Data Flow Advisory Panel
Report to the SKA Board

17th March 2016

Paul Alexander	University of Cambridge, UK
Ian Bird	CERN
Gary Davis (chair)	SKA Office
Miles Deegan (secretary)	SKA Office
Jack Dongarra	University of Tennessee and SEAC
Brian Glendenning	NRAO and SEAC
Jasper Horrell	SKA SA, South Africa
Melanie Johnston-Hollitt	Victoria University of Wellington, New Zealand
Chris Loken	University of Toronto, Canada
Niruj Mohan Ramanujam	NCRA, India
Sergio Molinari	INAF, Italy
Peter Quinn	ICRAR, Australia
Sarah Pearce	CSIRO, Australia
Nick Rees	SKA Office
Russ Taylor	University of Cape Town, South Africa
Michael Wise	ASTRON, The Netherlands
Meng Zhao	NAOC, China

1. Introduction

The Data Flow Advisory Panel (DFAP) was initiated by the SKA Board at its meeting in July 2015. The remit of the Panel is to provide the Board with advice on how to optimise the data flow system for the SKA, which is at present undefined, to ensure that science results can be efficiently extracted from SKA data. The Terms of Reference of the DFAP are attached at Annex.

Because the DFAP is large, diverse and geographically distributed, a Steering Group was set up to guide its work. The Steering Group members were Alexander, Davis (chair), Deegan, Quinn, Taylor and Wise. The Steering Group met on 5 occasions to explore the issues and the parameter space, to generate an initial set of documents for review and discussion by the Panel, and to plan for a face-to-face meeting of the DFAP.

The face-to-face meeting was held at the SKA Office on 3–4 March 2016, and most members of the DFAP were able to attend either in person or by videoconference. This report is based in large part on deliberations at that meeting.

In addition to the named members of the DFAP, three observers participated in the face-to-face meeting: Rosie Bolton (SDP), Antonio Chrysostomou (SKAO) and Keith Grainge (SaDT).

2. The Challenge

2.1. Background

In July 2013, the Board approved a set of top-level principles governing SKA operations (SKA-BD-10-13). Three of these principles were as follows:

20. The SKA Observatory will calibrate SKA data and make science-ready data and ancillary products available to the users.
21. The SKA Observatory will provide an archive with a data management system to support data-intensive astronomy.
22. The SKA Observatory will provide user support and tools to enable exploitation of the SKA data.

At the same meeting, the Board imposed a cap on the construction cost of SKA1, and in so doing defined the scope of the project to include the generation and storage of data products by the Science Data Processor (SDP). No provision was made for the distribution of data to users, nor for computational facilities to enable users to undertake further data analysis, both of which are mission-critical if the SKA is to deliver on its scientific promise. It has been widely assumed, but never specified, that these functions would be provided from regional[1] funding rather than from project funds. The expectation that the solution would include a network of regional centres was clearly stated in the DFAP's Terms of Reference.

[1] Throughout this report, we use the terms "region" and "regional" generically to refer to an SKA Member State, an area within a State, or a grouping of States; thus "regional" could be replaced by "national" in the case of a single State.

2.2. Current provision

The scope of SDP is currently under review with the SKA Office in the context of capital and operational cost estimates. This work is urgent but will not be concluded before the DFAP's submission deadline. For the purposes of this report, therefore, the DFAP assumes that the SDP will be delivered as presently scoped. If this should change (for example, if some processing currently scoped within SDP is moved outside the project) then the DFAP's recommendations should be reviewed.

The SDP will ingest data from the correlator at an average rate of 1.5Tb/s (across both telescopes). A set of standard data products, appropriate for the range of anticipated science investigations and comprising mainly images and catalogues, has been defined. Visibilities will in general not be retained nor provided, with one exception: the highest-priority project for SKA1-LOW, the Epoch of Reionisation (EoR), requires visibilities for effective foreground removal. The rate of production of science data products by the SDP is estimated by the SDP Consortium[2] as follows:

SKA1-LOW excluding EoR:	3 Gb/s
EoR only:	22 Gb/s
SKA1-MID:	9 Gb/s
TOTAL:	34 Gb/s = 370 TB/day = 130 PB/yr

It is evident that the EoR experiment dominates the data rate.

The SDP architecture includes, in addition to the pipelines that generate these products,

(a) a Long-Term Preservation (LTP) component to store these products for as long as considered necessary (TBD), and
(b) a Delivery component for making data products available to the outside world using VO-compliant protocols.

