Report of the SKA Siting Options Working Group

Executive Summary

The Site Options Working Group (SOWG) was established by the SKA Organisation Board of Directors to investigate whether viable dual-site implementation options exist for the SKA (Phase 1 and 2) and, if possible, present a preferred option. The Board required that the preferred option (or options) should be capable of delivering the SKA science case while making best use of the existing investments and characteristics of both sites and being financially viable. The SOWG was not tasked, and did not consider, whether these options should be preferred to single-site implementations. The SOWG considered only scientific, technical, and programmatic issues, including cost and implementation risks.

SOWG members were appointed by SKA Organisation Members and have acted in delivering their recommendations in the best interests of the Project in the widest sense. In some cases, it is acknowledged that there are areas of national interest, these have been noted in the discussions. This report also notes where the conclusions are those of a majority rather than a full consensus view. The work of the SOWG builds on that of the SSAC and earlier work and discussions about implementation options that had been discussed and analysed by the SSEC and its working groups. The SOWG did not attempt to re-review or re-do the evaluation undertaken by the SSAC. The report of the SSAC, and analysis within it, did form part of the basis for the analysis of dual-site implementation options, though the SSAC itself did not explicitly consider the consequences of a dual-site implementation.

A dual-site implementation, in which the two sites host different technologies operating at different frequencies, appears capable of delivering the SKA Phase 2 science case. There is no identified scientific advantage of a dual-site implementation for the SKA Phase 2 over a single-site implementation. For SKA Phase 1, incorporating the MeerKAT and ASKAP Precursors and related infrastructure offers a distinct scientific advantage. The SOWG noted that the total investment in the precursors which would then become available to the SKA project was in excess of €300m. However, a dual-site implementation for either SKA Phase 1 or Phase 2 introduces additional capital costs, as well as additional operating costs and programmatic risk. An important aspect of a dual-site implementation is the introduction of the concept of the “SKA Observatory”—a single facility that may have multiple locations (e.g., RSA and/or ANZ and/or other locations for global or regional science centres).

A detailed assessment of the additional costs for a dual-site implementation for SKA Phase 2 could not be conducted. However, the dominant costs for each SKA element (low-frequency aperture array, mid-frequency dish array, and the advanced instrumentation program (AIP) technologies—mid-frequency aperture arrays, phased array feeds, wide-band-single-pixel feeds) and the directly associated infrastructure appear, to first order, independent of implementation model. The SOWG noted from the SSAC report the difference in cost of provision of power infrastructure on the two sites giving a “high” advantage to RSA and a “low” advantage for other infrastructure owing largely to the lower labour and material costs. The SOWG further noted from the SSAC report that with the exception of recurrent power costs the operation costs on the two sites were “virtually the same”. In addition to site-specific cost differential, there would be additional costs for infrastructure for a dual-site implementation compared to a single-site implementation, but these additional costs, while not insignificant, are likely to be less than the overall cost uncertainty for the SKA. The operational cost of a dual-site implementation would also be somewhat higher. The operational budget will be dominated by the required power provision which will be very comparable between a single- and dual-site implementation, but with costs dependent on the differential cost of power on each site.

The SOWG noted that the SSAC report also concluded an advantage for RSA in the four “Implementation Plans and Cost Factors”; while ANZ were favoured in all of the six “Other selection factors.” A majority of the SOWG noted that a dual-site implementation for an “SKA Observatory” could potentially offer risk mitigation to the Project across all ten of these additional factors.

A majority of the SOWG identified the following preferred option for a dual-site implementation of SKA Phase 2.

	SKA Element	**Location**
Low-frequency aperture array	SKA2_Low	ANZ
Mid-frequency dish array	SKA2_Mid_Dish	RSA
Mid-frequency aperture array	SKA2_AIP_AA	RSA or ANZ

The minority opinion was that it was premature to identify a preferred dual-site option for SKA Phase 2. This preferred option is consistent with the detailed analysis presented in the SSAC report and, in particular, matches the mid-frequency capability to the Southern Africa site where the tropospheric stability is better. A majority of the SOWG expressed a slight advantage for co-locating SKA2_AIP_AA and SKA2_Mid_Dish based on technical considerations, including the potential to cross-correlate dishes and mid-frequency aperture arrays. However, retaining options for the AIP component of SKA Phase 2 at this stage offers advantages in terms of risk mitigation. None of these considerations constitutes an *a priori* reason to select a dual-site implementation in preference to a single-site implementation.

