Expert Panel on the Provision of Data Transport

Assessment of the Response to the RFI from South Africa

November 4th, 2011

I. From Receptors to the Data Processor

I.1. Feasibility

Is the proposed solution a logically possible proposition?

Yes

The basic feasibility question is: can an optical network of this complexity be built in a harsh environment?

Detailed proposals exist in the response for both the core area and for the remote receptors.

The proposal provides sufficient detail on how the network will be constructed that there is a good probability that feasibility has been demonstrated. Credible construction techniques, equipment, and fiber splicing and maintenance techniques that appear to be appropriate to the environment have been described.

In the innermost 30 km they propose co-locating the fiber and electrical feeds in common trenches and this seems like an obvious way to save cost. In urban areas in the U.S. this is common even when the electrical feeds are high voltage (e.g. IEEE-1590 Recommended Practice for the Electrical Protection of Communication Facilities Serving Electric Supply Locations Using Optical Fiber Systems).

Beyond 30 km the electrical feeds will use aerial distribution, but the proposal is to keep the fiber in trenches for stability reasons. This is a good idea generally, and probably essential for the timing system fiber as wind influence on aerial fibers can cause 10s of ps propagation time changes over a few seconds as the fiber is moved around by the wind.

The proposed splice cabinet solution appears to be well thought out, and their proposal to use special trucks with inflatable tents that are placed over the splice huts and maintained at positive pressure with filtered air during splicing is an indication that careful thought and experience are informing their proposal. (This general idea is used in the U.S. for fiber cable repair during bad weather conditions.)

I.2. Credibility

Has the information presented come from a reliable source with sufficient expertise?

Yes, with some reservations for the cross-border connections

Credibility is indicated most importantly by the fact that two radio telescope projects (MeerKAT and KAT-7) are already installed and operating on the site. Details such as the splicing cabinets suitable for harsh environments that are intermediate between the instruments and the Data Processor are described in some detail as is the equipment that will be used to actually perform the splicing in the harsh environment (e.g. the trucks equipped with positive air pressure tents that can cover and protect the splice cabinets during the splicing operations).

Based on this experience they propose subcontracting all of the construction work to commercial entities. Using this approach they can choose contractors who have a demonstrated ability to produce results in the South African environment. The level of information presented does not permit an in-depth assessment of the **risks involved** in establishing the cross-border connections. There is – naturally – a lack of experience in this area at the ambition level of the SKA project.

I.3. Cost Estimates

The cost estimate for installing fiber in the core site of €[redacted]/km appears to be realistic. One direct experience comparator is a new fiber install in suburban / rural, but not desolate, areas in North America where the cost was approximately €[redacted]/km of installed fiber. Given the substantial differences in physical, political, and economic environments, the 1:1.5 agreement can be considered a "close" match that supports the S. African estimate.

The cost estimates are complete except for lacking a cost estimate for the passive optical network. See also under Gaps.

I.4. Capability

Does the proposed implementation deliver the capability that is required?

Yes

The South African proposal is based on engineering designs and cost estimates from companies that have demonstrated an ability to do the work that they are proposing. This is seen as an existence proof and therefore demonstrates credibility. That is, it demonstrates that the work can be done and the capability delivered. The actual vendors will presumably be chosen by a public tender process that takes into account both cost and ability to do the work.

I.5. Gaps

Gaps in the site's response to the RFI or deviation from specification

There is a strawman estimate for active WDM equipment in the amount of approximately €[redacted] (which is a reasonable estimate for active equipment assuming certain technology advances that are neither conservative nor aggressive). However, this number is not in the budget summary.

One thing that appears to be missing is the cost of passive wave division multiplexing (WDM) equipment that surely will be required to multiplex multiple data channels onto single fibers in order to keep the fiber count to a workable number. The passive equipment will be considerably less expensive than active equipment, but still involves high power line drivers ("long-reach optics") for the interfaces between the telescope instrument modules and the passive WDM (optical multiplexers), the many WDM devices required to get the telescope signals onto the fiber, and some (probably substantial) number of optical amplifiers that will likely be needed to reach the whole 180km area. So even using a passive system the cost will not be negligible, and should be included in the budget. This is considered a gap.

Granted that the final design and state of technology will impact this, but enough is known about the basic parameters of the source and fiber that something should have been said about this issue.

Apparently missing from the **risk analysis** is the potential for damage to the fiber infrastructure from flash floods. In American deserts, flash floods due to infrequent, but not non-existent, heavy rain can cause deep gullies to form in a matter of a few hours, washing out everything in their path to depths of several meters.

I.6. Sequencing

Does the plan facilitate a smooth rollout?

Cannot be assessed due to lack of information.

Very little is given in the proposal in the way of time estimates, so nothing can be said about sequencing.