It is important to note that the LTP is not a science archive: its role is to serve as a backup for the SDP data products.

The data rate estimate above is for standard data products. Although visibilities will in general only be provided for the EoR experiment, this is motivated by cost and there is nothing in the SDP architecture that prevents visibilities (or other intermediate data products) from being provided to users if there is sufficient scientific justification for so doing and if resources for this are made available. The SKA will offer the capability of delivering non-standard data products under the umbrella of "custom experiments".

[2] The volume of science data varies strongly over the range of anticipated science projects, and the data rate is therefore very sensitive to the scheduling of the different projects on the telescope. The SDP Consortium has used a strawman schedule provided by the SKA Office to derive the estimate provided here. PI projects, which will take up a significant fraction of the telescope time, have unknown data volume requirements and have not been taken into consideration in generating this estimate. The data rates presented here are therefore subject to some considerable uncertainty.

2.3. Absent from current provision

Full scientific exploitation of the SKA requires that a research ecosystem be in place for efficiently translating the large data volume into science results. The DFAP is aware of other observatories whose productivity has suffered from the lack of such an ecosystem.

This research ecosystem must have three components that are not within the current scope of the SKA1 Observatory:

1. *Computational capacity for re-processing and science analysis*. Users will not, in general, have sufficient capacity in their home institutions. The nature of the science analysis will vary widely from project to project: source extraction, Bayesian analysis, consolidation of catalogues, visualisation, etc. The computational requirements are similarly variable from case to case and it is therefore difficult to estimate the overall capacity requirement. Any divergence from the current plan to provide visibility data only for the EoR experiment (§2.2) would also have a significant impact. In the absence of any other information, an upper limit can be derived by assuming, in aggregate terms, that the science analysis requires the same amount of processing as the science pipelines in the SDP, i.e., of order 100 PFlops. More work is required to specify this requirement, and the corresponding short-term storage need, more accurately. (For reference, many high-performance computing centres offer more than 5 PFlops; the current state of the art is 10–15 PFlops; we anticipate 100-PFlop centres becoming available within the next decade.)

2. *Storage capacity for archiving SDP and derived data products*. Science archives, in which standard and advanced data products comprise an HST-style data heritage accessible to the worldwide community, are multipliers for total science output. We note here the need for a clear policy statement from the Board, based on principle 21. A full science archive in which basic and advanced data products and their provenance are preserved and made available to users is a non-trivial undertaking, involving not just computational facilities but people. The current policy for Key Science Projects (KSPs) requires that the KSP teams place their derived data into such archives, and this policy could also be extended to PI projects. The DFAP recommends that the SKA Office develop a data archiving policy statement for the Board's approval. An overall storage capacity requirement can be estimated by assuming the advanced data products to be equal to the data rate produced by SDP; archiving both would require 2 x 130 = 260 PB/yr (but see §6).

3. *Local user support*. Although this is straying beyond the remit of the DFAP, we raise it here since it is an essential element of the ecosystem (principle 22). The operational model of the Observatory includes a central user support function: this is where the progress of observations will be monitored and the detailed knowledge of the telescope behaviour will reside. The Observatory cannot, however, provide functions such as training and one-to-one support for proposal

preparation and data analysis. Bearing in mind the diversity of the SKA user base, these functions are best provided on a regional basis.

It is evident from this list that the aspirations represented in the Board's operational principles (§2.1) cannot be met within the current scope of the project.

2.4. Data transport considerations

The DFAP received a presentation from the SaDT Consortium on the feasibility and cost of transporting the SKA1 data volume around the world from the Host Countries (Australia and South Africa); comparison was also made with the Worldwide Large Hadron Collider Computing Grid (WLCG) project, which provided an extremely useful reference point. The conclusions from these considerations were:

- that transporting data from the Host Countries to the SKA Member States by 100 Gb/s fibre networking links is both feasible and sufficient; and
- that the cost of data transport, estimated by SaDT at US$5M/yr, is small enough that it should not drive or constrain the Observatory's data flow strategy.

This cost estimate is the lower bound of the range provided by SaDT, US$5–15M/yr; achieving it requires early engagement with the National Research and Education Network (NREN) partners. There are also technical issues that are straight forward but need to be addressed. The DFAP recommends that the SKA Office commence this engagement promptly.

