Rev K

Report of the Review Panel
SKA Signal Transport and Networks
Concept Design Review
June 28-30, 2011

Table of Contents

1 Summary

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

- The preparation for the review produced documentation which adhered to a high standard and was distributed before the meeting. The presentations given during the review itself were well received, although, in a few cases, the information presented was not provided in the documentation.
- The panel believes that the present scope of the STaN Domain as defined needs to be reviewed at the system level.
- All of the solutions presented appear feasible. But estimates of cost and power consumption were very preliminary.
- The project needs to define SKA1. There was confusion about including PAFs for the definition phase of the project. If the PAF requirements for data transport are included, then they will drive the design. The panel recommends that the PAFs be included in the SKA1 design.
- Data transport and timing distribution solutions should be grouped by distance. The project should put forward a set of options for each distance rather than by technology or data type.
- The requirements for RF over Fiber appear to be less developed than alternatives and there was less implementation experience. This area will need to be focused on during the next phase to bring it to the same level as the others.
- Taking into account the preceding comments, the panel believes the Signal Transport and Networks element is ready to move into the Definition Phase.

2 Introduction

The SKA Signal Transport and Networks (STaN) Design Review (CoDR) was held on June 28th and 29th 2011, at Jodrell Bank Observatory, UK. The STaN presented the progress to date to determine if STaN had met the Project's CoDR milestones. The CoDR panel consists of four members external to the project with experience in timing standards, signal transport and networks and fiber optic systems. In addition the panel was also joined by the project engineer from the SPDO. See Appendix 1 for panel membership and appendix 2 for the panel charter. In a verbal report (supported by slides), the panel's initial observations were fed back to the SKA Signal Transport and Networks element, representatives from the SPDO, the presenters and the observers on June 30th. This report outlines and further details the observations and recommendations made by the panel and assesses whether the STaN element has achieved sufficient maturity to move ahead into the next phase.

The panel was in full agreement in its assessment of the STaN and our comments are listed in this report.

3 Preparatory Documentation

The preparatory documentation, consisting of 18 documents totaling over 650 pages, was distributed to the review panel a few days before the meeting. The documents are of high quality and technically detailed. The review panel acknowledges the level of effort required to prepare the documents and thanks the staff in the SPDO and the contributing organizations for this standard of preparation.

4 Scope of the Signal Transport and Networks System Element

The STaN “domain” or system element, as currently defined by the SPDO, contains several disparate parts. It appears that the principal connection between them is more technology-based than based on a logical analysis of the system. It almost appears that “anything that uses fibre” is part of the STaN element.

For example, RF-over-fibre technology is potentially useful as part of an AA-low station, or part of the RF chain in the Dish Array, two completely different parts of the system. Also, the delivery of coherent timing signals to the receptors (either for Local Oscillators or for ADC clocks) is a technology that is used to achieve required signal coherence across the array – it could be a system element on its own. The fact that fibre runs are probably part of the synchronization solution is not the main driver in parsing the system into elements. Equally the provision of accurate time for comparing pulsars with current time standards is a completely different aspect of the system. The Panel feels that this should be reviewed at the system level, so as to ensure clear boundaries to the STaN element or its logical successor(s).

The Panel also feels that it is essential that the SPDO (and later the SPO) retain and enhance its expertise in every technology area now considered part of the STaN domain. If there are changes made at the system level, it will probably be necessary to increase access to independent technology experts, to provide sufficient coverage. These experts will have to operate across the system wherever the relevant technologies are utilized. It is clear from this review that the work of contributing organizations cannot be properly evaluated, reviewed or tracked unless the SPDO retains people with the experience and expertise required to do so in each area. This is especially important since there are often competing claims or technologies vying for a role in the system or that must be continuously evaluated as they are carried forward as options. In some areas, access to the relevant experts may not require a full time commitment. The important thing is continuous coverage by persons who are knowledgeable of how the system is developing and who can provide independent advice on the technology.

Recommendation 1: The project should review the STaN element at the system level, to ensure that a logical hierarchy is created. This could result in the assignment of aspects now considered STaN to other elements.

Recommendation 2: The project should review its level of support for STaN technologies to ensure that there is sufficient coverage of the technologies utilized in the SKA system among SPDO/SPO staff, or accessible to the SPDO/SPO sufficiently regularly to provide continous coverage. This will become even more important as the SPO develops.

5 Overall Progress

It is clear to the Panel that the Signal Transport and Networks concepts are adequate to meet system requirements as now described, are sufficiently mature and in most cases use well understood technology.