II. From the Data Processor to the Super-Computer Centre, Co-located at the Astronomy Site.

The proposal is based on co-locating the super-computer centre with the data processor at the SKA Core Astronomy site.

II.1. Feasibility

Is the proposed solution a logically possible proposition?

Yes

The Proposal is based on locating the supercomputer centre at the SKA Core Astronomy site with the Data Processor, this makes the connection between the two straightforward as only passive fiber is required.

II.2. Credibility

Has the information presented come from a reliable source with sufficient expertise?

Yes

As mentioned above the connection is based on short distance passive fibre.

II.3. Cost Estimates

The cost will be negligible.

II.4. Capability

Does the proposed implementation deliver the capability that is required?

Yes

The several engineering proposals that cover this work provide the required capability.

II.5. Gaps

Gaps in the site's response to the RFI or deviation from specification

There is no analysis of the potential impact on the operation of the telescopes created by potential radio frequency emissions from the super-computer centre. Nor is the cost of any mitigation of this **risk** evaluated.

II.6. Sequencing

Does the plan facilitate a smooth rollout?

Cannot be assessed from the material

Very little is given in the proposal in the way of time estimates, so nothing can be said about sequencing.

IIbis. From the Data Processor to the Super-Computer Centre, Remotely Located Alternative

Although the proposal is to co-locate the super-computer centre with the data processor at the SKA Core Astronomy site, an alternative solution where the super-computer centre is located remotely has been presented – including a detailed cost analysis – in the response. In the interest of the overall evaluation – the Australian-New Zealand proposal is based on placing the super-computer remotely – the panel has also assessed that part of the response.

IIbis.1. Feasibility

Is the proposed solution a logically possible proposition?

Yes

With the supercomputer located remotely, presumably at the SKA Main Office in Cape Town, the connection is still "straightforward," but very expensive.

The 400 Tb/s connection between the Data Processor and the remote supercomputer requires some 42 x 96 channel x 100 Gb/s/channel DWDM systems on each of two independent paths for redundancy. This is an enormous amount of optical transport equipment to maintain, manage, and operate. For example, it substantially exceeds all of the optical transport equipment used in all three of the major U.S. national R&E networks combined. So while engineering the system is probably straightforward, operating it reliably may raise **substantial risks**.

IIbis.2. Credibility

Has the information presented come from a reliable source with sufficient expertise?

Yes

As with the core area network, the engineering and construction tasks will be subcontracted. Commercial vendors have been approached for a preliminary estimate and there is no reason to believe that they cannot accomplish what they have proposed. There has not been any competitive bidding at this stage and clearly this should take place in due time.

IIbis.3. Cost Estimates

The cost of the 42 x 2 (redundant paths) x 2 ends (assuming no regeneration, so 1 system at each end of the fiber runs between supercomputer and Data Processor) is provided by NSN (Annex G.23) based on 400 Gb/channel systems. The cost provided (taking the bi-directional system for comparison purposes) can be compared to today's cost of such a system. A rough calculation of the cost of such equipment as it is being deployed in the U.S. today (assuming the same levels of discounts by both suppliers to R&E institutions in both countries) and based on 100 Gb/s DWDM optical systems gives a capital cost of roughly USD [redacted]. Even assuming that 400 Gb/s/channel results in 25% of the cost (which it will not), this gives USD [redacted] (about EUR [redacted]). This compares to the EUR [redacted] in the vendor proposal used in the response as a basis for the cost estimates.

So, assuming that the proposal is for the 42 x 2 x 2 systems noted above, the pricing provided, and the resulting budget, seems optimistic. Under the assumption of a 50% price reduction in the 5+ years to acquisition brings the comparator somewhat into line with the NSN estimate, though the factor of 4 price reduction for 400 Gb/s/channel is optimistic (historically, a factor of four increase in bandwidth results in a factor of 2 price reduction).

Figure 4.7 ("SKA Network Expansion ..") in the proposal implies that there are five elements in the 400 Tb/s fiber network: 1) A 115 km fiber build from Data Processor to Kronos; 2) PoPs at Kronos and Carnarvon; 3) existing long haul fiber; 4) PoPs in Cape Town, and; 5) SANReN metro fiber in Cape Town.

Risks: There are ambiguities due to the lack of detail in the proposal budget. For example, 1) are there regeneration terminals needed in the Data Processor to supercomputer path? If regeneration is needed (and the distances are right at the edge of where they might or might not be needed) then another complete set of WDM terminals

will have to be housed in a commercial PoP. This is expensive. In the U.S., cost of this would be something like USD [redacted]/yr/rack and the power to run a fully loaded chassis. 400 Tb/s requires 42 of these x 2 for the second (redundant) path. This could be something like USD [redacted]/yr in additional operational cost. 2) Is the 115 km build included in the Infraco quote? 3) Is the SANReN metro fiber included? Etc.