For SKA Phase 1, the formal inclusion of the Precursors into SKA Phase 1 together with full utilisation of site infrastructure leads to implementation options that offer a scientific advantage of enhanced sky surveys compared to the baseline SKA Phase 1 planning. The SOWG received from the Precursors an inventory of infrastructure (including the Precursor instruments) that is installed now or expected to be installed on completion of the Precursor programmes. The SOWG assumed that this infrastructure would be available in full for potential inclusion into the SKA Project; this assumption represents a programmatic risk.

The specific dual-site implementation options that appear viable for SKA Phase 1 are the following

		Option A	Option B
Dish Array	SKA1_Mid	RSA	RSA
Low Frequency AA	SKA1_Low	ANZ	RSA
Survey instrument	SKA1_AIP_Survey	ANZ	ANZ

A majority of the SOWG did not express a clear preference between these two options.

Both options deliver the core SKA Phase 1 science case (as described in SKA Memorandum #125 and quantified in the Design Reference Mission for SKA Phase 1), while providing enhanced science return via sky surveys. The notional SKA Phase 1 realizes a mid-frequency dish array using 250 antennas. In both options, MeerKAT and ASKAP would be incorporated into SKA Phase 1. MeerKAT will provide a high-sensitivity array with 25% of the sensitivity required for SKA Phase 1 in the mid-frequency range, and, if combined with 190 SKA Phase 1 antennas, an SKA Phase 1 dish array meeting the mid-frequency science requirements for sensitivity would be realized on the RSA site. ASKAP will provide a high survey speed array demonstrating the science return from the field of view expansion technology of phased array feeds (PAFs), and, if combined with up to 60 SKA Phase 1 antennas equipped with PAFs, an SKA Phase 1 survey capability meeting or exceeding the mid-frequency science requirements for survey speed would be realized on the Australian site. The inclusion of PAFs in Phase 1 SKA is consistent with the overall Phase 1 plan described in SKA Memorandum #125 and the engineering plan outlined in the SKA Project Execution Plan. In both options, the performance of PAFs at an acceptable cost would have to be demonstrated successfully at the review scheduled for AIP technologies in the current SKA Project timeline. In either option, a low frequency aperture array must be constructed in full as per the SKA Phase 1 baseline to deliver the crucial high-redshift hydrogen science as described in SKA Memorandum #125.

Option A provides full use of Precursor infrastructure, excellent use of the Precursors themselves, and a natural build-out option from SKA Phase 1 to Phase 2 if a dual-site implementation option for SKA Phase 2 were to be adopted. Option B provides an optimal use of Precursor infrastructure and antennas in RSA, less so in ANZ where data transport provision in particular is under-utilised, and potentially leads to SKA Phase 2 implementations in which there is redundancy of a subset of SKA Phase 1 hardware. A definitive conclusion regarding whether an adoption of a dual-site implementation option for SKA Phase 1 would imply a dual-site implementation option for SKA Phase 2 could not be reached by the Working Group.

Both options are likely to have an additional cost associated with them compared to the baseline SKA1 concept via the inclusion of PAFs; the estimated range is €9–15m, depending upon whether ASKAP dishes are retrofitted with latest PAF technology. This will be offset by possible savings of €7-8m realized by not equipping antennas at the Australian site with single-pixel feeds. The re-use of existing infrastructure at the sites is likely to reduce some of the cost and mitigate programmatic risks associated with a dual-site implementation for SKA Phase 1. The SOWG's best estimates of the additional cost for options A and B, including additional resourcing at the SKA office, are very approximately €21m - €39m and €17m - €29m respectively, both less than the overall cost uncertainty for SKA1 and of order 8% the SKA1 nominal budget. These figures are based on a range of assumptions described in more detail in Section 2.2 (p14-16). Additional annual operating costs per annum are likely to be dominated by power costs which are estimated to be in the range €5m - €7m for option A and €4m for option B.

Incorporation of the Precursors into the SKA Project increases the overall risk to the SKA Project. This risk can be mitigated to some extent if appropriate management arrangements are adopted early. The majority of the SOWG favour a model in which the Precursor projects are regarded on the same terms as other self-funded technical contributions to the SKA Project. The Precursor projects are formally incorporated within the SKA Organisation at a very early stage, preferably the start of the preconstruction phase, with the aim to fully integrate the Precursors (engineering, management and science) into the SKA Project by the time a positive funding decision for SKA Phase 1 is achieved.