Some concerns were expressed in the DFAP meeting about network links in and out of China. The WLCG experience in this area indicated no cause for concern, as data transport to and from a data centre in Beijing had not encountered any difficulty. The DFAP's view is that the existing links, coupled with continued growth and integration over the next decade, should enable full Chinese participation in the network. If problems are encountered, it will of course be possible to manually load data onto discs for shipment, although this should be avoided if at all possible.

3. The Solution

3.1. Top-level options

Having established the elements that are absent within the current project provision, the Board must now choose between two modes of implementation: centrally funded through the SKA1 construction project, or distinct from the project and using regional funding. The merits of both alternatives are considered in the following paragraphs.

Adding the missing elements to the project scope would lend itself to a centralised solution, in which the project would be responsible for securing and operating the necessary computational resources. One potentially attractive option might be to locate such facilities in the Host Countries, co-located with the SDP facilities; we have not, however, conducted an analysis of this or any other option. Although the change in scope would increase the cost of the project very substantially, the advantages of this approach are significant: this solution

would likely be less expensive when viewed in aggregate, and it would be under the direct control of the Council.[3]

The other alternative is to rely on regional funding for the provision of the missing elements. The advantages of this approach are:

- it recognises the understandable preference of any country to invest in infrastructure within its own borders;
- resources can be tailored to the local needs of the regions, which are diverse across the SKA partnership;
- it offers the opportunity for early community engagement with the SKA, not only for astronomers but also for data scientists and software and computing experts; and
- it offers the opportunity to leverage existing computational infrastructure across the SKA partnership.

We recommend the Board give full consideration to the top-level choice between central and regional provision of critical functions. Both offer significant advantages and also pose significant risks. Because the Board has already expressed its preference in the Terms of Reference, the regional model will be assumed for the remainder of this report.

3.2. Network of Regional Centres

The above considerations lead to the concept of a network (or alliance or federation) of accredited Regional Centres (accreditation is discussed in §5). We adopt the generic name SKA Regional Centres (SRCs) although, following the ALMA example, we recognise that individual SRCs may wish to adopt different names for their own purposes and within their own regions. We recommend that this nomenclature be adopted from this point forward.

The essential functions of the network are:

1. to provide SKA users with access to SKA data, in compliance with the SKA data access policy[4]. It is not envisaged that individual users will obtain data directly from the SDP facilities in the Host Countries. Given the rapid development of cloud storage solutions and international networking, the physical location of stored data is no longer a serious constraint on the design of this network[5];

2. to provide SKA users with access to computational resources for processing, archiving and science analysis. It is expected that this will include a blend of existing and new infrastructure, and that such infrastructure may not be dedicated to the SKA, but shared with other computational activities. The capacity requirements are estimated in §2.3. It is envisaged that at least one complete copy

[3] This paragraph applies only to items 1 and 2 in the list of missing elements, i.e., processing and storage; it is self-evident that local user support must be provided locally.

[4] SKA Member States will have data access rights as a matter of policy irrespective of the provision an accredited SRC. It is anticipated that most, if not all, Member States will either host an accredited SRC or collaborate in a multi-national, regional SRC.

[5] The EoR experiment may be an exception to this statement due to the very large data volume (§2.2).

of the standard data products produced by the SDP will be archived across the network, in addition to the advanced data products generated by science teams;

3. to provide a federated environment which allows transparent data access across the network to enable science extraction to be achieved efficiently by all members of the SKA community. It is the DFAP's view that software tools for common tasks such as visualisation, co-addition, stacking, etc., should be developed by the community rather than centrally imposed. The SKAO will work with the SRCs to support the development of standards where appropriate (such as data format and access protocols, run-time environment commonality, and user authentication methods) to encourage interoperability across regional centres to allow efficient sharing and collaboration; and

4. to provide users with local user support, as described in §2.3.

The SKA user base is large and diverse, and the DFAP recognises that the regions may wish also to provide additional, non-essential functions through their local SRC. Some examples of possible value-added functions are:

- providing a point of contact in each region for relations with governments and with the SKA Observatory;
- outreach, communications, publicity, etc.;
- expertise in data management, particularly in support of the KSPs; and
- development activities, including software development and development of new technologies for potential upgrades of SKA1 or towards the design of SKA2.

Such additional functions are optional and at the discretion of the regions.