Recommendation 3: The panel judges the overall progress of the STaN element of the system to be sufficient and ready to move into the Definition Phase.

Many of the presentations described specific implementations, although no prototypes had been tested. It will be important in the next phase to evaluate which of these implementations require physical prototyping, so that this effort can begin.

The panel observed that the Project will be moving from PrepSKA to governance under the Project Execution Plan (PEP) (SPDO to SPO). This plan requires a more rigorous approach in which work will be carried out by work-package contractors. It is important that the STaN participants and the project be capable of providing sufficient definition (documentation, requirements and other specifications) to effectively contract out subsequent work. This may require a more structured approach than is currently under way, but there are insufficient details available to provide specific advice in this area.

5.1 STaN Requirements

Development of STaN requirements is still a work in progress. It is important to maintain momentum in refining the STaN requirements to reflect the system level requirements. The panel recognizes that the system level requirements will continue to be refined and increase in number as the project continues. It is the responsibility of the STaN to follow these developments and incorporate the relevant specifications.

The panel views the provided STaN documents, which make reference to the overall system level requirements, as a first draft and the assumption is made that these documents will be brought up to date as the StaN moves on to the next stage of the project. The STaN should develop a method of referencing the system requirements while specifying the domain design specifications.

The current documents do not make a clear delineation between SKA1, SKA2 and extensibility requirements. The STaN should emphasize the issues associated with extending the SKA1 system to SKA2 in the STaN area, particularly including PAFs in the design.

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

Recommendation 4: The project should emphasize further development of requirements and the flow-down of system requirements to the STaN area. The organization of requirements may change by implementing Recommendation 1:. The requirements documents should clearly separate requirements for SKA1, and those for extensibility to SKA2.

5.2 Clarification and Reduction of the Number of Options

As shown in Figure 1, the STaN plans to take forward a number of potential solutions to the Sub-System Requirements Review Phase. The completion milestone of the Sub-System Requirements is a review, (STaN-SRR). It is believed that the decision to reduce the set of Candidate Options will be presented at that review.

The Panel observed that some of the presented material looked like results from existing projects, so it was not clear which of the options were to be carried forward. During the feedback, however, it seemed that all the presented options were potential candidates and are likely to be carried forward.

The STaN domain should aggregate similar options and then obtain clear statements of the cost, risk and feasibility options and put together plans to pursue these options. Parameterized models should be assembled to allow a rapid response to project wide decisions such as site or array configuration selections.

The Panel is concerned that carrying forward too many options could be counterproductive. There are also concerns about the efficiency of the collaborative process. The Panel feels that it would be desirable to reduce the number of options as early as possible and that this could be achieved through encouraging collaborations that will allow the merging of some options. The SPDO should provide leadership in this area.

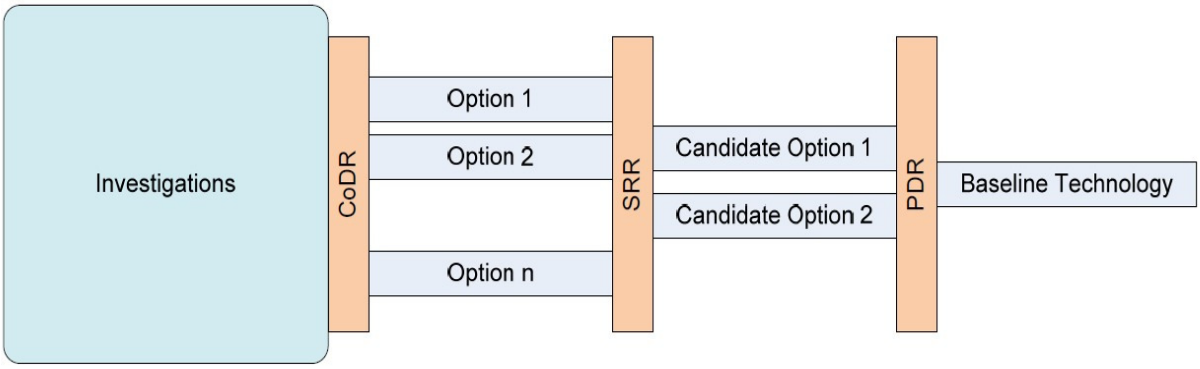

Figure 1

Recommendation 5: The project should carry out a review of the various options being considered for carrying forward, reduce the number if possible, and merge the effort contributed to very similar options into more cohesive single efforts. This should be done in parallel with the steps outlined in Recommendation 1:.

