



Memo 84

Report of the SKA Operations Working Group

K. Kellermann et al

10/06

Report of the SKA Operations Working Group

v6.2

K. I. Kellermann, T. Cornwell, D. DeBoer, R. D. Ekers, M.
Goss, A.J. Green, P. Hall, J. Tarter, R. Vermeulen

October, 2006

Report of the SKA Operations Working Group

Table of Contents

Executive Summary

1. Introduction
2. Operational Model
 - 2.1 Overview and assumptions
 - 2.2 Interface with construction – commissioning timeline
 - 2.3 Tendering and contracts - outsourcing
 - 2.4 Regional and national centers – Centers of Expertise
 - 2.5 Human resources
 - 2.6 Innovation
3. Governance, management and oversight
 - 3.1 International aspects of the SKA
 - 3.2 Funding model for contracts – construction and operations
 - 3.3 Coordination of geographically dispersed activities
 - 3.4 Impact on national radio astronomy programs
 - 3.5 Inclusion of new partners
4. Technical Issues
 - 4.1 Maintenance and repair schedules and resources
 - 4.2 Renewal and upgrading of equipment
 - 4.3 Computing systems and software
 - 4.4 Power requirements
 - 4.5 Communications and fiber links
 - 4.6 Infrastructure support
 - 4.7 Operational health and safety issues
 - 4.8 RFI Environment
5. Operations and maintenance costs
 - 5.1 Cost-performance trade-off
 - 5.2 Comparison with existing arrays
 - 5.3 First estimate of operating cost
 - 5.4 Maintenance models
 - 5.5 Legal considerations
 - 5.6 Funding cycle considerations
6. Impact of Science drivers and outcomes
 - 6.1 Science innovation and impact
 - 6.2 User support and observing modes – multiple users
 - 6.3 Observing time allocation
 - 6.4 Data products and archives
 - 6.5 Student training and participation
7. Summary and Recommendations

Report of the SKA Operations Working Group

Executive Summary

The Square Kilometre Array (SKA) Operations Working Group (OWG) was formed to examine operating scenarios for the SKA and to explore how the operational model might impact its design and construction. It is clear that the design of both the hardware and software systems of the SKA will be strongly influenced by the operational model for the telescope, so these considerations are timely. Since the design, governance and siting are still to be finalized, this document only addresses high level operational issues.

Consideration of the budgets for current telescopes produce a robust estimate for operational costs of 10% of capital, split between routine operations and maintenance (5%), renewal and upgrades (2-3 %) and user support (2-3%). Since we are already using innovative concepts to reduce the capital cost per square metre, meeting even this fractional operating budget will require an operating culture substantially different from that of current large telescopes.

The design of the SKA should minimize some combination of construction and lifecycle costs, so short-term savings in construction costs which lead to increased maintenance and renewal costs are to be avoided. As cost models for the different designs are developed we need to consider parallel cost models for operations. However, experience is that government and institutional funding is easier to obtain for capital investment than for ongoing operational costs so this issue must be addressed at an early stage. Noting that the telescope will come on-line incrementally and that the telescope will undergo continuous evolution during both its construction and operating phases, we suggest a flat funding model. Assuming a 10 year construction period and capital cost of €1 billion and then a 10% operation budget, the SKA will require about €100 million annually over its lifetime.

We suggest the SKA Facility adopt the multi-tier support system proposed for the Large Hadron Collider. The multiple “tiers” would be arranged around the telescope: Tier-0 is the SKA itself; Tier-1 consists of large national scale Regional Centres; Tier-2 comprises smaller Regional Centres, mostly dedicated to analysis; Tier-3 are institute-level Centers; and Tier-4 are individual researchers. It may be attractive to have some Centers dedicated to particular Key Science projects such as pulsars or large-scale HI surveys. When setting up these Centers we need to consider the impact on and the threat to the existence of current national and institutional facilities, which are crucial for the development of the next generation of astronomers and instrument builders.

Pressures to share operations support activities across the contributing nations, which are customary in international scientific projects, may compromise the reduction of operations costs. International collaboration will require integration of components and sub-systems from different countries, and coordinating the geographic spread of activities will be a challenge. In some other international astronomy projects, the multiple lines of control and funding have resulted in misunderstandings and cost

increases. Based on what we have learned from other international scientific projects, it will be important that the operation of the SKA Observatory be vested in a Director and SKA staff with a minimum of political restrictions. Some system of *juste retour* returns in proportion to contribution, will be desired but this must be done in a way which avoids potentially inflated contracts from partners who expect their entitlement. The process of contract allocation must also avoid compromising the “open skies” policy for scientific programs which is the declared strategy for the SKA. For staff appointments, we suggest a system similar to that used for the European Space Agency (ESA) Space Telescope Science Institute (STScI) personnel, where member countries would second their staff to serve under the SKA Director. The host country will also need to facilitate the movement of people and materials across international boundaries and this will involve many details which must be negotiated in advance.

Since the SKA operations budget will be dominated by personnel costs, it is particularly important to look carefully at this component. A high level of automation and reliability will be required and the number and complexity of moving parts should be kept to a minimum. Performance tradeoffs may need to be considered by the Engineering Working Group (EWG) and the Science Working Group (SWG) to reduce operating costs. For example, reducing the reliance on cryogenics and using more robust Low Noise Amplifiers (LNAs) may result in reduced sensitivity. Further studies of the SKA configurations will also be needed to minimize the impact of failed elements or stations. For the central complex of antennas, the processing facility and the stations located within a few hundred kilometres of the core, a staff of at least 230 engineers and technicians and 30 PhD scientists, appear to be needed. For each of the more remote stations two to four individuals will be needed depending on location and access. These staffing levels are comparable to that currently needed to operate the VLA+VLBA (Very Large Array and Very Long Baseline Array, operated by NRAO), so considerable efficiency of operation will be required to reduce staffing to this level for the much more complex SKA.

In order to meet the construction cost targets, it is likely that the SKA will need to exploit the consumer and commercial markets. To protect against the obsolescence of consumer items, we may need to insure that the construction budget provides a lifetime supply of spare parts for long-lived components, and that software be portable to the next generation of computing hardware.

Power requirements to operate the SKA are likely to cost at least several tens of millions of Euros per year, with the operation of correlator/beam former and movable antenna elements being the prime drivers. Options for power provision, including solar, wind, and geothermal sources, should also be explored in relation to operational issues. As the technology for integrated circuits advances, the power requirements also grow. Just as power, and not Moore's law, is now starting to limit technology developments, a cap on the available power may ultimately be an SKA performance limiter.

Software plays a key role in all modern radio telescopes, especially those using synthesis imaging, where there is rich potential for innovation. Unfortunately, large software projects do not always meet their completion, budget or functionality objectives. The SKA will need to weigh carefully the distribution of limited resources between developing algorithms needed for the most challenging projects and providing convenient access for non expert users. Software development is very

expensive and must be included in planning the construction budget in the same way as the development of antennas, receivers, and signal processing systems. However, there will be a continual need for new algorithms and new software so there must be adequate provision for this in the operations budget. In many cases the lifetime of SKA software will exceed that of the computing hardware so it will also be important that the resources be made available to make software easily portable to new computing hardware.