The conclusion is that while the proposed budget might be reasonable, it is really not possible to confirm this for sure with the level of detail provided.

IIbis.4. Capability

Does the proposed implementation deliver the capability that is required?

Yes

The several engineering proposals that cover this work provide the required capability.

IIbis.5. Gaps

Gaps in the site's response to the RFI or deviation from specification

Risk: The operational aspects are potentially non-trivial and are not discussed.

IIbis.6. Sequencing

Does the plan facilitate a smooth rollout?

Cannot be assessed due to lack of information

Very little is given in the proposal in the way of time estimates, so nothing can be said about sequencing.

III. From the Super-Computer Centre to Data Centres in Other Parts of the World

III.1 Feasibility

Is the proposed solution a logically possible proposition?

Yes

The objective of this part of the system is to transfer data from the supercomputer building to data centers elsewhere in the world where scientists will use these data in specific R&D activities related to the overall objective of the SKA Project. According to the RFI for Candidate SKA sites the globally located data centres require data rates of approximately 100 Gbps. In this document it is stated that this will be standard IP routed data, but the globally located science centre will be fixed and so point-to-point connections are suitable for this link. (It is assumed that a network operator will undertake operation of the network on behalf of the SKA).

The proposed solution is feasible with existing technologies and with the technologies expected to exist when the systems enters into operation (a few years from now) we expect that no major engineering problems will exist in the implementation of the proposed solution. However, it is expected that many of the data centres around the world to receive traffic from the SKA project rely on the local National Research and Education Network (NREN) for the main transmission capacity. **Risk:** The response is based on the network out of SKA being completely outsourced to a commercial provider and should therefore have details as to how this operator plans to have connectivity at 100 Gbps with the relevant NREN networks, since it is not trivial to cross network boundaries at these speeds.

III.2. Credibility

Has the information presented come from a reliable source with sufficient expertise?

Yes

The regulatory environment of telecommunications in South Africa is described as "vibrant" in the response and has contributed to the introduction of a number of operators. However, the experience in the telecommunications liberalizations in other parts of the world shows that it is very difficult to break the dominance of old incumbents. Especially in leading edge services, where competition is scarce, the effects of liberalization in availability of services and their costs.

The recent investments in submarine optical cables in Africa (East and West Coast) are extremely significant events and represent a major improvement in the international connectivity that impacts this project. Indeed with these optical cable infrastructures adequate connectivity to other parts of the world can be achieved (Europe, Americas, Asia and Pacific Rim and, of course, Africa).

Concerning the ageing of the fibres and the technological evolution of optical equipment the South African Response to the SSG Request for Information:
Queries and Response of 17 October 2011, namely sections 4.3 and 4.5 clarify the proposal in an adequate way.

The additional responses provided on 21 October 2011, in section 4.8 (Constraint in Transmission Protocols) clarify the proposal.

Conclusion: The proposal addresses most of the relevant topics in an adequate form but additional details may be necessary due to the time distance for the implementation and the lifetime of the project.

III.3. Cost Estimates

The cost estimates presented form a good basis for the response to the RFI and fit inside the budget, but do not cover all the costs that may be incurred during the lifetime of the project and will have to be revised once time gets closer.

Due to the long lifetime of the solutions, it is expected that major technological improvements will occur, resulting in possible cost reductions and this should be taken into account (this is mentioned in the reply to the questions posed by the Expert Panel).

III.4. Capability

Does the proposed implementation deliver the capability that is required?

Yes

The proposal covers most of the relevant technical, organizational and financial aspects of the project and, together with the telecommunication providers, they have the capability to implement and operate the systems.

III.5. Gaps

Gaps in the site's response to the RFI or deviation from specification

The proposal mentions IRUs of 20 and 50 years for different parts of the communication infrastructure. There are different modalities of IRUs and we consider it is necessary to detail several aspects related to the contractual options, namely, replacement of aging fibres (if aerial fibres are used in parts of the infrastructure special care should be taken). Although a lot of emphasis is given to the submarine cables the proposal lacks some detail in the overall fibre segments.

The slides provided with additional information related to the SEACOM Final Proposal (included in the PDF of 17-10-2011 16:58 – SEACOM Final proposal, dated 31 Aug 2011) exclude several costs and these should be carefully analyzed concerning their impact in the investment and operating costs of the infrastructure. An estimate of these additional costs must be performed for different scenarios, in order to be used in a **cost/risk** analysis.

The proposal need further clarification on how peering will be achieved at the proposed speeds with the relevant places in other parts of the world where data connectivity at 100 Gbps to other NRENs is necessary. Again with time this will evolve under pressure from the research community.