1 Background, assumptions and approach

1.1 Introduction

The Site Options Working Group (SOWG) was established by the SKA Board of Directors to investigate whether viable dual-site implementation options exist for SKA (phase I and II) and if possible present a preferred dual-site option to the next members meeting subject to the following constraints:

- The preferred option should be capable of delivering the SKA science case
- The preferred option should make best use of the existing investments and characteristics of both sites.
- The additional costs compared with a single site option should be part of the group's consideration, and should be estimated
- The group should comment on implementation risks (and advantages) associated with this approach, and where possible identify mitigation strategies
- If the group cannot identify a viable option that meets these criteria then they should report that fact.

The work of the SOWG builds on that of the SSAC and earlier considerations of dual-site implementation options considered by the SKA Science and Engineering Committee (SSEC), the Science Working Group (SWG) and the previous SKA Director. Although the SOWG explicitly did not re-review or re-do the detailed evaluation undertaken by the SSAC, the detailed report and analysis by the SSAC of the site factors are used by the SOWG in the analysis of implementation options. In two areas the nature of the work being done by the SOWG necessitated consideration of issues outside the scope of the SSAC, but in areas related to the SSAC evaluation of implementation plans and costs: these were the location of the data processor, which was considered as an intrinsic part of the implementation optimisation, and updated information on array configuration.

The SOWG noted that in considering a dual-site implementation option for the SKA ultimately there are three constituencies who need to be considered:

- The global science community, either currently involved or those potentially involved in the future. An outcome from the SOWG has to provide a compelling scientific way forward.
- SKA Organisation Members, current Members (candidate sites and beyond) and potential future Members. Specifically for future Members, the project had to remain an attractive proposition for investment and engagement.
- The people managing the project, and the engineering community delivering the detailed design – in considering options, SOWG should consider that the project will be manageable and ultimately deliverable.

1.2 Summary of site-specific infrastructure

A critical input to the considerations of the SOWG was the re-use of infrastructure at the two sites which is expected to be in place at the end of construction of the precursor instruments ASKAP and MeerKAT. In addition, the SOWG explicitly considered options for the possible incorporation of ASKAP and MeerKAT into the SKA. The low-frequency precursors (PAPER and MWA) were not of sufficient size compared to the low-frequency component of SKA1 that they represented a significant opportunity advantage and therefore were not considered further.

The SOWG received from the precursors an inventory of infrastructure (including the precursor instruments) which is expected to be installed on completion of the precursor programmes. For the purpose of this analysis the SOWG assumed that this infrastructure would be available in full for potential inclusion into the SKA project; this assumption clearly represents a programmatic risk. However most of the infrastructure is in place in Australia, some is in place in South Africa and work is underway to deliver the remainder.

Summaries in tabular form of the infrastructure and its value to the project as presented to the SOWG are given in the appendix. Critical elements of the infrastructure were regarded as particularly important for assessing site implementation options. The following table provides a summary of these key elements:

		RSA	ANZ	Note
Roads	Euro/km	16000	16000	1
Power Installed	MVA	7 + 3.75	2.1	2
Installation 132 kV + per MVA	M€	34 4	N/A	3
Additional Power per MVA	M€ / MVA	N/A	3	4
Power recurrent cost per MW	M€/yr	5 - 10	12.5	5
Data lit	Gb/s	10	40	6
Additional data	k€/ 10 Gb/s	280	0	7
Control building		SKA1 ready	SKA1 ready	8
Accommodation		SKA ready	SKA ready	8
Computing capacity (Racks)		130	76 (+50)	9
A/T	m^2/K	260	65	10
Survey speed	$m^4/K^2/deg^2$	$5\ 10^4$	$1.3\ 10^5$	11

Notes

1 Cost of minor roads comparable on both sites therefore not a discriminator
2 Total power available at each site: for further calculations assume a power factor of unity hence MVA are equivalent to MW. For RSA 7 MVA is available from the grid with 3.75 MVA available from on-site generation.
3 Additional cost to upgrade to 132 keV power supply – only applicable in RSA and needed for SKA2 only. €34m is the cost of the 132 kV line and substation; €4m per 10 MVA 33 kV feeder line to the core and skirt regions.
4 Cost per MVA for increasing power capacity. Figures are for ANZ only and assume the availability of additional solar power units and appropriate storage. Increments of power at similar cost points may be available in RSA (design decision).
5 Recurrent cost of power quoted as MEuro/yr per MW. The figures are those used in the cost modelling of the SSAC which the SOWG has followed for comparison including the range the SSAC used for RSA.
6 Each site has installed fibre – this is the data rate capacity expected or in place now from each site.
7 Cost of increasing the data rate capacity from each site to Cape Town / Perth for 10 Gb/s. The capital cost for RSA reflects the need to purchase additional capacity; there are no equivalent costs in ANZ. In both cases additional components are needed which are taken to be the same cost on the two sites.
8 These elements are SKA ready on both sites.
9 Capacity to house computing equipment in pre-existing shielded buildings on site. In ANZ, 76 racks will be behind 190db (now measured) of shielding, an additional ~50 racks could be accommodated elsewhere in the building but with only 95db (now measured) of shielding. In RSA all 130 racks will be behind 80db of shielding (excluding shielding due to adjacent hill and earth surrounding the buried facility).
10 Sensitivity of precursor.
11 Survey speed of precursor at L-band, scales as λ^2