The scientific return from the SKA will depend on operational flexibility but this comes at a cost and it will be important to strike a balance between support for non-routine experiments and the advantage of reduced operation costs associated with long term legacy type programs. Providing the, now vogue, end-to-end software capability may not be feasible or fundable for the wide variety of programs envisioned for the SKA. Instead, we suggest that the most cost effective approach may be to provide limited observing and processing modes which can satisfy the majority of programs, while using experts to assist with the more challenging programs.

Data management has traditionally been a challenging area for radio observatories and the SKA will be no exception. However, by the time the SKA is completed it will be able to take advantage of the developments in the International Virtual Observatory to facilitate the use of archived data from the SKA. But, this means that the SKA must have the resources to promptly provide the IVO with reduced data and interfaces.

An “open skies” policy with peer reviewed assessment of the scientific quality of proposals will give the best science returns for the SKA. We assume that nations investing in the SKA will reap direct benefits by involving their industries in ongoing construction, new instrumentation, and operations rather than through the allocation of observing quotas. But we still need to discuss whether all the participating parties will subscribe to this philosophy. These considerations also have implications for the mechanism for new partners wishing to join the SKA after construction is completed.

To move forward with many of the recommendations in this report, it may be necessary to form new operations groups within both the EWG and the SWG.

Report of the SKA Operations Working Group

1. Introduction

The Square Kilometre Array (SKA) is a radio telescope planned to have a sensitivity one to two orders of magnitude better than existing radio telescopes with a wide field of view, high angular resolution, and operation over a wide range of radio wavelengths. The SKA will be built and operated with international funding on a site chosen to optimize the scientific returns. The design and construction of the SKA must consider operational costs from the very beginning. Short-term savings which lead to increased maintenance and renewal costs are to be avoided.

The SKA Operations Working Group (OWG) was formed to examine operating scenarios for the SKA, and in particular to determine how an optimum cost-effective operation of the SKA might impact the design and construction of the SKA. Although the nature of SKA operations will depend somewhat on the design of the instrument, some general remarks can be made even at this early stage. When useful, for specificity, we assume the Reference Design described in SKA Memo No. 69, although many of the remarks will apply to other centrally condensed arrays containing a large number of elementary antenna elements, such as dipoles, tiles, steerable paraboloids extending up to a few thousand kilometers from the central core.

The governance of the SKA is not yet decided and the site is also not yet chosen. Given these uncertainties, this document sets out an initial framework for the operations model, raises issues which need to be resolved and outlines a management structure. It is clear that the design of both the hardware and software systems of the SKA will be strongly influenced by the operational model for the telescope, so it is important to address these issues in a timely way.

Current experience is that government and institutional funding is easier to obtain for capital investment than for ongoing operational costs. This is not the case for industry, where deferral of costs is preferred. The SKA project will involve industry, so the business and budget planning needs to include a synergy of these two cultures. Funding cultures differ from country to country. The European Southern Observatory (ESO) model which guarantees by international treaty a fixed budget from year to year may be attractive.

Since design and development costs as well as annual operating costs are likely to be about ten percent of total construction costs, and since it is likely to take about a decade to construct the SKA, we might think of the SKA as a continuous evolution, with no sharp demarcation between construction and operations.

Where actual costs are discussed, we use 2005 Euros (€).

2. Operational Model

2.1 Overview and assumptions: The design, governance and site of the SKA are presently still being developed. Given these uncertainties, this document reports on

high level issues that must be considered for the operations model. We assume that there are no major site dependencies in the operating model, although there will be some specific issues such as radio frequency interference (RFI), security, employment conditions, procurement policies, import/export constraints and insurance cover that may subsequently arise.

Ideally the design of the SKA should be optimized to minimize the total of construction and lifetime operating costs. Furthermore, pressures to share operations support, which are customary in international scientific projects, may compromise the effort to minimize operations costs. It is likely that the telescope will come on-line incrementally and that construction will take at least a decade to complete. The operating lifetime of the SKA is unclear, but might be of the order of 50 years. The telescope will undergo continuous evolution during its construction and operating phases. Currently operating radio telescopes have a total operating cost (including repairs, maintenance, renewal, upgrades, and user support) of roughly 10% of capital cost. For a decade of construction time, which is approximately evenly spread over the period and which then merges into operations, the outcome is relatively flat funding over the lifetime of the SKA of about €100 million annually. Building in stages is powerful risk mitigation.

Preparation of a detailed operational model is beyond the scope of this document. However, it is essential that an operational plan be implemented in conjunction with the telescope design and construction plans. During the design and development phase of the SKA, the engineering group should include one or more scientists with responsibility for planning for the SKA operations

Since the SKA operations budget will be dominated by personnel compensation, it is particularly important to minimize the number of personnel involved. This means a high level of automation and reliability. It is assumed that operations will be controlled from a center near the core of the SKA, although not precisely co-located to avoid self generated RFI. There are also likely to be regional and lower tier centers with a variety of functions, managed and funded at some level by the participating nations (Section 2.4). External contracts and the use of commercially developed major components are also likely.

Some of the key questions to be addressed for the operational model are:

- What is meant by fully operational?
- What fraction of equipment failure is acceptable before the science goals are compromised?
- What are the single-point failures that must be repaired and what are the redundancies, which could mean that some components might never be fixed or made operational?

If there are particular systems for which repairs will be highly labor intensive and expensive, then it may be better to build and install an adequate number of spare systems which are then allowed to degrade over the lifetime of the telescope. However, we note that some operating experience must be acquired to determine the “mean-time between failures” (MTBF) of various subsystems and to adjust the provision of spares accordingly. The operations model must include strategies for

graceful failures and this will impact on how contracts for different aspects of operations are let.

The principal sections of the operating model relate to:

- Routine operation of the telescope, scheduling and conducting observations (simultaneous and multi-mode) and some level of data management and analysis, such as RFI mitigation. Telescope operations will be conducted from a central control area. Actual observing will be under computer control at each station, based on a script generated either by the scientist or by SKA operations personnel. Individual antenna elements will feed their performance information into a central station hub, so that the main duty of the telescope operator will be to monitor the station rather than the individual antennas. Based on current telescope experience, reliable and continuous operations will require five support staff to perform tasks including detailed supervision of the inner part of the array, monitoring of remote stations and sub-arrays, management of dynamic scheduling, data quality control, management of the correlator and handling of callouts for equipment repair. Additionally, a number of technicians would need to be available for on-call support, although not necessarily scheduled for a shift. The number will depend on the reliability of the various elements. To cover these positions continuously (including illness and vacations), six people are required for each position, making a total of 30 operators. This does not include any assistance with preparing observing schedules or data analysis beyond routine quality control, either during or post-observations.
- Repairs and maintenance, including emergency repairs. Routine maintenance will be split between major maintenance and overhaul which will occur only infrequently and minor activities which will occur more frequently. The major maintenance activities will require that systems be routinely taken out of operation, say about every three years. Thus, at any time some fraction of the SKA will be out of operation. If possible, an initial complement of replacement spares should be included as part of the construction costs and where practical such spares could be installed from the beginning. The operating budget needs to include a continual complement of spares at a level commensurate with their failure rate. With a large number of individual antenna elements, it will not be possible nor necessary to keep all antennas operating at all times. Outages are inevitable and not unacceptable provided that the non-operating antennas are suitably distributed. For the SKA, we need to consider two levels of operational failure: individual antenna breakdowns and outages of entire stations. We anticipate that no emergency visits be made to a station for outages of individual antennas. Instead, stations may be visited for after-hours repair only if, for example, 10% to 20% of the stations in a given annulus from the center (0–35 km, 35–100 km, etc.) are down. For present purposes, we define a station to be down if at least 10% to 20% of its individual antennas are not operational.
- Development of new systems and upgrades of instrumentation and computing facilities. If we assume that about half of the SKA cost will be in electronics instrumentation, then a 2 to 3 percent annual investment in renewal and

upgrades will mean complete replacement in 15 to 20 years. However, computing hardware will need to be replaced more frequently. A detailed study of the trade-offs between using general purpose computers and dedicated hardware is needed as well as a comparison between cost effective consumer personal computers (PCs) with their rapid obsolescence and the unavailability of customized replacement parts.

