e.g. in searching for binary millisecond pulsars time resolution of 100 microseconds may be sufficient for detection, but for precision timing, resolution as short as 0.2 microseconds is specified (Document 1a, WP2-040.030.010-TD-001, High Level SKA Signal Processing Description, p26), thus implying 500-1000 times wider reconstructed bandwidths for timing as compared with the channelization used in searching. The bandwidth reconstruction process is straightforward, requiring an FIR filter and inverse FFT, but it may affect the upstream sampling rates and pre-beamforming channelization polyphase filter characteristics.

Appendix III: Panel Membership

Robin Sharpe (chair)
External advisor
Ex Philips Semiconductors
Winchester, UK

Peter Dewdney
SPDO
University of Manchester
Manchester, UK

Andre Jongeling
Jet Propulsion Laboratory
Pasadena
CA USA

Phillip Stanley-Marbell
Research Staff Member
IBM Research
Zurich

Louis J. Morales
iSine Inc.,
Boston
MA USA

Tim Hankins
New Mexico Tech,
NRAO
NM USA

Appendix IV: Panel Charter

The CoDR will be conducted to evaluate:

- The overall progress,
- Whether the technical adequacy obtained during the concept phase is at a sufficient level of maturity to allow the Signal Processing Element to move into the next phase,
- Whether all Signal Processing Element aspects of the project have been covered and where gaps exist, whether adequate measures have been identified to address the shortcomings.

The expected outcome of the review is the establishment of the signal processing concept baseline by conclusion of the signal processing level concept phase. Following the successful conclusion of the review the next phase, the Signal Processing Element definition phase, will be initiated.

More specifically the Review Panel is requested to consider the following questions:

1. Are the requirements complete, and sufficiently defined for this stage of the project?
2. At the concept level, is the element/subsystem presented capable of meeting the requirements?
3. Have interfaces to other aspects of the system been adequately identified and defined at this stage of the program?
4. Are the options proposed to be carried forward credible and are the presented data and information in support of each option credible?
5. Have all the necessary aspects of the specific element/subsystem been considered and addressed during the review or are there gaps and/or shortcomings?
6. Does the risk profile appear reasonably detailed and assessed for this stage of the program?
7. Do the stated risk controls and proposed mitigations appear reasonable and executable?
8. Is the overall plan (including the identification of the tasks, effort, resources, costs, schedule and risk mitigation needed) to complete the subsequent project phases credible?

These questions are to be within the context of SKA1 but with consideration of extensibility to SKA2.

Appendix V: Review Documentation

The CoDR Documentation Package consists of the following documents:

Doc No	Title
1a	Signal Processing High Level Description
1b	Searching For Fast Transients with SKA PHASE 1*
1c	Pulsar survey with SKA phase 1
2	Technology Roadmap
3a1	Software Correlator Concept Description 001
3a2	Software Correlator Concept Description 002
3b	GSA Correlator Concept Description
3c1	ASKAP Correlator Concept Description SKA2
3c2	ASKAP Correlator Concept Description SKA1
3d	A UNIBOARD Based PHASE 1 SKA Correlator and Beamformer Concept Description
3e	CASPER Correlator Concept Description
3f1	SKA1 ASIC-Based Correlator for Minimum Power Consumption—Concept Description
3f2	SKA2 ASIC-Based Correlator for Minimum Power Consumption—Concept Description
3g	AA CORRELATOR SYSTEM CONCEPT DESCRIPTION
3h	Central Beamformer Concept Description
3i	Station Beamformer Concept Description
3j1	GPU Processing for Real Time Isolated Radio Pulse Detection
3j2	GPU Processing for Pulsar Search
3k	PAF Beamformer Concept Description
3l	FPGA DeDispersion Concept
3m	FPGA based pulsar searching
3n	Benchmarking pulsar search backends
7	SKA Signal Processing Requirements
8	SKA Signal Processing Costs
9	SKA Signal Processing Risk Register
10	Signal Processing Strategy to Proceed to the Next Phase
32	Software & Firmware Strategy

Additional reference documents were available through the Project's website at:-

http://www2.skatelescope.org/public/