2 Implementation options

The SOWG holds the very strong view that maintaining the scientific vision of the SKA concept is of paramount importance: this underpinned all of the considerations of the SOWG. Furthermore the SOWG noted the conclusions from the SSAC report which are of particularly direct relevance to the SOWG's work namely that while the SSAC report recommended Southern Africa as the preferred site they also noted that either site could host the SKA and that the SSAC explicitly did not consider the consequences of a dual-site implementation.

In order to approach the consideration of dual-site options the SOWG adopted the concept of the "**SKA Observatory**"[1] – the SKA Observatory is a single facility which may utilise multiple locations (e.g. RSA and/or ANZ and/or other locations for global or regional science centres). Implementing the SKA as multiple independent organisations would not maintain the scientific vision of the project.

With these preliminary considerations a small number of options exist for SKA2:

Option: Single site

Option: Frequency split: e.g. put SKA-low on one site and SKA-mid on the other

Option: Functional split: telescopes optimized for different science programs, notably survey-optimized telescope and sensitivity-optimized telescope

The SOWG also adopted the following additional constraints on its considerations:

- All collector technologies will operate simultaneously within the facility (a formal science requirement) in both SKA1 and SKA2
- There should be no fundamental change to the SKA concept or the defined engineering plan as outlined in the PEP – this therefore defines conceptual SKA collector deployment blocks (see Appendix) as a mid-frequency dish array; a low-frequency aperture array and two AIP technologies of mid-frequency aperture arrays and dishes equipped with phased array feeds.
- The existence of the precursors (ASKAP and MeerKAT) and site infrastructure, which is either installed or planned to be installed at the completion of construction of these precursors, represented a possible opportunity advantage for SKA1 only: the scale of SKA2 greatly exceeding that of the precursors or SKA1. The low-frequency precursors (PAPER and MWA) were not of sufficient size compared to the low-frequency component of SKA1 that they represented a significant opportunity advantage for SKA1 and therefore were not considered further.

The SOWG concluded that the only possibly viable options for SKA2 which could deliver the science case were the single-site solution or dual-site solutions in which the split was on frequency/technology grounds. However the opportunity advantage of the existing precursors means that a functional split is also potentially viable for SKA1. The SOWG further concluded that

[1] The concept of an Observatory controlling or operating spatially distinct infrastructure in pursuit of an overall scientific mission or program has significant precedent within both astronomy and physics.

even if a viable dual-site implementation for SKA1 exists that does not mandate a dual-site implementation for SKA2.

The conceptual blocks that were considered therefore were:

○ SKA1_Low:	Low-frequency aperture array
○ SKA1_Mid:	Mid-frequency dish array with high sensitivity single pixel feeds
○ SKA1_AIP_Survey:	A component of SKA1 delivering high survey speed at mid frequencies based on PAFs (part of the AIP and noted for potential inclusion in SKA1)
○ SKA2_Low:	Low-frequency aperture array
○ SKA2_Mid_Dish:	Mid-frequency dish array with high sensitivity single pixel feeds, and possibly equipped with PAFs (part of AIP)
○ SKA_AIP_AA:	Mid-frequency Dense Aperture Array (part of AIP)

The AIP components are not guaranteed to be available but are subject to future engineering decisions based on cost/performance as detailed in the PEP. Only PAFs are considered in the SKA engineering plan for inclusion in SKA1. The SOWG note that during SKA1 a production-readiness mid-frequency aperture array demonstrator is part of the SKA engineering plan, but since this is not part of the science delivery plan for SKA1 the SOWG do not consider its siting in detail.

An SKA2 implementation option on a dual-site cannot split these core building blocks without compromising the SKA science case.

2.1 Dual-site implementation options for SKA2

The SOWG believes that dual-site implementation options for SKA2 which deliver the full SKA science case exist based on a frequency split of collector. By this we mean that the main elements of the SKA (mid-frequency dish array, low-frequency aperture array and mid-frequency aperture array and each having a separate core in the SKA conceptual design) form building blocks which could be split between sites, but that each building block cannot be split and still deliver SKA science. There is no identified *scientific advantage* of a dual-site implementation over a single site implementation for the SKA observatory. Furthermore there are cost and programmatic risks which we discuss in more detail that are associated with adopting a dual-site implementation.