- Support for users at all stages of observing will be needed. Regional Centers could provide assistance with proposal preparation, data processing, and analysis. Such Centers are being set up for the Atacama Large Millimetre Array (ALMA), and already exist for several space astronomy missions, which could serve as models for the SKA. A policy to provide financial support to users to carry out the science on major new facilities should be evaluated and could include support for senior investigators, post-doctoral fellows, students, travel, computing and publication costs, as well as any special instrumentation needed for particular experiments. It is assumed that user support should be at a level of 2 to 3 percent per year of construction costs and this expense might form part of the contribution of the participating countries. SKA science support would provide expert assistance to users as well as make a contribution to system developments and upgrades.

2.2 Interface with construction – commissioning timeline: The design and construction of the SKA must consider operational costs from the very beginning. Short-term savings which lead to increased maintenance and renewal costs are to be avoided. In order to insure smooth transition between construction, commissioning, and operations, it will be important to identify an Operations Manager with full responsibility at an early stage. Construction engineers will undertake initial technical commissioning of the instruments but the operations staff will have responsibility for scientific commissioning and early “shared-risk” observing.

Consideration of the budgets for current telescopes produce a robust estimate for operational costs of 10% of capital, split between routine operations and maintenance (5%), renewal and upgrades (2-3 %) and user support (2-3%). We would like to break this paradigm, but we have no concrete suggestions at present as to how we can achieve any substantial reductions in operating costs. We note that since we are already expecting to use an innovative design to reduce the cost per square meter by a large factor, meeting an annual operating budget of 10% of construction costs will involve a major challenge. As a cost model for the different designs and their construction is developed, it is essential we have a parallel cost model for operations to track how different modifications and developments influence operations. Part of this calculation is the inclusion of a realistic lifetime for the telescope and system lifecycles for the major sub-components. For example, both hardware and software computing facilities, the correlators, the antennas, feeds and beamformers all need evaluation.

2.3 Tendering and contracts – outsourcing: Some tasks such as simple infrastructure maintenance and non-essential technical tasks can be outsourced. However, it has been the normal practice at most radio observatories to use in-house staff for all essential operations support following the perception that this is the most cost effective practice. In contrast, Japanese optical and radio observatories and the

NASA Deep Space Network essentially outsource all operations and this strategy is known to be very expensive. However, for both Japan and NASA it satisfies a mandate to support industry.

Some of the reasons underlying the perception that outsourcing is too expensive and not the preferred model, include the profit motive of external companies, international union and labor regulations, excessive servicing and redundancy if the external contractor has been accustomed to maintaining close to 100% fault-free operations or if there is no understanding of the overall telescope and its purpose, which may occur if a contractor is servicing only a limited area of the instrument. Of special concern is the critical response for single-point failures and the most complex customized systems such as the correlator.

A sociological issue arises if we opt for substantial out-sourcing as this is likely to favor local companies in the host country – local employees will be easier to recruit. If in-house operations are preferred, then it should be easier to include staff from the international partners under the participation agreements that will be in place. Use of local employees to service the remote stations and to provide infrastructure support may be the most efficient option.

2.4 Regional and National Centers – Centers of Excellence: Centers of Excellence and national SKA Science Centers will be needed both to maintain the involvement of national partners and to support local scientists with the preparation and review of proposals, the analysis of data, training of students, and possibly the allocation of financial support. A strong scientific staff including postdoctoral fellows and students will be important. However, for this model to work, there needs to be adequate support for travel and an accepted freedom-of-movement culture.

The multi tier support system being adopted for the computing associated with the Large Hadron Collider may be attractive for the SKA. This model¹ envisages multiple “Tiers” of support and interaction arranged around the telescope. The MONARC model speaks in terms of computing resources but the model is applicable to all support activities associated with a large, highly computerized system such as the SKA. Adapting their model directly to the SKA gives the following layering:

- Tier-0: the SKA, incorporating also a Tier-1 facility
- Tier-1: large national scale Regional Centre, expensive, multi-service
- Tier-2: smaller Regional Center, less expensive, mostly dedicated to analysis
- Tier-3: institute-level center satellites of Tier-2 and/or Tier-1
- Tier-4: individual researchers

This model requires that the SKA Center, being both Tier 0 and Tier 1, has the capabilities and responsibility to support most planned observations and certain standard services such as a priori calibration and editing for all observations. The other Tier 1 Centers will provide similar capabilities on a regional basis. Together with some Tier 2 Centers, they might provide expertise and facilities dedicated to particular uses of the telescope, such as pulsar observations, polarization studies,

¹ MONARC - http://barone.home.cern.ch/barone/monarc/RCAr_model_chitecture.html

archive mining, SETI, etc. Clearly there are many policy questions to be answered if there is to be ample flexibility in arranging nations' commitment and funding. For example, are the Tier-1 facilities autonomous or do they have deliverables back to the telescope or to scientists from other regions? We note that the delegation of some tasks to regional centers can add layers of complexity and increase the risk of miscommunication and increased costs. The authority and responsibility for decision-making must be understood from the start. It is not yet clear what operational capability and responsibility should reside at the Tier 1 Centers. It is possible that the responsibility for SKA operations could be handed on from center to center, so that each one only has primary control during the normal work day period, thus eliminating the need for shift work.

To provide focus, it may be attractive to have particular centers dedicated to specific scientific themes or Key Science projects such as strong gravity tests or large-scale HI surveys. This cuts across technical boundaries as there might be experts at each Center to develop specific tools and equipment for that Key Project. The role of these Science Centers might include assistance with proposal preparation, data analysis, development of specific software tools, archive hosting, etc, but it is not clear how many of these activities should also be undertaken in more general national facilities where there is a common language and culture.

2.5 Human resources: It is often desirable to retain the most productive members of the construction staff as the basis for the operations staff. This may be problematic if the construction of electronics and mechanical systems is undertaken in a range of countries and/or widely spread geographical locations. However, this option should be included in the timelines and work packages if at all feasible. It may be that regional centers which are set up for the construction phase, subsequently morph into operational, science and development centers. Training of staff should be progressive as natural attrition occurs. Employment opportunities in the host country will involve government approval of visas for the influx of skilled international staff.

