Power Considerations for the Square Kilometre Array (SKA) Radio Telescope

Peter J Hall

International Centre for Radio Astronomy Research (Curtin Institute of Radio Astronomy), GPO Box U1987, Perth, Western Australia 6845.
(email: peter.hall@icrar.org)

Abstract

The Square Kilometre Array (SKA) will be the world's most sensitive radio telescope and is expected to be fully operational at frequencies below 10 GHz by 2023. The SKA will extend over more than 3000 km but over half its collecting area will be located at one of two remote, radio-quiet sites in either Australia or South Africa. The instrument will collect and process vast amounts of information, and the provision of reliable, affordable electrical power over an expected 30-50 year operational lifetime is a major challenge. This paper outlines some of the main issues and mentions a few exemplar innovations in the area of SKA power.

1. Introduction

The Square Kilometre Array project [1] has undergone a guided process of specification and system design refinement since 2001, when various end-to-end concept documents for the telescope were produced. One of the most significant shifts in the project was in 2004-5, when the scientific attractiveness of wide fields-of-view (tens or hundreds of deg^2) was recognized widely. Incorporated rapidly into instrument design goals, this aspiration increased greatly the amount of information to be collected, distributed and processed relative to earlier thinking. In turn, the information expansion significantly increased the projected power demand of the telescope, to the point where the instrument's scientific capability will be limited by the availability and price of power at the sites being considered, either in Australia or Southern Africa. In the last five years power provision for the €1.5 billion (2007 currency) SKA has become a widely discussed challenge, partly as a result of a Power Investigation Task Force (PITF) which was formed to raise awareness of the issue and, in particular, to forge links between the electronics-dominated SKA engineering community and strategic thinkers in the power industry. There is now recognition that power supply to the SKA is almost entirely an industry issue, while demand minimization falls largely in the purview of the SKA Program Development Office (SPDO) and its likely successor, the SKA Project Office (SPO). There is, however, a strong interplay between supply and demand considerations, not least because total-cost-of-ownership (including capital and operating costs) comparisons must include not only supply-side variants but also SKA technology implementation alternatives. The recent SKA Project Execution Plan identifies four major stages relevant to power provision; the first of these, to be undertaken in 2011-12, invokes greatly increased links between the SPDO and power industry analysts with a view to setting out key cost comparisons. Importantly, this linkage will also quantify important SKA-specific technical requirements, including specifications for radio-quiet supply and distribution systems.

2. SKA Power Estimates

Accurate estimation of SKA power demand will be an outcome of detailed design and prototyping undertaken in the 2013-15 project execution phase. However, simplified calculations [e.g., 2, 3] have produced demand estimates of the order of 100 MW, split roughly equally between the array proper and the associated high performance computer centre. Importantly, these estimates assume significant innovation and consequential demand minimization in e.g., receptors, digital signal processing, computing and environmental conditioning (notably cooling). Without such innovation, estimates of hundreds of mega-watts flow from an extrapolation of present-day radio telescopes. Fortunately, many of the innovations needed for the SKA are directly relevant to demand minimization in other domains, notably data intensive web, archiving and related applications. Notwithstanding the need to capitalize on commercial development, SKA proponents are active in pioneering a number of demand minimization strategies. The motivation is clear enough: with the widely-accepted figure of 10% of capital cost as a first estimate of the annual SKA operating budget, a 100 MW demand absorbs two-thirds of the budget (assuming €0.12 per kWh power cost). Savings in energy cost translate very directly into increased capacity

978-1-4244-6051-9/11/$26.00 ©2011 IEEE

for user support, instrumentation investment, etc. In an absolute sense, the long projected lifetime of the SKA (30-50 years) means that operating costs dominate initial capital costs by a large margin and, with power a large fraction of the operating budget, savings in power demand are highly significant over the life of the project.

From a power provider's viewpoint the SKA challenge is similar in some respects to that of powering a rural population centre, its outlying districts and remote settlements. Broadly, these regimes correspond of course to the SKA core(s), a region of perhaps 50 km diameter in which power is reticulated, and remote stations which require stand-alone power. Reticulation in the central regions is straightforward in principle but requires careful design to account for the rather complex SKA antenna placement pattern, and the need for an upgradeable system when antenna numbers increase between SKA Phase 1 (10% sensitivity) and Phase 2 (full telescope sensitivity below 10 GHz). Power provision at isolated, remote stations is a similar challenge to that of supplying outlying, off-grid rural settlements.

It is useful to consider the distribution of electrical load in more detail. In current thinking the SKA exa-FLOP computing centre will be located in a populated area, for reasons which include ready access to a ~50MW supply and the need to minimize radio frequency interference (RFI) at the telescope central site. A large amount of central site digital signal processing is still required, however, with one estimate [4] placing the associated demand at 12 MW. In terms of distributed loads, it is useful to look at receptor technologies separately and, for simplicity, to examine the "baseline" SKA system design. Here, dishes with wideband single-pixel feeds require of order 10 kW each, including cryogenic refrigeration (for at least low-noise amplifiers) and environmental cooling. Sparse aperture array stations may require about 80 kW each, including environmental cooling for station electronics. A basic SKA consisting of 2000 dishes and 250 sparse aperture array stations then requires 40 MW. Within this representative array 40 remote stations (groups of ten or so dishes) may be needed, each station needing a 200 kW supply. Thus, of the 40 MW total, 8 MW of stand-alone remote area power supply may be required, although current SKA configuration studies aim to place as many stations as possible near grid supplies.