5.3 The Schedule

Inconsistencies in the order of major reviews in the schedule seem to be present (see the presentation on Network Infrastructure Concept Description), especially between the STaN area and the site engineering. For example the CDR for site engineering precedes the PDR for STaN. These two things are strongly interacting. It seems incongruous to review detailed site plans, while work is still being done on a detailed design of the networks. These should be more coordinated.

Recommendation 6: The project should review the system schedule, particularly with regard to the site engineering and STaN areas.

6 Digital Data Back Haul

The panel chose to consider the STaN data transport function (DDBH) in three different distance regimes: the central array signal transport and networks, the intermediate, and the long haul signal transport and networks.

6.1 Center Array Signal Transport and Networks

Different approaches have been presented by the teams to address the capacity and transport requirements associated with the different type of elements to be deployed in the field.
Considering the phases of the project, the panel recommends that further studies be conducted to

evaluate the effects of the distance to be covered versus the required data capacity to be transported. Based on the proposed solutions, it seems there could be opportunities to install more optical fiber in the central array thus avoiding signal amplification and regeneration as much as possible. In addition to this, the impact of additional fiber in the central array should be taken into account during the cost evaluation of the solutions used to transport the signals from the dishes and the other antennas in the central array.

Recommendation 7: For the shortest and possibly mid-range distance regimes, the project should undertake a study of various solutions as a function of different distance regimes and amount of data to be transported. It appears that the most cost effective design approach would emphasize "fibre in the ground" over very high bandwidth with expensive modulator components. In other words, it is probably not worth while, in a new system, to pack too much information into each fibre.

An attempt to leverage on commercial solutions to format and transport the signal, the use of DWDM, 10GE wavelengths or 100GE wavelengths should be evaluated versus the distance which has to be covered between the pedestals and the correlator. The fact that all the dishes and the antennas will not be installed all on the same day, should be factored into the analysis to determine which is the most efficient way to handle the complexity in terms of connectivity and number of optical paths.

A solution which involves the adoption of a fully fledged commercial solution (leased bandwidth) is not likely to be optimum for the short and possibly mid-range distances (essentially any part of the DDBH network that is likely to be wholly owned and operated by the project). This is because a full commercial solution will require interfaces, switches and other equipment that may never be used in a single direction high data-rate system. For this situation, developing specific hardware using COTS components (lasers, modulators, etc) is more likely to be an optimum solution.

Recommendation 8: The project should develop architectures that incorporate custom designs utilizing COTS components, so that such architectures can be contrasted with fully commercial solution.

6.2 Mid and Long Haul Signal Transport and Networks

At this level the panel noted different approaches.

It has been appreciated that all the different studies which have been presented recommended to leverage on standard interfaces and, as much as possible, on transport capacity available on "public" networks.

In this case as for the Central Array one, the panel recommendation is to evaluate the effect of the distance versus the required capacity to be transported from the pedestals to the correlator. In this case it will also be key to estimate the required capacity over time to be able to leverage functionalities and wavelengths bit rates expected to be available by the time the Mid and Long Haul Network will be required.

In the timeframes under consideration it is expected that 100 GE technologies will be mature so STaN should not be considering 40 GE technologies as the highest speed option.

Recommendation 9: Long haul and possibly some mid-range DDBH will have to use public networks. The project should keep track of the technology roadmaps for this area, so that the best choices can be made when they are actually required. Outside experts, familiar with industry trends, may be needed to assist with this work.

6.3 Impact of Phased-Array Feeds

There was some confusion about whether PAFs can be included in SKA1. This section of the report assumes that PAFs may be included in SKA2, but not in SKA1.

The late introduction of phased-array feeds (PAFs) on dishes will likely have a huge impact on the DDBH network. Even though PAFs are not expected to be sufficiently mature for SKA1, it may be necessary to assume that they will be used in SKA2, and to design the SKA1 DDBH network to be extensible to a large increase in data traffic. Otherwise, it may be very difficult to retrofit the system with sufficient fibre, especially in the densely packed central area of the array.

Recommendation 10: As part of the extensibility aspect, the project should carry out a DDBH design that anticipates PAFs in the system, and to look for innovative solutions to minimize the total cost in scenarios in which they are used or not used in SKA2. Possibilities include the installation of extra fibre, conduits instead of direct burial, larger termination huts, more wavelengths per fibre, etc.

7 Synchronization & Timing

These comments are based on WP2-030.070.000-TD-001 and the associated presentation.

The requirements for time and frequency transfer within the SKA array are undemanding (by top-level time and frequency metrology standards) for the first phase (SKA1) but may be more significant for the longer-distance links in the second phase (SKA2). This is because the maximum observing frequency of 3 GHz (SKA1) is to be increased to 10 GHz (SKA2) and does not require phase stability beyond the few picoseconds level at a few tens of seconds.