The central complex of antennas extending to several tens of kilometers, the central processing facility, and approximately one hundred stations within a few hundred km of the SKA center can be operated and maintained by central staff. These personnel will provide administrative services, buildings and grounds maintenance, antenna, cryogenic, electronic and computer maintenance, as well as some instrument development. From current paradigms, we estimate these tasks will require about 150 engineers and technicians, 30 research scientists, 10 technology experts to support computing hardware, 20 for systems work and communications, 10 to support monitor and control systems, and another 10 technicians for data management. Some of these individuals might be located at one of three or four regional service centers to provide support for the antennas located beyond about 100 km, but less than a few hundred kilometres from the central facility. For each of the more remote stations two to four technicians will be needed. Assuming an average of three full-time equivalent (FTE) staff per extended region, of the order of a few hundred kilometers beyond 350 km from the center, this means an additional 240 individuals. Administrative personnel probably should number at least 10% of the total number of employees and would include a fiscal/business division with procurement, management, human resources, and secretarial support. This gives a total personnel requirement of about 520 to fully staff the SKA Facility, using the assumptions listed.

2.6 Innovation: Provision of large national radio astronomy facilities has often had the effect of reducing the number of smaller observatories, including university-operated telescopes (Section 3.4). While understandable in terms of rationalizing total spending, this trend is unfortunate since these smaller facilities allow the next generation of astronomers and engineers to gain vital hands-on experience. The freedom to experiment and learn from mistakes, without risking an entire observing program or the disruption of a large array, is essential in fostering innovation in radio astronomy. Maintaining a strong student population is of crucial importance to the SKA and it will be a challenge to retain opportunities for students to undertake original research projects in the context of a multi-national instrument with several large key projects dominating the scene.

Allocating a fraction of SKA resources to training and innovation may partially offset any negative effects of domination from a large facility. Apart from exploiting observing parallelism in the SKA (via wide fields-of-view or multiple beams), allowing innovators access to parallel branches of central data streams in order to commission and test new instruments is likely to be beneficial. On the engineering side, training and innovation programs associated with Tier-0 or Tier-1 facilities (Section 2.4) and limited access to the SKA (perhaps via one of several beams), or one of its end-to-end engineering test beds, would build hands-on capability. An intriguing possibility is to use SKA technology to build simple low-cost radio telescopes, making small-scale radio astronomy ubiquitous at the senior high school and undergraduate levels. Coupled with the more advanced training initiatives mentioned, this could actually stimulate the supply of nascent radio astronomers.

3. Governance, management and oversight

3.1 International aspects of the SKA: The international nature of the SKA project will impact both the construction and operational phase. Unlike most other international scientific programs, “the SKA was born international.” However, issues of “real politik” must be recognized and there is already a divergence in activities supporting national and regional interests.

Based on what we have learned from other international scientific projects, it will be important that the operation of the SKA Observatory (SKAO) be vested in a Director and SKA staff who have the responsibility and authority normally associated with such activities, such as control over the operations budget, with a minimum of political restriction on how and where the funds are spent. While the funding agencies will certainly want to influence the organizational and financial structure, it is important that their impact be limited to broad fiscal and management oversight. This implies a well defined legal structure to the SKA that is internationally recognized for its governance and management excellence.

3.2 Funding models for contracts – construction and operations: Some system of *juste retour*, returns in proportion to contribution, will be desirable for the contracts for construction and operations. However, this is not the strategy to be adopted for allocating observing time, where an “open skies” policy has been agreed because it will maximize scientific returns (Section 6.2). Proportional returns must be balanced

against the potential for inflated contracts from partners who expect their entitlement. We envision that industries from contributing countries may wish to become incorporated in the host country to facilitate undertaking contracts with the SKAO. In addition, we note the apparently successful system in use at STScI, whereby ESA provides a fraction of the STScI staff who continue to be employed by ESA and thus continue their social benefits, but serve under the STScI Director. Such a system may be appropriate for the SKA construction as well as operations, with both staffing and instrument work packages distributed among the partners to minimize the flow of funds among the partners.

A particular burden will fall on the host country to facilitate the movement of people and materials across international boundaries. There will be issues concerning (1) security, (2) employment conditions and staff safety, salaries, and work policies, (3) visas, immigration and residency factors, (4) legal and contractual conditions, (5) insurance, medical cover and work permits, (6) customs and tax regulations, among many others, all of which must be negotiated in advance of any commitment to a specific site or host country.

3.3 Co-ordination of geographically dispersed activities: The SKA is an international collaboration and as we converge on the design it is essential that the most cost effective and efficient facilities are selected for the different components and sub-systems. The coordination of the geographic spread of activities will be a challenge as we must ensure that there is no unnecessary duplication of effort and that the different components of the SKA interface effectively. Other international astronomy projects have had to deal with multiple lines of control and funding coupled with proportional distribution of contracts according to participation, and this has created great difficulties, misunderstandings, and cost escalation.

3.4 Impact on national radio astronomy programs: The construction and operation of the SKA with its unprecedented large budget for a ground-based telescope will threaten the existence of current national and institutional facilities. This will be particularly unfortunate if the SKA is dedicated to a narrow range of scientific programs, such as an HI machine to study the Dark Energy problem. Building, funding and operating the SKA could absorb almost all the available resources supporting the current centimeter-decimeter facilities. However, we must keep the next generation of astronomers and instrument builders engaged and able to undertake their research and training programs. How we manage this transition and the full impact of having one SKA to serve the whole radio astronomy community is unclear. In other words, how do we keep the SKA from strangling individual countries' national and university facilities while being realistic about the total resources available?

3.5 Inclusion of new partners: Consideration needs to be given from the start to the mechanism for allowing new partners to join the SKA project after construction is completed. What will be the buy-in requirements? If an "open skies" policy for access is adopted, will there be any incentive for new partners? What will be the incentive for any government to provide major financial support if their astronomers may participate openly in major science programs. The important benefits from Research and Development outcomes, and the opportunities for industry and future commercialism must be publicized effectively.

4. Technical issues

4.1 Maintenance and repair schedules and resources: Extrapolation of the operating costs of current large radio astronomy facilities leads to a prohibitively expensive funding profile for the SKA. Hence, the operating culture will need to substantially differ from that of the current paradigm. Reliance on cryogenics to achieve good sensitivity must be reduced, and the number and complexity of moving parts should be kept to a minimum. The balance between the added expense required to provide self-checking and self-repairing capabilities, and the benefits from increased reliability and the operational cost savings achieved with the implementation of these systems, needs to be carefully evaluated.

The primary driver for the location of array elements will be to optimize the characteristics of surface brightness sensitivity, angular resolution, and achievable dynamic range compatible with any geographical and political constraints. It will also be important to consider access to roads, fiber and support staff to minimize their impact on operating costs.

The use of cryogenic cooling for receiver front ends may place a significant burden on the cost of maintenance and repair, although there is evidence that new cryogenic systems may be available with a MTBF of the order of a million hours. Ultimately, it may be cost effective to accept the somewhat degraded sensitivity of un-cooled systems, especially at the longer wavelengths, and compensate with increased collecting area.