4. Resourcing the Solution

The SRCs are to be regionally resourced. How this occurs is within the domain of each region. We expect funding models to vary according to structure, opportunities, aspirations, etc. Although the SKA Office will provide information and supporting effort as required, it is the responsibility of the regions to self-organise and secure the resources to put the SRCs in place.

There are nevertheless three major issues for the Board to consider.

The first is whether the Board is comfortable with the concept that functions that are critical for the scientific success of the SKA are not under the Board's direct control, but are provided by the regions. In the ALMA case, for example, essential functions are supported through the observatory's operations budget, whereas value-added functions are supported by the regions. This issue was raised in §3.1 but bears repeating here; in the final analysis it will be for individual Members to ensure that there is adequate provision for their own user communities. This risk can be mitigated to a degree through an appropriate governance arrangement (§5).

The second is whether such contributions from the regions should be recognised in some form, e.g., share in the project. This is beyond the remit of the DFAP but we recognise the

implications of the issue and we recommend the Board give it serious consideration. The Operations & Access Working Group of the IGO negotiations discussed this but did not reach a conclusion.

The third is whether, having identified a network of SRCs using regional funding as the preferred solution, any of the associated costs might more appropriately be brought under the observatory operations budget. Options might include critical archival storage, software development or intercontinental networking costs. The current assumption is that the operations budget is restricted to the operation and maintenance of everything within scope of the construction project; on instruction from the Board, however, that definition could be changed. This is again beyond the remit of the DFAP and it needs to be considered in light of the overall funding context.

5. Managing the Solution

The essential functions listed in §3.2 apply to the network as a whole. A key aspect of the network concept is that individual SRCs may choose, according to their local needs, only to provide a subset of those essential functions. Recognition of the diverse aspirations and capacities of the SKA regions is essential to the success of the scheme.

Translating the essential functions for the network into requirements for individual SRCs is therefore somewhat complex. The approach we recommend is the one adopted by the WLCG project, in which individual centres pledge on a periodic basis to provide processing and storage resources. The pledges would vary from region to region but it is anticipated that the pledged resources across the network would deliver the required overall capacity.

The requirements on individual SRCs would then be:

1. to comply with and enforce the SKA data access policy (currently being negotiated);
2. to comply with and enforce the SKA data archive policy (yet to be developed, see §2.3);
3. to deliver the agreed processing/storage resources;
4. to make data products, which may be uniquely stored in one region, available on request to other SRCs in the network (and thereby to SKA users in other regions);
5. to meet minimum SKA standards around critical functional aspects (security, interoperability, etc.); and
6. to provide a minimum standard (TBD) of user support.

We recommend that all SRCs be accredited by the SKA Observatory, and that they sign a Memorandum of Understanding (MoU) with the Observatory in order to confirm their acceptance of these requirements. The accreditation process would provide a necessary opportunity for confirmation that the offered resources meet the required standard. Accreditation would then permit full integration into the SRC network, including access to the central user support functions of the Observatory (§2.3).

In the short term, we recommend that the SKA Office set up an SRC Coordination Group (SRCCG) to take this concept forward. The remit of the SRCCG should include at least the following:

- further developing requirements for the network as a whole (§3.2),
- further developing minimum requirements for individual SRCs,
- developing the pledging and accreditation processes,
- developing the MoU,
- monitoring progress in the regions towards securing resources,
- engaging with the NRENs, and
- carrying out a series of data challenges to test and verify network communications, access protocols, data management, etc.

As the process moves forward we would expect the SRCCG to evolve into a somewhat more formal coordination body to implement the terms of the MoUs. In the longer term we also see a need for an oversight body to provide assurance to the Board that the essential functions are being provided as planned and that the network remains fit for purpose to deliver the scientific goals of the SKA.

We recommend that the regions commence work immediately to identify the appropriate resources, and that the SKA Office provide support as required. In the WLCG case the process from project inception to a fully operating data network took 7 years, and although the time scale should not be quite so long for SKA the work nevertheless needs to commence without delay if we wish to have the SRC network in place for the start of full science operations in 2023.

We note that there is an immediate and continuing requirement for data management effort in the Office, and that this effort is currently not available. We recommend that the Director-General consider how this might be arranged.