Recommendation 11: The panel recommends that the STaN Domain pay close attention to developments in time transfer techniques in the time and frequency community, as this is a rapidly evolving field.

In the SKA, the local oscillator (LO) function depends on a choice to be made in the antenna signal processing. It may be used for frequency down-conversion but, more likely, it will serve as the reference for the signal ADC. The concept design requires quartz crystal LOs with short-term fractional frequency stability of up to 3×10^{-13} at 10 s, for compatibility with the eventual upgrade to 10 GHz envisaged for SKA2, which are synchronized to the system master oscillator over unstabilized links.

Recommendation 12: The panel recommends that more attention be paid to alternative techniques emerging from the time & frequency community such as distribution of the master oscillator over stabilized fiber links. There will likely be a cost and complexity trade-off between the removal of the need for many quartz crystal oscillators and the requirement for active stabilization for the fiber links.

7.1 Maser clock/central timing facility

The master clock or SKA timing facility will need to have more attention paid to it. One of the key scientific objectives of the SKA, starting with SKA1, is the detection of gravitational waves by the monitoring of an ensemble of pulsars. This requires a long-term, stable and accurate timing reference. We encourage the STaN domain to look beyond the idea of having multiple masers at the central timing facility simply for redundancy, and look towards exploiting this to generate a more stable synthetic timescale – effectively to generate an SKA system timescale which is UTC(SKA). If the science question is "are pulsars more stable than the best atomic clocks?" then SKA must be able to compare pulsars to the best atomic clocks.

Recommendation 13: The project should re-confirm a science requirement to compare pulsar timing with the best available long-term timing references, and when such a facility should play a role in the system design. If the science requirement is confirmed, the project should plan accordingly.

7.2 Concept design

The concept design (WP2-030.070.000-TD-001) is predicated on LOs providing the phase-stable reference for each antenna or group of antennae. The stability requirements are such that these LOs need short-term fractional frequency stability of up to 3×10^{-13} at 10 s (for compatibility with upgrading of dishes to 10 GHz in SKA2).

It is currently proposed that within the central 1 km diameter core timing and frequency reference distribution be done one-way without any need to measure or compensate for delay fluctuations.

For the outer antennae in SKA1, up to 100 km, it is proposed to adopt the solution currently being implemented for e-MERLIN. e-MERLIN uses out-and-back measurement of delay fluctuations and lacks active compensation.

For the outermost antennae in SKA2, which may be up to 3000 km from the core, independent clocks (probably passive or active hydrogen masers) are proposed.

7.3 Comment on the concept design

The central cores, which in SKA2 will become very densely populated (SPFs in particular), can be treated on a "building" model with rather short path lengths. (It is not evident that the master oscillator/clock will be located at the center of the core region and indeed the cores of the SPFs and AA-lo arrays are not co-located, so consideration needs to be given to the distance from the location of the master clock.)

Application of the e-MERLIN model for the intermediate and outer SKA1 zones would certainly work, but it is likely to be expensive in LOs given the large number of antennae involved. Frequency distribution over stabilized links may in fact be less prohibitive in cost. The out-and-back delay measurement system proposed already generates the error signal for stabilization of the link.

The remote antennae in SKA2 will have digital data transmission to the correlator over fiber but these links are likely to be in the form of leased capacity from commercial or academic networks. Leased capacity would probably rule out use of these links for frequency/time synchronization, so high quality LOs will be needed at these remote locations. These could be masers, to provide long-term stability, or high quality quartz oscillators steered by GPS or other global navigation satellite system.

7.4 Detailed Comments by Paragraph

This section identifies three potential methods for distribution of synchronization and timing signals: one-way, two-way with path compensation and independent clocks. These methods encourage the distance grouping which we have encouraged elsewhere in the STaN Domain. Unlike DDBH, LO synchronization is not blind to the antenna type because of the different detection frequency ranges of the SPF and AA-lo antennae.

7.4.1 Functional requirements

The distinction between the coherence requirements within the array and the absolute timing requirements is not always clear. It would help to refer consistently to the former in units of phase and the latter in units of time.

The SKA has four distinct time and frequency requirements:

1. Coherence between array elements (phase stability across the array)
2. Telescope pointing (relation of system time to UT1)
3. Data time stamping (relation of system time to UTC)
4. Pulsar timing (relation of system frequency reference to the SI second)

The telescope pointing (0.6 s) and data stamping requirements (45 µs) are undemanding. The requirement for accurate absolute timing only applies to the master clock, meaning that absolute time does not need to be distributed over the array. Only the phase coherence condition needs to be met. The Master Clock and the relationship of SKA system time/frequency to external timescales does not receive much attention in this document but is important for the pulsar timing experiments.