3. Special Requirements for SKA Power Systems

While the SKA power challenge is in large measure a generic one, the radio telescope application imposes some special considerations. Perhaps the most inconvenient from the perspectives of both the power providers and the SKA system designers is the need for low RFI. Undoubtedly, many otherwise-attractive, cost-effective and efficient options will be rejected because of the difficulty of meeting stringent RFI and derived electromagnetic compatibility (EMC) requirements. While final SKA specifications are not yet available, representative limits at candidate central sites in Australia and South Africa set received power density limits of order -230 dBm Hz^{-1} at 1 GHz in the array central region. Intra-system EMC standards are still under development for the SKA but it is clear that efficient distribution of clean power within electronics racks (perhaps via direct-current links) will be as challenging as the supply of mains power to those installations. For the latter, infrastructure plans envisage an underground reticulation network within the SKA central region, giving a degree of attenuation to radiated high-frequency signals carried on power conductors. Notwithstanding care in choosing and installing SKA power systems, it is likely that at least some static inverters and associated control circuitry will be a part of any contemporary solution, requiring the power industry to examine closely power-electronics topologies and installation methodologies.

Load profile and supply reliability specifications for the SKA will be important inputs to industry designs for the power infrastructure. The SKA will operate essentially continuously giving power providers little opportunity to use off-peak storage and similar strategies. On the other hand, the SKA has rather more flexibility to define its operating modes than many commercial installations. A simple example might be a decision to limit dish slew rates, or slew concurrently only selected sub-arrays of antennas. In terms of power supply reliability, it is highly likely that an affordable SKA will require a hierarchical specification. Some parts of the telescope, such as the central signal processing and off-site computing, will require very high integrity supplies if electronics lifetime and science output is not to be compromised. Other parts, such as an individual group of antennas, can tolerate less supply reliability, providing this is factored into the SKA system design from the outset. In practice, operational strategies and reliability specifications will enter the power system design as cost trade-offs via, for example, the need to cope with large peak/average loads or to provide redundant infrastructure.

Protection, monitoring and control are factors requiring additional thought in the SKA context. With the SKA approaching an all-electronic telescope, it is potentially vulnerable to the propagation of damaging surges, whether from errant power systems or natural phenomena, such as lightning. Many parts of the SKA will be isolated galvanically from each other via optical fibre links but an inadequately designed power reticulation system would be an effective medium for transferring catastrophic damage. Apart from provision of active protection strategies, the tension between good signal grounding practice, compliance with electrical safety grounding standards and best-practice lightning grounding will have to be resolved in an arid environment, with potentially high soil resistivity. In a related consideration, SKA power monitoring and control will be relatively sophisticated, if only because of the geographical diversity of the array. Fortunately, the nature of the instrument provides excellent signal connectivity on which to superpose relatively low-bandwidth power system data traffic. More challenging perhaps is the implementation of back-up systems for antennas forced by communications failure to operate in stand-alone mode, possibly in the face of a depleting remote area power supply.

4. Solutions and Pathfinders

Much of the SKA system design being undertaken and prototyped relies on innovation in commercial and consumer electronics. For example, Liebsch [5] outlines a few of the relevant demand minimization technologies likely to be applicable in computing and digital signal processing sub-systems. These include dynamic variation of parameters such as clock speed and supply voltage and the use of power efficient, but less flexible, programmable circuitry in association with general purpose central processor units (CPUs). Another promising path, designed to mitigate the energy use inherent in transporting very large volumes of data, involves dynamically reconfiguring processor engines based on e.g., field programmable gate arrays (FPGAs). Still other solutions may use structured application specific integrated circuits (ASICs) to obtain near-ultimate power efficiency while preserving a degree of design flexibility. In broad terms, it is the inverse relationship between flexibility and power efficiency (operations per watt) that will force SKA designers to formulate carefully a hierarchical flexibility specification, sub-system by sub-system.

SKA pathfinder telescopes, particularly precursors such as MeerKAT and the Australian SKA Pathfinder (ASKAP) built on candidate SKA sites, are developing important innovations in power systems, as well as providing invaluable site-specific insights for infrastructure deployment. Optimal power delivery strategy differs between the two candidate central sites, if only because of the difference in proximity of grid power. At the MeerKAT site South African project engineers, working with ESKOM, are demonstrating a radio-quiet transmission line connection to the national grid; this is based on a carefully over-specified 33 kV feeder followed by underground reticulation from a point 5 km from the MeerKAT instrument. In Australia, CSIRO and Horizon Power are demonstrating a fossil-fuel and solar hybrid, off-grid, power plant for ASKAP; the intention is to increase progressively the renewable energy fraction of the plant's output. Both the MeerKAT and ASKAP supplies are of order 1 MW and will provide experience for scaling to SKA Phase 1 and beyond. In new, low-cost aperture array designs being developed as part of the European-led SKA Aperture Array Verification Program (AAVP) the feasibility of solar-powered active antennas, with an associated super-capacitor storage, is being investigated [6], both as a cost-effective alternative to reticulated solutions and as a way of galvanically isolating highly distributed receiving elements. Cooling for electronic systems is a major concern at candidate SKA sites, and ASKAP and MeerKAT are studying interesting new approaches to this challenge. For example, MeerKAT is demonstrating evaporative pre-cooling to increase significantly the coefficient of performance of some of its refrigerative coolers. An elegant electronics cooling approach using chilled, reticulated water is also being trialed. In Western Australia, CSIRO and its partners are investigating a novel geothermal cooling solution for the Pawsey Centre for SKA High Performance Computing, an approach likely to reduce greatly Centre operating cost.