It was not possible to determine meaningful differential costs for a dual-site implementation for SKA2: even an approximate estimate of differential costs would require a conceptual design for the dual-site implementation which was beyond the scope and time-constraint of the SOWG process. . The SOWG noted from the SSAC report the difference in cost of provision of power infrastructure on the two sites giving a “high” advantage to RSA and a “low” advantage for other infrastructure owing largely to the lower labour and material costs. The SOWG further noted from the SSAC report that with the exception of recurrent power costs the operation costs on the two sites were “virtually the same”.

The SOWG however note the following:

- Overall cost differentials between the sites are analysed in the SSAC report. The SOWG limited itself to considering only further differential costs resulting from not co-locating all aspects of the SKA on a single site.

- The manufacturing cost (excluding infrastructure costs) of each full SKA2 component (dishes, low-frequency aperture arrays and potentially mid-frequency aperture arrays) is assumed to be independent of site-implementation, just scaling as total amount of hardware and add no additional differential costs beyond the site-specific deployment costs addressed by the SSAC.
- There will be additional infrastructure requirements for utilising a dual- versus single-site. Some infrastructure costs are directly associated with a given component of the telescope (e.g. foundations, bunkers, local power reticulation, local data transport, and local roads). These elements of the infrastructure do not provide any differential costs between a single-site and a dual-site implementation. Other components of the infrastructure must be duplicated between a single- and dual-site implementation (e.g. access roads, airstrips (which already exist), on-site support). Power and data transport would need to be duplicated in a dual-site implementation, but the costs of this duplication are reduced to some (possibly considerable) extent due to splitting the required scale of the provision between the sites. The data processor will need to be duplicated, but the size of processing facility on each site is dependent on the detailed design. For only one component (SKA2_Mid_Dish) is infrastructure and power provision located beyond 200km from the core and this again reduces the cost of duplication. The cost of a dual-site implementation therefore has increased infrastructure costs compared to the single-site. Detailed costing is beyond the scope of the SOWG analysis, however the SOWG estimate that the differential cost of a dual-site implementation is likely to be a relatively small fraction of the total infrastructure costs and a small fraction of the overall SKA cost and current cost uncertainty. The detail of analysis the SOWG was able to perform however means that adopting a dual-site implementation for SKA2 would introduce some additional cost risk to the project.
- A dual-site implementation introduces an additional risk in limiting the scope for design optimisations for delivery of the maximum science within a specified cost envelope.
- Cost of operations will increase using a dual-site implementation for the observatory. Observatory staff will be needed on two sites, however this should be significantly less than a factor of two since there will be a technology split between the sites (expert engineering would therefore be site specific), and it would ***not be necessary*** to duplicate certain aspects of the observatory operations. Further as is clear from the SSAC report it is expected that the operations budget will be dominated by the required power provision which to first order is the same between a single or dual site implementation with costs dependent on the differential cost of power on each site

Taken together the SOWG concluded that there is an additional cost and cost risk of a dual-site implementation option, but the SOWG identified no specific issues that demonstrate at this time that a dual-site implementation was not viable on cost grounds.

The dual-site option does mitigate some project risks. The SSAC report analyses site-specific risks and below further programmatic risks are discussed in the context of a dual-site implementation. The SOWG note the following:

- By utilising a dual-site model for the SKA Observatory there is some mitigation of site-specific risks available to the project.

- The transition period from SKA1 to SKA2 will be substantial during which there will be significant build-out of collectors and infrastructure. This construction work will impact the SKA plan of a full observing programme with SKA1 during this phase. Utilisation of two sites during this period may provide operational options that mitigate the risks of building out and observing simultaneously on a single site.

The members of the SOWG disagreed on whether it was helpful to specify the details of dual-site SKA2 implementations at this stage. A minority of SOWG members argued that a corollary was not to specify a preferred full-SKA dual-site implementation option at this stage due to the level of project uncertainty.