Note 4 to Table 1: carrier phase GPS (or GNSS) methods will definitely be required to achieve 1 ns time accuracy.

In the SKA the local oscillator (LO) function depends on choices made in the antenna signal processing. It may be used for frequency down-conversion but it also serves as the reference for the signal ADC. The SKA is designed to process radio signals over the entire observing bandwidth up to 3 GHz (10 GHz in SKA2).

7.4.2 Optical methods

This section does not adequately distinguish between techniques being developed in the time and frequency metrology community. The work mentioned is all frequency transfer and has the goal of remote comparison of primary frequency standards and new-generation optical atomic frequency standards, at a level of frequency stability commensurate with the stability of these standards.

The three following techniques should be more distinguished.

RF/Microwave Frequency Transfer - Similar to the French work mentioned, this is done by amplitude modulation of the optical carrier. Functionality has now been extended to 10 GHz with improved transfer stability [Lopez et al., Appl. Phys. B 98 (2010) 723].

Optical Frequency Transfer – the optical carrier itself is the transmitted reference frequency. This yields the best stability and is expected to be the preferred method for frequency comparison of optical atomic clocks but does require optical frequency combs at each end of the link to convert from the transmitted optical frequency to the atomic clock frequency. [Optical or RF/Microwave Newbury et al., Opt. Lett. 32 (2007) 3056; Grosche et al., Opt. Lett. 34 (2009) 2270; Lopez et al., Opt. Exp. 18 (2010) 16849 provide a representative sample of the field]

Combined RF/Microwave and Optical Frequency Transfer - achieved by a transmission of mode-locked pulse trains (i.e. an optical frequency comb).

The optical carrier method is unlikely to be viable in the SKA context but either of the RF/Microwave techniques might be applicable.

The development of techniques for time transfer over fiber has largely followed on the back of these frequency transfer techniques.

7.4.3 GPS timing

Reference is made to 100ps precision achieved with GPS dual frequency carrier phase time transfer methods. It should be noted, however, that the accuracy is limited to around 1ns. These methods greatly reduce multipath effects, and the dominant limitation is receiver instabilities.

GPS timing is discussed more fully in the PrepSKA document “WP2: LO and Timing Requirements” (AppDoc20_LORequirementsvers1.pdf), Section 6, p. 10. Comments on this may be useful. The figure of 10ns for direct timing with GPS represents the best possible performance with a high quality, calibrated receiver, 20ns - 50ns is more typical. A dual frequency GPS-disciplined oscillator likely to achieve 5ns – 20ns relative to GPS time, but we are not aware of any commercially available devices. Single frequency GPS common-view can achieve 5ns – 10ns uncertainty on a time scale/clock difference over short baselines (< 100km). Dual frequency GPS common view can achieve a 1ns – 5ns uncertainty on clock differences (The BIPM uses 5 ns as the “typical” uncertainty.) This requires proper in-situ calibration of both receivers together with all cabling and distribution equipment. Note the measurement is of the difference between two clocks (or time scales). Even if one is a UTC(k) time scale, the UTC-UTC(k) offset introduces additional uncertainty. (UTC is a post-processed time scale published monthly for a particular date and is available from between two to six weeks after the event.)

7.4.4 Active Hydrogen Masers

NPL and USNO should be described as national measurement institutes rather than timekeeping services. If BIH refers to the Bureau International de l'Heure note that it was dissolved in 1987 when its functions were taken over by the BIPM. The BIPM does not operate any clocks or maintain a real-time physical time scale.

8 RF-over-Fiber

RF-over-Fibre (RoF) has potential applications for both the SKA Dish Array (SPFs and PAFs) and for the AA-low stations. They may also be applicable to AA-mid stations (dense AAs) as well, although such an option was not presented at the CoDR. In all cases, RoF is seen as having performance advantages over the use of copper cable to carry RF signals, and cost advantages over the placing of analogue-to-digital converters (ADCs) very close (ie, within a few cm to a meter) of the LNA-receiver combinations. The option of placing ADCs very close to the receivers is referred to as “immediate digitization” in the following discussion.

The option of using copper cables for carrying RF over more than a few meters is not thought to be viable in the Dish Array, either for performance or for cost reasons, and will not be discussed further here. For AA-low, copper cable is a viable alternative, especially if the copper can also be used for local power distribution.