At most levels, hardware failures will result in a temporary degradation in performance, until the fault is located and repaired or replaced. However, more serious will be major component failures, which may cause additional irreparable harm. An example is the damage that may be inflicted on sensitive front ends and delicate feed coupling structures due to the thermal cycling associated with cryogenic faults.

Very widespread expensive damage can be inflicted by lightning strikes and associated power surges. While good power and other infrastructure design practice can minimize the vulnerability to the secondary effects of a lightning strike, it is notoriously difficult to provide protection against direct, or nearby strikes, on antennas employing sensitive receivers. The potential for widespread disaster is so great in the SKA core, or within a station, that receiving systems will most likely need to make a small tradeoff in sensitivity in order to be more robust against static damage. This is especially true for aperture array patches which will employ many thousands of receivers. While lightning is the obvious threat, other charge separation processes such as dust or rain induced static, have already caused problems at several of the dry candidate SKA sites.

4.2 Renewal and upgrading of equipment: A normal part of any regular maintenance program is to provide an adequate supply of spare parts in quantities related to their MTBF. For an instrument like the SKA, with many components duplicated many hundreds of thousands of times, isolating and repairing failures will be time-

consuming and consequently expensive. At the very least, specialized test equipment must be developed as part of the construction phase to minimize the debugging phase. Spare parts, even spare antenna elements could be built into the system to automatically replace failed components as a result of continual automated monitoring of performance. Repairs could then be made when convenient, such as at the time of planned upgrades.

It must be understood that once a delay in repair time is established it will be very difficult to recover, so a policy for what is to be repaired and what requires urgent response is needed. Planned obsolescence also needs an accurate assessment of the failure rates for critical items as well as a clear listing of redundant components. Fully flexible scheduling, which deals with system failures will require a considerable investment in software. A reference software architecture that minimizes operational costs and allows maintenance and renewal of systems is essential. A change in the maintenance culture could see spare parts installed from the beginning with failures left progressively unattended.

In order to meet the construction cost targets, it is likely that the SKA will need to exploit the consumer and commercial markets for antennas, receiver and signal processing components, and certainly computers. To protect against the obsolescence of consumer items, careful attention will be needed to insure that the construction budget provide a lifetime supply of spare parts, and that software be portable to the next generation of computing hardware.

4.3 Computing systems and software: Software plays a key role in all modern radio telescopes, especially those using synthesis imaging, with its rich potential for innovation. Unfortunately, most large software projects do not satisfactorily meet their completion times, budget limits or full functionality. While radio astronomy is not unique in the late delivery of software, the size and complexity of the SKA project, and its dependence from the start on innovative algorithms and digital controls, highlights the need for software to receive careful attention. It will be very important to write realistic, affordable, and stringent software specifications for the initial operation of the SKA and to ensure that there are the same controls for the delivery of software as for the delivery of hardware items. Software development is very expensive and must be included in planning the construction budget in the same way as the development of antennas, receivers, and signal processing systems. It must not be deferred to the operations phase.

On the other hand, as with hardware, there will be a continual need for the development of new algorithms and new software to implement them, which means there must be adequate provision in the operations budget to support this activity. At the end of construction, there should be a fully implemented and tested software package capable of supporting the planned scientific goals. The planning and development of that package, including the specification, needs to be undertaken with the recognition that good software is probably best tested and released incrementally in the same way as is done with hardware development, including upgrades and renewal. If data distribution and processing is carried out by regional centers, a coordinated approach to the production and upkeep of archives and adequate funding for maintenance is required. Properly resourced software support and ongoing development is crucial.

Because of the immense investment required for the development of SKA software, and the need for continual upgrading and maintenance, the lifetime of SKA software will often exceed that of the computing hardware on which it is implemented. Therefore, it will be important that the resources needed to make software easily portable to new computing hardware be made available at the time of development, and not deferred until later. However, this must be balanced with the potential need to use specialized hardware to meet some of the more demanding performance specifications, which may require the use of that hardware for specific processing for some (small) number of years.

4.4 Power requirements: The power needed to operate the SKA is likely to cost several tens of millions of Euros per year, with the operation of an advanced correlator and movable antenna elements being the prime drivers. There are two principal options for power generation – public facilities or self-generation; hybrids of the two may also be considered. The risk of reliance on external power sources is that there is no long-term guarantee of supply or cost. However, self-generation also adds a cost vulnerability if fuel prices rise. It will be useful to investigate alternative power sources such as wind, geothermal, or solar power. Power requirements for the construction phase and for long-term operations may well be substantially different so sourcing the peak and continuous needs for both stages must be determined. Since provision for peak loads is expensive, the SKA should be designed as far as possible to minimize the peak-to-average power consumption ratio. Hybrid systems appear to be the most cost effective way to deal with the storage problem common to wind and solar power energy generation.

As the technology for integrated circuits advances, the power requirements and internal leakages also grow. Now that power demands and not Moore's law are starting to limit technology developments a cap on the available power, whether directly or via the capacity to actively extract heat, may ultimately be an SKA performance limiter.

4.5 Communication and fiber links: Operation of the SKA will demand very high bit rate data transmission over long distances. While commercial units and leased fibers may be more cost effective than installing fiber networks, consideration of life-cycle costs may indicate that a more prudent decision is for the fiber to be owned and installed by the SKAO as part of the construction phase. This is almost certainly true for distances up to about 100 km. However, the possibility of government subsidies, commercial use of SKAO owned fibers, or access to existing dark fiber may determine whether it is more cost effective to lease or to purchase fiber on longer links. Security and control on leased fibers are additional factors that need to be considered.

4.6 Infrastructure support: As well as the major costs of power and communications links, there will be substantial infrastructure costs such as the provision of staff housing and transport options, construction of the central and regional Centers. Good coordination of infrastructure between the regional Centers will be needed to ensure transparent interfaces and interoperability while preserving national characteristics.

4.7 Operational health and safety issues: There will be national regulations in the host country that will impact on the schedules and strategies for all operations, particularly maintenance and repairs. Standards for construction of various components and subsystems will also need to be uniform or at least conform to a set of guidelines.

4.8 RFI environment: Much effort has been expended on the mitigation of unwelcome radio frequency interference. However, the first line of attack against RFI will be to prevent it. Technical personnel will be required to track down sources of interference, both internally and externally generated, as well as to maintain the political protection of the SKA site from space and ground-based transmissions. Monitoring of RFI at the site must start early, before construction and continue throughout the lifetime of the project in order to follow the progression of both internally and externally generated interference.

5. Operations and maintenance costs

5.1 Cost-performance trade-off: An important challenge for SKA designers involves the balance between construction and life-cycle costs, and system performance. Traditionally, radio astronomers have worked with a small number of extremely advanced (“bleeding edge”) subsystems in order to extract the maximum performance from individual system blocks. The SKA will incorporate through economic necessity a good deal of consumer-grade technology. In this domain, the “bleeding edge” is an uncomfortable place to sit. For example, the failure rate of the newest field programmable gate arrays (FPGAs) can be worse than for previous-generation devices. On the other hand, the hugely improved integration levels of the newest devices are very attractive in minimizing such characteristics as printed circuit board area and connector count, which is highly desirable in the quest for greater reliability. It is probable that for most parts of the SKA the optimal path will be to use technology near, but not at, the extreme cutting edge, with the recognition that this position will involve a greater number of hopefully easily maintainable units to achieve the required performance. Nevertheless, with the huge number of subsystems involved, errors in misjudging this balance point could have a significant effect on the operations and maintenance budget, which is a potent argument for a thorough study of this issue (incorporating experience from the pathfinder projects) as part of the design process.