5. Renewable Energy and Site Selection

The hope has long been that the SKA project can be an exemplar in terms of its reliance on renewable energy, both from the perspective of minimizing the instrument's carbon footprint and, possibly, insulating it from the inevitable increase in fossil-fuel power costs. Renewable power could be delivered directly to the SKA central site from a fairly local (low-RFI) generating facility or, alternatively, could be "wheeled" via a grid which incorporates renewable power components. Both candidate SKA sites are, in fact, exceptionally well placed to use solar-based generating solutions, provided cost and storage issues are adequately addressed. Early consideration [e.g.,7] of solar technologies, including relatively new concentrating solar voltaic systems, indicate that such solutions may be feasible, given a willingness to invest in a suitable-scale development program and, most probably, to subsidize the unit cost of SKA renewable power, at least for the period of the instrument's operational life in which fossil-fuel alternatives are priced relatively attractively. The key question is where the investment and subsidy might come from. As part of the SKA site selection process proposals for infrastructure provision and operation will be developed by Australia and South Africa. There is also strong interest in renewable SKA power in other parts of the world, especially Europe [8]. It could be that a trans-national effort to provide green power would yield excellent technical results, assuming sufficient political will to go down this path.

6. Conclusion

Provision of reliable power is essential for the scientific effectiveness of a radio telescope and was a challenge in the early days of instruments such as the Very Large Array (VLA) and Australia Telescope Compact Array (ATCA), even at relatively well developed rural sites. With the construction of ASKAP, MeerKAT and the Atacama Large Millimetre Array (ALMA), the SKA community is learning from efforts at providing affordable, reliable power at much more remote and inhospitable sites. Many demand minimization solutions will come from innovation in commercial electronics but others will flow from advances within the SKA pathfinders and allied facilities. Careful specification of SKA operational requirements (which flow, in large measure, from science priorities) will be especially important to SKA designers in terms of choosing systems with the greatest possible energy efficiency.

7. Acknowledgements

I am indebted to my co-chair (Dr Bernie Fanaroff), fellow members of the SKA Power Investigation Task Force (PITF), and industry and other PITF associates for insights into the SKA power challenge.

8. References

1. P. J. Hall, R. T. Schilizzi, P. E. F. Dewdney and T. J. W. Lazio, "The Square Kilometre Array (SKA) Radio Telescope: Progress and Technical Directions, " *The Radio Science Bulletin,* No. 326, Sept. 2008, pp. 4-19.

2. P. J. Hall, "Power Investigation Task Force - Introduction," Presented at PITF meeting, Feb. 2009, Cape Town. (Available: http://www.skatelescope.org/pages/PITF_meeting_17Feb2009/hall_intro.pdf)

3. A. J. Faulkner, "Aperture Arrays for the SKA: The SKADS White Paper", Presented at SKA2010, Mar. 2010, Manchester. (Available:http://www.skatelescope.org/indico/getFile.py/access?contribId=41 &sessionId=13&resId=0&materialId=0&confId=2)

4. P. E. Dewdney, "Overview of SKA Power Considerations," Presented at PITF meeting, Oct. 2009, Manchester. (Available: http://www.skatelescope.org/pages/PITF_23Oct2009/Presentations/PITF-091023_2.0_Dewdney.pdf)

5. T. Liebsch, "Power Trends Discussion," Presented at PITF meeting, Feb. 2009, Cape Town. (Available: http://www.skatelescope.org/pages/PITF_meeting_17Feb2009/liebsch.pdf)

6. P. J. Hall, "ICRAR Contributions to the Aperture Array Verification Program," *Widefield Science and Technology for the SKA: SKADS Conference 2009*, Limelette, Belgium, 2009, pp. 405-406.

7. I. Burdon and P. Willson, "International SKA Project: Review of Generation Technologies, Trends and Costs," Presented at PITF meeting, Oct. 2009, Manchester. (Available: http://www.skatelescope.org/pages/PITF_23Oct2009/Presentations/PITF-091023_6.0_Burdon.pdf)

8. European Cooperation in Science and Technology (COST), "Benefits of Research Infrastructures Beyond Science (The Example of the SKA)," *COST Strategic Workshop Final Report,* Rome, Italy, 2009. (Available: http://www.cost.esf.org/download/Final%20report%20COST%20SKA)