Within the context of the SOWG's remit to investigate dual-site implementations and having identified such an implementation was possible while still delivering the SKA science case, the SOWG proceeded to identify a preferred option. The majority of SOWG members were of the opinion that there was sufficient information available to identify preferred options for a dual-site implementation of:

SKA2_low	ANZ
SKA2_Mid_Dish	RSA
SKA2_AIP_AA	RSA or ANZ

This option was preferred on the grounds that there is medium advantage as identified in the SSAC report for RSA in terms of the troposphere; however, the SSAC also recognised that a risk exists around the proposed GSM suppression technology in RSA. The mid-frequency AA for SKA2 could be sited either with the low-frequency AA or the Dish array. The SOWG concluded that there was no conclusive scientific or technical argument to determine this siting option, but a majority of the SOWG were in favour of co-locating the SKA2_Mid_Dish and SKA2_AIP_AA (independent of the location of SKA2_Mid_Dish), noting:

- Co-location of mid-frequency survey instrument component with mid-frequency dish array offered some small advantage for follow-up observations of fast transients
- Co-location maintained the option of cross-correlation between the elements
- A design option presented in Memo 100 for the AA system made use of the frequency overlap between AA-low and AA-mid to give good sensitivity in the 300-450 MHz regime, however this does not require co-location.

The SOWG noted that none of these considerations constitutes an a-priori reason to select a dual-site implementation in preference to a single site.

2.2 Dual-site implementation options for SKA1

For SKA1, the inclusion of precursors explicitly within SKA1 and the reuse of infrastructure detailed in Section 1.2 offer opportunities and potential scientific advantage over the baseline SKA1. The SOWG noted that the total investment in the precursors which would then become available to the SKA project was in excess of €300m.

Even if the SKA Observatory adopts a dual-site implementation for SKA1 this does not force a dual-site implementation for SKA2: this follows from the relative scale of SKA1 and SKA2.

A systematic analysis of SKA1 implementation options consistent with the boundary conditions adopted by the SOWG is given in the next section. The analysis and motivation for the final two options is discussed here.

As already noted the existing infrastructure and precursors are significant compared to SKA1 requirements. Their inclusion into an overall SKA observatory as part of phase 1 not only offers a scientific advantage, but also rationalises the overall structure of the SKA project, fully integrating precursors into the SKA programme. The SOWG noted that the continued existence of an independent precursor project in the long-term on one or both sites to be utilised for SKA development represents a significant project risk in itself.

The core elements for delivery of SKA1 science are a mid-frequency dish array with a sensitivity of $1000m^2/K$ and a low-frequency aperture array of sensitivity of $1500m^2/K$ at 131 MHz (see Appendix). The SOWG determined that these sensitivity requirements must be maintained in order to deliver the core science case for SKA1 of the evolution of hydrogen and the study of gravity via pulsars. For the dish array the SKA1 system design has 250 dishes each with a sensitivity of 4 m^2/K and 50 AA stations of 31 m^2/K sensitivity. Splitting the AA collector between sites would have a serious negative impact on science delivery and was therefore not considered by the SOWG. The cost of delivering SKA1 is as yet uncertain, and will be determined over the 4-year pre-construction phase. The SOWG felt it should not take irrevocable engineering decisions to fit within the nominal €350m cost-cap before the detailed design and costing have been completed.

Options for incorporating precursors

The precursors are both dish arrays with MeerKAT providing a projected sensitivity at L-band of 260 m^2/K and ASKAP having 65 m^2/K but additionally optimised for survey science with the inclusion of PAFs giving 36 beams on the sky.

The SKA1_mid sensitivity requirement for the dish array can be met by including the 64 dishes of MeerKAT in SKA1 combined with a further 190 SKA dishes. Both the total sensitivity and number of antennas (important for determining baseline coverage determined by the configuration) are achieved with this combination.

- Note that the A/T for a MeerKAT dish plus feed is almost identical to the SKA specification despite the smaller dish diameter due to greater cooling of the MeerKAT feed.

The SKA1_Low sensitivity requirement can only be met by constructing the full 50-station low-frequency aperture array: as noted above neither MWA nor PAPER provide significant A/T and they do not meet other SKA1 requirements. Some reuse of infrastructure, such as fibre infrastructure, may be possible.

An opportunity then exists to combine the remaining 60 SKA dishes either with MeerKAT or ASKAP. The SOWG compared these two options considering the case in which 60 SKA dishes on the ANZ site are equipped with Phased-Array Feeds (PAFs). The figure of merit used for the comparison was the survey speed (SS) between the two implementations (the assumed performance figures are given in the appendix).