5.2 Comparison with existing arrays: The SKA Reference Design has some commonality with existing arrays, particularly the requirement to operate reliably a number of antennas in relatively harsh environments. Although the optimization between capital and maintenance costs will almost certainly be different for the SKA, we at least understand the essential nature of the operations and maintenance challenges. Notwithstanding this, the SKA presents us with many unknowns. For example, the power – performance trade-off in phased array feeds, and the possible static immunity versus sensitivity compromise in aperture arrays, are as yet unaddressed. We expect the maintenance issues associated with the newer technology will only be determined by the SKA pathfinders. With infrastructure, it is clear that items such as energy, long-distance communications (particularly on leased links subject to rate variations outside the control of the SKAO) and roads will be

significant drains on the operations and maintenance budget. However, it is unknown at present which components will dominate costs. Hence, infrastructure cost modeling will be an important part of the SKA Phase 1 Design Studies.

5.3 First estimate of operating cost: Despite the uncertainty in the operating and maintenance model, there have been first order estimates of the SKA annual operating budget. One analysis, first presented in the 2002 USA Large Number Small Diameter (LNSD) White Paper, is based on the NRAO operating model and an assumed personnel requirement. As described in Section 2.5, the total annual personnel count amounts to over 500 FTEs, dominated by more than 400 engineering and technical positions. Significantly, about 240 engineering FTEs were devoted to “distributed” engineering positions, underlining the need to design for reliable, unattended, operation of SKA stations. The resultant annual operations cost is then approximately €70 million, including some component for re-investment and user support.

To first order this estimate is probably realistic, although there are many unknowns in the SKA model. One obvious twist is the possible need to service at least two antenna types, which may be constructed by different organizations with different development, documentation and maintenance training structures. An important part of the SKA Phase 1 Design Study must be a comprehensive operating model, including cost, for the whole telescope.

5.4 Maintenance models: It is clear that SKA sub-systems must be comprehensively modular: any sensible “on-the-ground” maintenance model overwhelmingly invokes module-style servicing, even to the replacement of a complete phased array feed, aperture array tile or similar assembly. Modularity is also central to the idea of an upgradeable instrument, although for a fully electronic telescope, upgrades will necessarily involve bulk replacement of far more of the actual system. There are key challenges for options such as the splitting of phased array feeds and associated electronics, which could be essential for minimizing subsequent servicing and replacement costs.

An important maintenance decision will center on whether faulty modules are repaired at all, either by SKA or contract technicians. With a consumer electronics model for much of the telescope one might expect a higher fraction of throw-away modules than for present-day instruments, where the throw-away culture is rare. Nevertheless, the SKA will have many high value sub-systems which will need to be repaired. Optimization of the discard-repair balance will be important and will involve analysis of the full cost of items such as labor, specialized test jigs and similar equipment, complete assembled modules and major components.

With the large scale and long lifetime of the SKA, two aspects of the maintenance model may be somewhat different from models used for current instruments. Firstly, it is probably impossible to retain a complete spares kit for the whole lifetime of the SKA. While antenna and related mechanical systems may require an extensive spares stock, many electronic sub-systems will be most effectively maintained in the long-term by the substitution of redesigned modules. This makes the definition of specifications and interfaces especially critical. While the initial need for precise definition is obvious, the inevitable push for increased performance with time will

make rigorous, continuous management of the telescope hardware and software configuration essential. It is conceivable that derivatives of current management tools used in the telecommunications and related industries may be attractive to SKA engineers tackling the configuration management challenge.

A second distinguishing aspect of the maintenance model is the likely use of many spare sub-systems as “hot” spares. These sub-systems will be online and augmenting the base telescope capacity from the initial installation. The effect of failures will then be minimized by automatic resource switching until performance has degraded gracefully below specification, or until blanket replacement of the whole sub-system occurs during upgrade phases.

5.5 Legal considerations: Whichever model is adopted, maintenance and spares provision for the SKA will represent a significant fraction of the total capital and operations cost of the instrument. Therefore, contractual arrangements for these aspects must be folded into a wider procurement model negotiated between regional partners, government agencies and industry stakeholders. The procurement model will need to address issues such as long-term asset ownership, transfer of intellectual property rights from the developers to maintainers of both hardware and software systems, depreciation allowances and differing procurement rules across various regional trade jurisdictions and countries.

5.6 Funding cycle considerations: One of the challenges for funding the SKA comes from the diversity of partners in the consortium. This diversity is one of the strengths of the project, but also a complication. With both direct government and other funding there will be a variety of policies and regulations, variable funding cycles and opportunities. Even within individual countries there are different funding regimes, with some cycles following the calendar year and some locked into an economic fiscal year. Regional Centers will probably be funded and staffed locally and hence, governed by local fiscal policy. There is also the seasonal mismatch between the Northern and Southern hemispheres which will have some influence on when staff are available and contracts negotiated. As part of the governance of the SKA, a model needs to be agreed for the process of gathering financial contributions and the subsequent auditing and reporting requirements, without adding enormously to the administrative load on the overall project. While funding cycles will be an issue during construction and commissioning, it is particularly important that sufficient advance thinking is done to guarantee a viable continuing model for operations, including maintenance and re-investment.

6. Impact of science drivers and outcomes

6.1 Science innovation & impact: The scientific return from the SKA is likely to be optimized if maximum flexibility can be maintained. However, flexibility incurs a cost and it will be necessary to balance the provision of multiple options and cost. As with other large research facilities, it may be difficult to facilitate the exploration of novel ideas with the SKA. It will be important to allow scientists to exploit fully the capabilities of the SKA through non-routine experiments, unrestricted by complex administration, while also adopting the efficient but naturally conservative use of

Time Assignment Committees (TACs) to optimize observing schedules. In particular, there will be reduced operation costs associated with long term legacy type programs.

6.2 User support and observing modes – multiple users: The distribution of user support between the central Control Center and the Regional Centers needs to be carefully evaluated. A hierarchical structure of support may be optimal with expert users directly interfacing to the telescope. As the extent of the software complexity becomes apparent it may be that we cannot support a wide variety of observing modes or data analysis options, which is the model adopted for the Low Frequency Radio Array (LOFAR). In any case, it seems that most observers are conservative in their approach and content to follow a prescribed recipe for analysis.

Simultaneous observations for independent but compatible projects will increase the analysis time and complexity. There needs to be appropriate support for calibration, quality control, data reduction and archiving. How this is implemented needs consideration and sophisticated modeling to optimize scientific outcomes.

6.3 Observing time allocation: The SKAO will need to give careful consideration to the balance between Principal Investigator-style programs and legacy or project driven programs as this will largely determine the culture of the SKAO, which will have an impact on operations. The best science will come from an “open skies” policy where observing time is allocated by peer review of the scientific merit and qualifications of the proposers, regardless of national or institutional affiliation. We assume that nations investing in the SKA will reap benefits through investment activities and through their involvement in the future management of the facility, rather than through the allocation of observing quotas.