For dishes, the immediate digitization option is likely to yield the best performance, if suitable packaging options can be found for ADCs. The packaging must provide a suitably controlled environment (especially temperature), and must shield the receptor from EMC produced by the ADC and associated digital equipment, including packetizing and optical modulation.

8.1 System Architecture Considerations

There is a strong interaction between the RoF option and the architecture of system infrastructure. The following describes this situation and a recommended approach.

The most efficient use of RoF is to aggregate RF signals from many LNA-receiver combinations in a local shelter or hub, at which the optical signals are received, digitized and packaged for distribution over the digital data backhaul (DDBH) system to the correlator. There are variations on this principle: a) for dishes with SPFs, the hub would be a small building servicing several nearby dishes; b) for dishes with SPFs and PAFs, the hub might be at the base of the dish (either in the dish tower or in a low building a few meters away); c) for AAs, the hub would service nearby array elements and could simply be more like boxes than buildings, and the signals from the hubs would go to larger building containing the beamformer equipment. All of these options have been presented in different ways in three documents covering RoF systems. Figure 4 of WP2-030.050010-SSDD-003 (CSIRO) discussing clearly displays sub-options for PAFs.

Thus the use of RoF will have large implications for the architecture of the system infrastructure that supports the collection of signals from dishes (ie, the routing of the DDBH fibres, the hub infrastructure, the number of fibres servicing dishes, etc.). There is less impact on system infrastructure for the AA-low stations. However, the number and distribution of hubs could be closely tied to the particular needs of a RoF signal collection system.

In another system architecture implication, RoF avoids the requirement to distribute timing signals to the clock input of the ADCs in the immediate digitization case. Note that if a heterodyne receiver design is adopted, requiring a Local Oscillator signal, then the cost of providing a clock signal to the ADC is greatly reduced.

The adoption of RoF solutions in the cases noted could have almost irreversible effects on the design of system and network infrastructure, once the design is at an advanced stage. Thus it is critical that the cost and performance of RoF solutions and the system implications of using these solutions be very well understood beforehand. If RoF solutions cannot be verified sufficiently early, then system and infrastructure design options will have to be carried forward until appropriate choices can be made.

On the other hand it may be possible to find system and infrastructure study architectures for which RoF or immediate digitization options can be easily carried forward.

In another twist, the construction of hubs (small buildings) may be required for the distribution of electrical power to dishes. (Hubs will require power themselves, as well.) On a much smaller scale, this may be convenient within AA-low stations. Therefore, the system architecture should take into account whether hubs are required for power distribution, and if so, whether they could be expanded for RoF aggregation and whether there is overlap between the ideal distribution of hubs for power purposes and for network purposes.

The operating costs of RoF; particularly power consumption, supporting more component types in the system configuration, and maintenance (reliability) should be considered. Power consumption may be an advantage for RoF over the immediate digitization option, although this may more likely for AA-low applications than for dishes. The two documents on RoF (WP2-030.050.010-TD-001 and WP2-030.050.010-TD-002) do cover these aspects at the concept level quite adequately. Significantly deeper analysis will be needed in the next phase.

Recommendation 14: A thorough assessment should be carried out to determine the impact on the system and network architecture by RoF in the aggregation of RF signals. This should be done separately for the Dish Array and the AA-low stations. The implications of sharing hub infrastructure with power distribution should be investigated at the same time, so that integrated architectures can be compared. If the performance of RoF systems for the SKA cannot be established sufficiently quickly (see Recommendation 15:), then the cost of carrying forward additional system design options should be considered in the overall cost balance. Thus the total cost of utilizing RoF vs immediate digitization should be considered. This is likely to be different for dishes and AA-low stations and should be analyzed separately.

This recommendation will most likely must be carried out at the system level, working closely with the STaN, AA-low and Dish Array element levels. The effort required to do so should not be expended until Recommendation 1: and Recommendation 5: have been carried out.

8.2 Performance

A RoF system is an analogue component in the RF chain. The required attributes are similar to those for gain stages, mixers and filters. A significant difference, however, is that the RoF component is physically and electrically long, and the fibre is likely to pass through a variety of uncontrolled environments, compared with a receiver system which is only a few 10’s of cm long and typically kept in a tightly controlled environment. Temperature changes and mechanical flexing (mainly for dishes, but potentially exposed cables in AAs could also be subjected to flexing) are the two environmental variables most likely to affect an RoF system.

The basic performance aspects of RoF are gain (loss), noise and linearity. These have been described thoroughly in the documentation. The analysis presented in the documentation and the review presentations should be checked more critically than was possible for the review panel. This is especially true for the noise and dynamic range analysis. The assumptions made in some cases seem questionable.