SKA1_Mid	N_{dish} = 254	A/T = 1000 m^2K^{-1}	SS (1GHz) = 10^6 $m^4K^{-2}deg^2$
SKA1_AIP_Survey	N_{dish} = 96	A/T = 275 m^2K^{-1}	SS (1GHz) = 2.7 10^6 $m^4K^{-2}deg^2$
SKA1_Mid (+)	N_{dish} = 310	A/T = 1250 m^2K^{-1}	SS (1GHz) = 1.5 10^6 $m^4K^{-2}deg^2$

The highest survey speed is achieved by incorporating ASKAP plus 60 SKA dishes into a survey component for SKA1 on the ANZ site. The engineering plan for the SKA requires a single dish design capable of holding single-pixel feeds and PAFs; the 250 dishes discussed in these options would be the same SKA dish.

This analysis assumes the availability of the PAF AIP receiver which is discussed below.

Cost of the AIP_Survey Component:
The AIP_Survey component is an additional component to the baseline for SKA1, although the inclusion of PAFs is considered in Memo 125 and the PEP. The inclusion of the SKA1_AIP_Survey component has the additional cost of outfitting 60 SKA dishes with PAFs. The estimated cost of a PAF plus beam former supplied by CSIRO is €150k giving a total additional cost of €9m. Retrofitting ASKAP with all updated PAFs in the SKA1 timescale, if deemed to be necessary will add an additional €5.4m. These costs would be offset by potential savings on 60 single-pixel SKA feeds, a conservative estimate for which is ~€6-8m. There will be some additional infrastructure costs for having a second dish core although this is minimised by keeping the number of SKA antennas deployed fixed. There would be an additional correlator required for the AIP_Survey component: the SOWG did not have available estimated costs for this correlator. The AIP_Survey component will also have a processing requirement. We return to these considerations below when discussing implementation options.

Power and Data:
The power requirements and data rates post correlator or beam former for each element are (see appendix)

Component	Power (MW)	Data rate (Peak) Tb/s post correlator
SKA1_Mid	1	2.5
SKA1_Low	3.7	0.4
SKA1_AIP_Survey	0.54	1
Correlator for SKA1_AIP_Survey	0.5	N/A
Science Data Processor (max)	4.7	

The power requirement for the SKA1_AIP_Survey correlator was taken to be comparable with that of the SKA1_Low correlator based on the similarity of input data rates. The power requirements for the science data processor are based on the maximum performance requirement which is associated with one of the core science experiments on the SKA1_Mid (dish) array – pulsar search.

The SOWG were unanimous of the view that the additional benefits of incorporating an SKA1_AIP_Survey element (made possible by incorporating ***both*** MeerKAT and ASKAP) make SKA1 dual-site implementations incorporating this option favoured in comparison to other implementation options (see Section 3). The added value of a high survey element has considerable scientific merit:

- It meets the requirements in the SKA1 DRM for high survey speed for the HI absorption experiment in the frequency range 0.8 – 1.6 GHz – the survey speed requirement is not met with the baseline SKA1 implementation.
- Having both SKA1_Mid and SKA1_AIP_Survey offers considerable additional scientific opportunity for SKA over and above the core science case while not compromising the latter.
- It offers the potential to incorporate both precursors bringing the full range of science programmes of SKA1 and the precursors under a single project.

Risks:
There is a cost risk for incorporating both precursors in a dual-site implementation.

- There is some mitigation of the cost risk by the optimal reuse of infrastructure however this is dependent on the detailed design.

There is a programmatic risk that the expected infrastructure and precursors (including performance) are not available on the timeline for SKA1

- This risk is mitigated by the precursor projects being ahead of the SKA project; the infrastructure is in place in Australia and being procured now in South Africa.
- Early inclusion of the precursors in the overall SKA organisation further mitigates this risk (see section 5).

Incorporation of MeerKAT as planned will use different feeds to SKA1 dishes with a different frequency coverage

- MeerKAT dishes could be fitted with SKA feeds at a cost of approximately €7-8m.

The survey element is conditional on the availability of the PAF. This is part of the AIP programme and the decision to incorporate PAFs within SKA1 will be made as part of the engineering process currently scheduled for 2014. The SOWG is unanimous of the view that the Project's engineering plan must be followed and therefore taking forward a plan including PAFs as an element represents a programmatic risk. This risk is mitigated by:

- The astronomical capability of PAFs is being demonstrated by ASKAP and APERTIF.
- The major risk is in the cost/performance of PAFs: the figures used here represent in the view of CSIRO a realistic and achievable cost/performance based on current measurements and engineering projections

Implementation options for SKA1

Two implementation options emerge for a dual-site SKA1 in which the two precursors are fully incorporated (see next section).

These are:

		Option A:	Option B:
SKA1_Mid	Dish Array	RSA	RSA
SKA1_Low	Low Freq AA	ANZ	RSA
SKA1_AIP_Survey	Survey instrument	ANZ	ANZ

The SOWG undertook a comparison of these two options against the range of criteria it had agreed to consider.