6. Discovery Data

According to the current design of the SDP, the standard data products will be tailored to each observing project's individual requirements (area, spectral resolution, etc.). The motivation for this is to reduce costs by only generating the data requested by the observing project. It is possible, however, for every observation, to generate data products to the maximum resolution in three dimensions (two spatial, one spectral) and in all polarisations or images at high cadence; rather than throwing away good data as in the current design, valid science data could be produced that would go beyond the scope of the science project being observed. This would permit serendipitous science, increasing the SKA's productivity and reducing repetition of observations.

The price for this increase in discovery space is a large increase in data volume. If it is assumed that the data volume is limited initially by the capacity of the 100 Gb/s network rather than by cost, then the total volume increases from 130 to 710 PB/yr. This has implications for data transport, processing and storage – all the same issues as discussed above, but magnified.

The DFAP is enthusiastic about the possibility of increasing the scientific discovery potential of the SKA. The implementation requires further examination before a clear recommendation can be made. Since the Office and the SDP Consortium are already reviewing the scope/capex/opex of the element, the DFAP recommends that this topic be added to that discussion.

7. Risks

The DFAP has identified the following risks in the proposed strategy that will need to be managed:

1. Although the processing and storage capacity requirements for the network as a whole have been estimated in this report, it will be for the regions to determine what capacity they are able to offer individually. We expect this to vary from region to region according to local context. It is possible that the aggregate offerings will not meet the overall requirements. We can think of no mitigation for this other than for the regions to commence work immediately to secure the required facilities.

2. The distributed network proposal gives the regions flexibility to implement solutions that meet local requirements and aspirations so long as they meet the minimum set of requirements set out in §5. There is a risk that regions will opt instead, for reasons of their own, for solutions that do not conform to the minimum requirements. (This is the generic risk associated with in-kind contributions: they rely on a shared motivation for collaborative, rather than local, success.) The accreditation process in §5 is intended to mitigate this.

3. Supporting the development and implementation of this network will require dedicated effort from the SKA Office, which is not at present available. Failure to identify this effort promptly will delay the implementation.

4. The cost of international networking will depend on adequate capacity being available and negotiated well ahead of time. Early engagement with the NRENs is recommended as mitigation.

5. It is anticipated that most, if not all, Member countries will wish to set up an SRC (or to set up a collaboration with another SRC) to make provision for their user communities. Under the regional model, however, this is not guaranteed and there remains a risk of some SKA users not having access to a local SRC. We can think of no mitigation for this; in the regional model, it is incumbent upon the Members to ensure that adequate provision is made.

8. Recommendations

The Data Flow Advisory Panel makes the following recommendations:

1. The SKA Board should carefully consider the top-level choice between central and regional provision of critical functions. Both offer significant advantages and also pose significant risks. Although the Board has considered these options in the past and has expressed its preference for the regional model in the Terms of Reference, the DFAP believes this choice should be formally confirmed at this time.

2. Assuming the regional solution is confirmed, then the SKA Board should endorse the proposed strategy of a collaborative network of SKA Regional Centres (SRCs) to provide the essential functions that are not presently provided within the scope of the SKA1 project. Membership in the network should require an accreditation process and should be certified through an MoU. Provision of resources to the network should be based on a pledging process.

3. The SKA Office should produce a data archiving policy document for the Board's approval at its next meeting. A clear policy in this area is an essential prerequisite in order to define the requirements of the SRC network.

4. The SKA Office should set up an SRC Coordination Group (SRCCG) to take this concept forward. The minimum remit of this Group is described in §5.

5. All regions should immediately commence activities towards securing resources for setting up local SRCs.

6. The SKA Office should commit staff effort in the area of data management. The initial requirements are: to develop the data archiving policy; to set up the SRC Coordination Group (SRCCG); to support the regions in securing resources; and to engage with the NRENs. The Office will need to participate in, and lead some of, the SRCCG activities. In the longer term, the Office will need to participate in the coordination and oversight of the SRC network. This effort is not currently available and we recommend the Director-General consider how this might be arranged.

7. The SKA Board should consider the wider implications of the regional funding model, including (a) whether the provision of SKA Regional Centres should be recognised in some form, and (b) whether any of the costs associated with the network of SRCs should be funded from the Observatory Operations budget.

8. The SKA Office and the SDP Consortium should consider, within their cost/scope discussion, all possible options for making the discovery data available to the user community.

9. Finally, we recommend that the SKA Board accept this report and disband the Data Flow Advisory Panel.

Respectfully submitted,
Prof. Gary Davis,
Director of Operations Planning, SKA Organisation,
Chair of the Data Flow Advisory Panel.