The best systems use expensive modulators, and a significant challenge is to better understand the potential for VCSEL laser/modulators, especially for AA-low, for which the cost of currently available alternatives is too high to consider further. The authors of WP2-030.050.010-TD-001 and WP2-030.050.010-TD-002 have put forward a plan to develop a technology roadmap. More expensive lasers and modulators may be suitable for dishes, but cost may still be an issue.

8.2.1 Performance Requirements

As noted already in Section 5.1, further development of requirements for the STaN element is needed.

Some of these requirements will affect potential RoF solutions. It will be important to cross-check the general requirements for the STaN element to ensure that sufficient information is available to understand specific RoF requirements.

8.2.2 Performance for the Dish Array

In addition to the usual performance parameters for analogue components, the most important aspects of RoF for the dish array are:

1. Complex gain variations and their rate of change: rapid gain and phase changes will require frequent re-calibration. This is not likely to be a problem as long as the temperature of critical components is controlled. As long as the time-constant of gain changes is reasonably long (many minutes to hours), the gains will be solved-for in a self-calibration loop.
2. Linearity, affecting spectral dynamic range. The standard methods of analyzing this, used in the communications field (3rd order intercept), are necessary, but potentially not sufficient. The wide RF bandwidths of radio astronomy systems are likely to accept multiple interference signals. Intermodulation products are likely to result. Or if out-of-band interference signals are strong enough (or band-stop filters insufficient), then intermodulation products may also occur.
3. Noise, affecting linearity. RoF systems are quite noisy (e.g. see WP2-030.050.010-TD-001, Table 7, in which the AOL component has a noise figure of 19 dB or 800 K). This should not be an issue in itself, if a sufficiently strong signal is injected into the input. However, such a strong signal can mean that there is not enough headroom to reach the dynamic range goal. This is likely to be less of a problem for AA-low than for dishes.
4. Bandpass shape, stability and smoothness. This could be the most subtle aspect of the requirements to design for. For example, bending of the fibre through the axes of dishes can cause weak multiple reflections, which when added to the main signal could produce small changes to the bandshape. These could be rapid enough to be difficult to calibrate, or small enough that they cannot be calibrated. These kinds of small changes are potentially difficult to detect in the lab, unless equipment is specifically designed for to detect them. Uncalibrated chromatic variations could also “leak” into the imaging domain, making it difficult to achieve the required imaging dynamic range.

Quite a lot of preliminary work has been done on the above effects by ASTRON, MPIfR, NRAO and MeerKAT. At the concept level, there is potential. For example, tests on bending have been carried out on the Effelsberg telescope, which indicates satisfactory results for the specific setup, and much better than copper RF cables. However, it is not clear that this test is sufficiently sensitive for SKA purposes; for example the bending radius in SKA dishes is likely to be much smaller than for the Effelsberg antenna. Also, MeerKAT and NRAO have field-tested RoF on KAT-7, and considerable practical advice has emanated from this work. Again it does not yield definitive guidelines for the SKA to use. A potentially useful test would be the following lab setup: a mechanical bender that repeatedly changes the radius of curvature of one RoF path, while cross-correlating with a RoF path in which no bending occurs. Synchronously dumping the correlator with the bending motion should yield a very sensitive test. In most radio astronomy labs, this could be set up in a few days.

PAFs present a stronger case for RoF solutions because of the large number of ADCs that would be required. If PAFs are used primarily for HI spectroscopy, the requirements on imaging dynamic range will be somewhat more relaxed. Because of the numbers of signal paths, cost, power consumption and reliability will be of greater concern than for single-pixel feeds.

Recommendation 15: In the next phase, a thorough evaluation of the best performing RoF systems should be tested in the lab using special test jigs, if necessary, to bring out the critical performance aspects, continuing work that has already been done to focus more on the SKA dish array requirements. If possible, comparative field tests in which RoF is compared side-by-side with immediate digitization should also be carried out. This should be done as soon as practical, and if the performance can be assured, then the number of system architecture options carried forward may be reduced (see Recommendation 14:).

Recommendation 16: Further investigation of cost, power consumption, and reliability is needed in the next phase, so many RoF options can be evaluated. This is potentially more important for PAFs than for SPF applications, especially since there are four PAF architecture options being carried forward at present (see WP2-030.050.010-TD-003).