A single TAC is probably the most efficient model, but multiple TACs may allow more innovation. There is currently great variability in the granting of Director’s Discretionary time at observatories, and a policy will need to be established for such use. Finally, we note the need to allow for unexpected Targets of Opportunity (TOO) where a rapid response occurs and an existing scheduled program is overridden. It will be important to have a carefully constructed policy for TOOs to avoid general and unreasonable claim-staking of a particular class of object. The SKAO will also need to identify a proprietary period after which all data enter the public domain.

The most efficient use of observing time will result from a system of dynamic or queue scheduling where the user submits in advance an approved observing program, which is then run by the SKA operators when weather conditions such as wind, temperature and precipitation are optimal. However, dynamic observing does not allow for user intervention and may restrict the opportunity for serendipitous discovery. It also puts additional burden on the support staff to develop the necessary software to implement and support dynamic observing. If the SKA operates only at relatively long wavelengths, the need for dynamic scheduling may be reduced but it may still be the most efficient observing mode. In general, it is assumed that real-time user interaction will be done remotely, but not often.

6.4 Data products and archives: The SKAO will be able to take advantage of developments in the Virtual Observatory capabilities to facilitate the dissemination of archived data. The developments include data formats, software interfaces and tools,

and secure access provisions. To allow the greatest use by the astronomical community, the archive should include both the raw data and multiple levels of processed data with a priori calibration, on-line editing, pipelined data editing and calibration, and pipelined images. In addition, observers should be encouraged to submit their own images to the archive as a further level of data product. It will be necessary to recalculate archive images periodically as data processing processes improve. Data products should therefore be given a version number, and knowledge of the processing algorithms applied should be maintained by the SKAO. Because of the vast datasets being produced, the question of what actually constitutes “raw data” and whether some compression and initial processing is inevitable before archiving must be addressed.

Data management has traditionally been a challenging area for radio observatories. Considering its sensitivity, field of view, spectroscopic, temporal, and spatial resolution, the SKA will provide data at an unprecedented rate. It has been estimated that ALMA will have a data rate of at least 100 Tbytes per year, which may be compared with the total of 20 Tbytes that have been accumulated to date for the Hubble Space Telescope. Data management at the SKA may be organized into areas of system administration, software development and maintenance and archival activities, including preparation and support of user friendly data archives.

Careful consideration will be needed to decide to what extent the data processing facilities should be located at the Control Center or be distributed among different geographic regions. Another question is how often does data processing equipment need to be upgraded? This could be a very large expense in the operations budget.

There has been a lot of discussion about the need to include end-to-end software capability in the design of new radio telescopes in order to allow the broadest user base to exploit the powerful new facilities. A decision must be made on how to distribute the limited available resources between developing powerful algorithms for the most challenging projects such as high dynamic range imaging for the study of the faintest radio sources, and providing convenient access to pre-made images by non expert users. The most effective approach may be to support limited modes of observing and processing for the majority of observing programs, with expert data analysts to assist the more challenging projects, rather than produce a fully automated flexible pipeline. However, since there are likely to be simultaneous multiple observing programs, the scheduling, completion of the observations, calibration, primary analysis and data quality control will need to be automated and under expert control. Customized operations and software might be supported by the regional centers especially if there is a potential need from a broad group of users. Delegating responsibility to regional centers does not reduce the overall costs, but may be attractive nevertheless.

6.5 Student training and participation: The future of the SKAO and its long-term scientific significant output will reside in the hands of the young researchers we train over the next decade. Managing the engagement of the next generation astronomers with the large Key Science projects that will probably dominate the SKA schedule, at least in the initial stages, will be complicated. The role and very existence of present national and university facilities is undetermined (Section 3.4).

7. Summary and Recommendations

- Lifetime operating costs of the SKA are likely to exceed the construction costs unless the facility is shut down after undertaking a number of specific science projects. Limiting the range of science goals could reduce the lifetime operating costs. Conversely, unrestrained ambitions for science outcomes and program flexibility could have a strong impact on operations and construction costs. The division of operating costs between the SKA Observatory (SKAO) and national and regional entities must be understood and agreed upon in advance, and provision must be made for the long term support of SKA operations.
- From existing telescopes a robust estimate for the operational budget is about 10% of capital, which is split between routine operations and maintenance costs of about 5% per year, plus a few percent per year for renewal and upgrades, and a few percent per year for user support. There is agreement that we should strive to break this paradigm, but it isn't clear if this is realistic, especially considering the intention that the construction cost should be significantly less than the traditional cost per square meter of collecting area, by an order of magnitude or more.
- It is clear that a new paradigm for operating as well as constructing the SKA must be developed if we are to achieve our scientific goals with the available resources. We do not yet have a model that can satisfy this requirement.
- It may be appropriate to think of the SKA as an evolving facility with no strict demarcation between construction and operations, but with a fixed annual budget which extends over its complete lifetime.
- It will be important to minimize the operations budget, which may require additional expenditure at construction time, but which will minimize the whole life-cycle costs. It is important that an operational plan and cost model be developed in conjunction with the telescope design, construction planning and funding profile.
- An “open skies” policy with full access to all qualified scientists independent of institutional, national, or regional affiliation will give the best scientific returns from the SKA. Allocation of observing time should be based on peer reviewed proposals without regard to quotas, financial, or in-kind contributions to the construction or operation of the facility. There must be adequate provision for Target of Opportunity observations and rules for proprietary access to data must be established.
- It will be important to have a well designed governance and management structure with full authority and responsibility invested in the SKA Director and with an Observatory staff responsible for day-to-day science operations, hardware repair and maintenance, software development and maintenance, and the development of new instrumentation. Care must be taken to mitigate the effects of salary inequities among staff from the different partner countries.

We suggest that the flow of funds among the partner countries can be minimized for construction as well as operations if there is a fair distribution of both staff and instrument work packages among the partners.

- The availability of multiple independently-steerable beams and the wide range of spatial frequencies planned in the design suggest that efficient use of the facility will require support of multiple simultaneous users.
- We anticipate that most, if not all observing programs will be carried out under dynamic or at least queue scheduling with little real time interaction with the scientist. Full time operational support by about five people will be required. Experience at current facilities suggests that to fill each position 24/7 will require 6 FTE's. Some sort of "turno" system may facilitate the implementation of support personnel. Alternatively, cycling of telescope control across regional centers may limit the support staff needed.
- Careful thought must be given to the split between proposal driven and project driven programs, since this will have a great impact on both the culture and costs of operations. Dynamically scheduled observing will be the most efficient but must be balanced with some opportunity for interactive observing, which may be required for some of the more innovative programs. The construction program must include the development of adequate software to support dynamic observing. To first order we anticipate that approximately an equal amount of time will be spent on short programs as on long programs. This means that the number distribution of observing time periods should be roughly proportional to the inverse of observing time granted.
- Particular attention must be paid if we are to minimize the maintenance cost of the SKA by providing sufficient redundancy, self checking, and self repair, particularly of single point failures. Cryogenics maintenance is particularly labor intensive and expensive, so the need for cryogenic cooling should be reduced where possible and used only to optimize performance. Careful attention to failure rates and repair/replacement protocols will be needed. Mechanical components such as servo systems can be major service items, so the number of "moving parts" should be minimized.
- Power will be a major cost driver, so every effort should be made to reduce power consumption, especially in the design of the correlator, antennas, and cryogenics. In the future it is likely that signal and data processing capability will be limited by power consumption and dissipation, rather than by clock rates and memory size. Beam-forming and other signal processing devices will also be cost drivers. Careful analysis of the comparative expense of achieving accurate pointing control by mechanical means or by signal processing needs to be evaluated, with consideration that the relative costs may change with time over the lifetime of the SKA.
- It seems likely that the basic operation of the SKA will involve more than 500 FTEs and that the cost will exceed 100 million Euros per year, including funds from national resources that will be needed to support the research enterprise that uses the facility.