Risk: The fact that the electricity grid network will be extended by means of a single 132kV line may represent a problem since if there is a failure in this line all the Astronomy Complex will be down for a period of time.

An analysis of the optimal solution to access all of the regions of the world (North America, Asia including China, India and Japan), including costs, will be needed as we get closer to the implementation.

III.6. Sequencing

Does the plan facilitate a smooth rollout?

Yes

We have not identified any problems in the proposal and what is presented looks coherent.

IV. Monitor and Control (M&C) Services

The response to the RFI did not include a full description of the M&C system. This is not considered a gap since the RFI could be interpreted as not requesting it.

IV.1. Feasibility

Is the proposed solution a logically possible proposition?

Cannot be assessed in depth due to insufficient information.

Because of the very limited level of detail given in the answer, it is difficult to give an in-depth assessment of the proposed solution. The provision of a 10 Gbps bidirectional link between the core and each remote station is straightforward, but the description is rather vague as it does not specify whether the link will be supported by a dedicated fiber pair or by a specific wavelength among others in the transmission equipment. These scenarios are all feasible, but the technical study is still to be done to address specific issues such as guaranteed maximum latency, SLA for real time application etc.

The proposed back-up solution, using VSAT terminals may be considered as well, however the switching from an optical support to a VSAT channel for M&C cannot be considered as straightforward as it will require to be revisiting all calibration steps (especially from the latency/jitter aspects).

Risk: It is hard to believe that a nominal 10 Gbps fiber link can be backed-up with a 512 kbps VSAT link.

IV.2. Credibility

Has the information presented come from a reliable source with sufficient expertise?

Cannot be assessed due to insufficient information.

We have no position on the credibility of the proposed solution, at this stage of the proposal, as too much information is missing. The credibility can be assessed only based on much more detailed information.

IV.3. Cost Estimates

No information is given on the cost estimates. The addition of a 10 Gbps circuit on the data transmission system is not to modify significantly the overall cost estimate.

The cost of using a dedicated fiber is not estimated.

IV.4. Capability

Does the proposed implementation deliver the capability that is required?

Yes

For the central part (≤ 180 km) the provision of a 10 Gbps link, either as part of the data system or via a dedicated fiber should not be a difficulty for the SA Telco.

IV.5. Gaps

Gaps in the site's response to the RFI or deviation from specification

There is **no information about the final solution** for the whole M&C sub-system.

IV.6. Sequencing

Does the plan facilitate a smooth rollout?

Cannot be assessed in depth due to insufficient information.

No information about sequencing of implementation of the M&C sub-system in the proposal, but it should be required in the early phases of the project deployment.

V. Timing/Synchronization Services

V.1. Feasibility

Is the proposed solution a logically possible proposition?

Yes

The only option considered in the proposal is to use hydrogen masers, which are highly reliable, and perhaps manageable through the M&C sub-system at the remote station. These devices can use GPS to align their time together (from time to time but it is necessary to monitor them continuously (check their status, reboot? etc...).

The central distribution of time/synchronization signals is not considered as being prohibitory expensive.

V.2. Credibility

Has the information presented come from a reliable source with sufficient expertise?

Yes

There is no serious issue with the implementation of the hydrogen maser as this is already existing in radio-telescopes, including the South African HeartRAO. However, the long distance between remote stations and SKA core may lead to specific problems, which still need to be identified or clarified.

The unique choice of hydrogen maser seems to be a quite conservative attitude compared to the alternate technology (fiber transport), which is evolving/improving very quickly. The current accuracy provided by masers, with GPS support for synchronization is, however, probably sufficient for the today needs.

Risk: Only when the final design has been made by SPDO and technical guidance can be given can the proposed solution be assessed properly.

V.3. Cost Estimates

The relative cost of hydrogen masers versus fiber transport has not been fully analyzed although it is explicitly stated that cost is the driver behind the choice of the maser solution in the proposal.

V.4. Capability

Does the proposed implementation deliver the capability that is required?

Yes

Not an issue for the proposed option, although more clarifications will be needed as the project design progresses.

V.5. Gaps

Gaps in the site's response to the RFI or deviation from specification

The Proposal relies heavily on commercial Telecommunications Providers, even for the inner part (≤ 180 km), thus and excludes managing the fibre directly and deploying dedicated optical components (in-line and edges) for T&S.

Even a hybrid solution (fibre transport in the core and maser in the remote stations) has not been considered.

V.6. Sequencing

Does the plan facilitate a smooth rollout?

Yes – but there is a risk

Risk: The choice of hydrogen masers may pose problems in the long run where fibre transport has the advantage of opening up to a much better accuracy/sensitivity, which may be required in the ultimate evolution of the project (SKA-2)..