8.2.3 Performance for AA-low

As outlined in WP2-030.050.010-TD-002 the potential for RoF for the AA-low stations is considerable. The requirements for AA-low are completely different from those for dishes. Performance of RoF is better than copper cable, especially if the complete bandwidth from 70-450 MHz is to be transmitted. The key issues are cost, reliability and power consumption. At present, only VCSEL devices are likely to be viable RoF solutions, because of the cost. (If similarly inexpensive lasers become available, then these should be investigated as well.) The numbers of devices required (~500,000 for SKA1) is sufficiently large that the development of specialized packaging could easily be justified. VCSELs appear to hold significant promise for AA-low stations, but considerable development is still needed, particularly in providing an integrated RoF solution that includes RF amplifiers with the optical components in a suitable package.

The main architectural difference for the AA-low stations between copper cables and RoF is that power distribution can be shared with signal transport for copper cables. This will have to be taken into account in comparing total station cost. Locally generated solar power is separately being investigated, but this option carries many additional risks at present.

Recommendation 17: Continue to develop the VCSEL-based options for AA-low stations so that realistic performance, cost, reliability and power-consumption figures can be made available. More expensive types of lasers should be abandoned, in favour of concentrating effort on VCSELs.

9 STaN Aspects of Monitor and Control

The STaN domain seems to include high-level aspects of the system that the Panel did not expect to be included (see presentation entitled SKA Monitor & Control Observations on Systems in production at ASTRON & NRAO, as well as a similarly entitled review document). This seems to be a system area currently in the STaN domain, but likely only belongs in the implementation of parts of the physical layer of Monitor and Control. As noted in Section 4, this should be clarified along with carefully defined interfaces.

Recommendation 18: The project should clarify the role of STaN in the system Monitor and Control area.

10 Comments related to specific questions

The panel was asked to evaluate the overall progress of the project using eight different criteria. These, and the panel's observations, are provided below. In cases where the subjects have been covered, cross-referencing has been used.

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

See Section 5.1.

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

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

Based on the information presented at the review, the STaN has many designs that will be developed in parallel in the Definition Stage. The requirements will have to be developed in much more detail for each of these designs to proceed.

A systems level requirement should be developed that specifies if time transfer between the SKA master clock ensemble and UTC labs is required over the long term. Vigilance is necessary to track current developments in time transfer technology to take advantage of recent innovations (see Recommendation 13:).

10.3 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. The location of internal interfaces is described to first order in functional breakdown form.

The panel recommends the M&C requirements at the System level (and hence interfaces) be defined early in the process (see Section 9 and Recommendation 18:).

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

See Section 5.2.

10.5 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:

- The STaN should track the RoF technology to ensure that if it is specified, it actually is major cost saver (see Section 8 and Recommendation 14:).

Although the Panel believes the following miscellaneous items have been addressed by the STaN, they were not presented in the documentation:

- Voice communication – VoIP should be considered as an inexpensive means of communicating.
- Video – Webcam support should be considered for the entire array as a means of enhancing security and maintenance activities.

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

The panel felt that the risk profile was too general for this stage of maturity. A few of the risks implicit and explicit in the various documents were not reflected in the Risk Register.

Recommendation 19: The project should improve and track the risk register for STaN. The top 10 risks should be highlighted so the STaN domain can track and act upon them.

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

All risk assessments in the document are vague. The STaN should recalibrate the probability and cost associated with each risk to introduce contrast. However, mitigation steps outlined in the risk document seem reasonable for this stage.

The process by which risks get retired was not completely demonstrated. The Panel assumes that this will be dealt with in the future (see Recommendation 19:).

10.8 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 in the STaN area. If the process for comparison of options could be specified, then the forward plan looks to be acceptable.

The STaN has a good strategy to move to the next phase. The Panel recommends that concepts put forward solely in presentations be documented before work starts on the next phase. A formal work breakdown structure and work package descriptions will be required soon, certainly well before the subcontracting phase. The StaN should ensure that documentation developed subsequent to this review results in a work breakdown structure.

Appendix I: Panel Membership

Steve Durand - Electronics Division Head, NRAO (Chair); sdurand@aoc.nrao.edu
Peter Dewdney - SKA Project Engineer, SPDO; dewdney@skatelescope.org
Stephen Lea - Senior Research Scientist, NPL; stephen.lea@npl.co.uk
Lorenzo Ghioni - Engineering Product manager, CISCO; lghioni@cisco.com
Tim Robinson - Operations & Technical Director, Net North West; tim.robinson@manchester.ac.uk

Appendix II: Panel Charge

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 Data Transport 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 a Data Transport concept baseline by conclusion of the Signal Transport and Networks Element concept phase. Following the successful conclusion of the review of the next phase, the Signal Transport and Networks 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.