- Current international telescope projects are moving in the direction of greatly increased accountability and increased administrative costs and burdens and this trend may be hard, if not impossible, to reverse. The ALMA project is but one of several examples.
- It is important that the source of operating funds for the SKA be identified even prior to construction. This should include pre-operations funding at a low level when the first components are build on site, which will ramp up to full operations support when construction is completed. Adequate resources and preparation for commissioning during the latter phase of construction and the early stages of operations will be critical, and it will be important to allocate sufficient funds from the construction project for the commissioning phase of the SKA.
- Equally important will be recurrent funds available to upgrade the SKA, taking advantage of new opportunities at a level of at least a few percent of the construction cost per year. This is critical for an instrument such as the SKA which will depend heavily on signal processing and is therefore very sensitive to the impact of Moore's law. This dependence suggests that the signal processing instrumentation will need to be completely replaced on a time scale of a decade at most.
- National or regional science centers might be established around specific science themes such as the designated key projects, but they will also need to provide more general support for all scientists from the region.
- A balanced appraisal must be made comparing the benefits gained from using consumer technology to minimize construction costs with the possible negative consequences for operations because of rapid obsolescence and unavailability of components.
- Consideration of life-cycle costs suggests that it may be prudent to install propriety fiber to service a large part of the array during construction phase rather than lease fiber for many years.
- Adequate software capable of supporting all the major scientific programs of the SKA must be produced during the construction phase and not be deferred to the operational stage, as has often become the current practice. Nevertheless, the capabilities of the SKA must grow with time to address questions that were not even conceived when the design was frozen. This will require continuing expensive software development.
- Where appropriate, SKA software which was developed at great effort and expense over many years should be easily portable to new generations of general purpose hardware. There may also be the possibility of importing software from other projects such as LOFAR but this will require early collaboration to optimize portability. It is important that software developers keep lifetime costs in mind and not take short cuts while under pressure to deliver, which might have negative long term implications.

- Detailed investigation of the SKA configurations will be needed to minimize the impact of failed elements or stations. The larger the number of antennas and the greater the redundancy, the more robust the array will be to random element failures, which means longer repair times may be tolerated. A detailed cost-benefit analysis of the impact of failed array elements based on the probability distribution of failed elements plus any remaining single-point failure items will be needed to optimize performance and to minimize operating costs. Over-design and/or redundancy may reduce or even eliminate the need for “emergency” replacement or repair. Scheduling procedures will need to be developed to mitigate the impact of system failures that may affect only specific types of observations.
- The nature of the SKAO products, which might include raw data, reduced images produced via end-to-end processing and a managed archive, must be defined in advance with adequate interfaces to the International Virtual Observatory. We suggest that while many routine observing programs can be supported with pipelined data reduction, this may not be realistic for the most challenging and innovative programs, which will best be supported by experienced skilled data analysts rather than the implementation of complex, expensive and difficult to define generic software.
- Much will be learned from the operation of SKA Phase I, which is planned to be built on the approved site.

Acknowledgements: We have benefited from talking with many people including Peggy Perley and Jim Ulvestad who have shared their experience with us in operating large radio telescope systems.

Appendix

Lessons Learned from Other Large Radio Telescope Projects

While we anticipate that the operation of the SKA will need to be substantially different from that of existing radio observatories, it is instructive to identify the main cost drivers for existing facilities so that we may identify new approaches to reduce the costs without significant impact on the performance and reliability.

VLA/EVLA/VLBA: The EVLA+VLBA is conceptually similar in the range of baselines, frequency, geographical dispersion and infrastructure needs etc that will be required for the SKA.. Some maintenance, such as HVAC is handled by local contract. There has been a tendency to underestimate the software support required. Software maintenance is more expensive if the development is contracted out. It is important to have a critical mass of competent people including scientists, engineers, and programmers on site to learn from each other. Systems must be immune to catastrophic damage from component failure. Construction funds must pay for an adequate supply of spares from the beginning. Management must be sensitive to local living conditions. New radio telescopes must recognize the strong desire to provide calibrated images to the broad multi-wavelength astronomical community via the IVO. Dynamic scheduling, robust and portable post-processing software including default pipelines, easy access to reliable archived data, and high level user support are all needed to optimize the scientific returns from the investment.

ATNF: The ATNF is about 10% of the scope of the EVLA+VLBA. ATNF staff consists of 21 FTEs which is relatively the same proportion as for the EVLA+VLBA. How will we operate the SKA with proportionally fewer people? The ATNF operating budget is about 15% cryogenics, 8% utilities, and 2% for improvements. Experience has been that after some period of time, it is difficult to find spare integrated circuit (IC) components, so new hardware must be bought or constructed, which has software implications. Software must be portable and the construction budget must provide a lifetime supply of spare ICs.

ALMA: ALMA will support a diverse community so there is pressure for ease of use by novices and experts alike. The issue of an “open skies” policy is not resolved. The different partners may have different policies. There will be both national TACS and a central TAC. It is important to consider the operations plan during the construction phase. ALMA will consider the commissioning team as a “user.” The ALMA Project Scientist will be in charge of commissioning. Operations will be under the control of a Director of Operations assisted by Heads of Administration, Science, and Technical Services. It is expected that only ¼ of the operating budget will be spent in Chile with 150 people on site. The goal is to minimize the number of on-site personnel. Scientists will do research in Santiago when they are not at the site. The ALMA experience suggests that it is important to start planning early for operations, with authority delegated to the Director, rather than under the control of a Board whose members may have conflicts of interest. Any resulting delays in decision making may have a big cost impact. Due to construction budget pressures, there has been a tendency to defer spares acquisition to the operations phase of the project.

A major cost driver for SKA operations will be the provision of power. Experience with commercial power and self generated power suggest that either may be risky. On one hand, Bolivia cancelled a project which ALMA was depending on for power, while the increased price of diesel fuel for generators threatens CARMA operations.

ATA: The operating philosophy will be to minimize the use of consumables such as electricity and oil, replace units in the field and repair only during only normal working hours. Most instruments will be removable from an antenna in 5 min. The ATA will operate in the “expert observer” mode and will provide for at least two simultaneous observing programs, one for SETI and one for astronomy, limited to the same pointing direction.

LOFAR: LOFAR is not yet operational, but there will be a lot to learn from the operational model adopted. In particular, the LOFAR facility will pioneer the use of multiple simultaneous observers, each with access to an independently steerable beam.