Science delivery and timeliness

- Both options enable the full SKA1 science case to be delivered and via the provision of the survey capability (made possible by inclusion of both precursors into the planning) give additional scientific performance to SKA1.

Infrastructure reuse

- Both options fully incorporate the precursors into SKA1
- Both options, to first order, would not require upgrading the data capacity to each site. This is achieved by siting the required processing at the Karoo site for RSA and utilising the additional existing data capacity in ANZ back to the Pawsey centre.
  - The estimated power requirement for the post-correlator processing is 4.7MW based on the peak data processing load from the pulsar-search experiment and is more likely to have a typical load of 3 MW (see appendix).
- Total on-site power requirements for each option are estimated to be:
  - Option A: RSA 5.7 MW ANZ 4.8 MW
  - Option B: RSA 9.4 MW ANZ 1.1 MW

  Option A would require an upgrade to the power availability at ANZ of approximately 2.5MW which is estimated to cost approximately €7.5m; option B provides optimal re-use of existing power infrastructure with no upgrade required on either site. The off-site processing facility in ANZ (Pawsey Centre in Perth) will draw power from the grid (see below).
- In option A the data capacity from the site in ANZ is fully utilised, but is under-utilised in option B.
- In both options the Pawsey centre in Perth provides additional computational capacity over a single site implementation option, but is likely to be more fully utilised in Option A, this is reflected in the different processor cost ranges in the table below.

Differential cost

- For both options there is a differential cost compared to baseline SKA1 planning of €9m for PAFs for the 60 SKA antennas or €14.4m if the incorporated ASKAP antennas are retrofitted with the same PAF design. Some of this estimated cost is possibly offset by the saving on 60 single-pixel feeds, which will be in the range of €7-8m, however a similar sum would be required to fit SKA1 feeds on the MeerKAT dishes if necessary. There is an additional cost for the correlator for the survey component. Both options require an additional processing facility. The additional processing centre requirement may be met by the availability of the Pawsey centre in ANZ.

- Option A has a differential cost compared to option B of approximately €7.5m for the additional on-site power at ANZ.
- There is an additional work load on the project office managing the detailed design and the necessary duplication of site-specific work. The SOWG estimate that at least a further 10 posts will be required in the SPO, at an estimated cost of €1m/yr.
- Option B only has longer baselines (out to 100km from the core) in RSA whereas option A has baselines of this length in both sites. There is some additional cost in implementing the longer baselines in two locations rather than on just one site, however the SOWG's analysis suggests this is a relatively small additional cost, estimated to be €7m, since the majority of the infrastructure cost is associated with the core out to approximately 10km for each of SKA1_Mid and SKA1_Low (see appendix). Only beyond this distance from the core do potential cost savings for co-location of SKA1_Mid and SKA1_Low become possible and little collector in SKA1 is at these or larger distances.

In summary estimated additional costs are:

Item	Option A	Option B	Notes
PAFs	€9m - €14.4m	€9m - €14.4m	Offset by saving on single-pixel feeds, of order €7-8m
Survey correlator	Not estimated	Not estimated	Same for both options
Power infrastructure	€7.5m	N/A	
Other infrastructure	€7m		Additional infrastructure for spiral arms, if 100 km baselines are required for SKA1_Low.
Additional costs in PEP	€4m	€4m	Additional staff in SPO
Processor	0 - €6m	€4 - €10m	Typical value for mid-size HPC facility. The range reflects unknown reuse of the Pawsey Centre.

Operating Costs

The SOWG estimated the recurrent cost of power provision for each option compared to the baseline SKA1. The cost model used by the SSAC was employed in which the cost per MW of power in ANZ was taken to be €1.25m/yr and for RSA a range was considered of €0.5-€1m/yr; this reflects the uncertainties in future cost of power provision as discussed in the SSAC report. The power consumption of the off-site processor facility (Pawsey Centre) required in either option in ANZ was taken to be 2MW reflecting the lower processing needs for the experiments linked to the SKA1 elements on the ANZ site in either option. The following table summarises the estimates of recurrent costs. For comparison siting the baseline SKA1 in ANZ or RSA would give a recurrent cost for power of €11.6m/yr or €4.7-9.4m/yr respectively.

Option	ANZ	SA	Total
SKA1 Baseline	€11.6 m/yr	€4.7 – 9.4 m/yr	
Option A	€8.4 m/yr	€2.8 – 5.6 m/yr	€11.3 – 14.1 m/yr
Option B	€3.8 m/yr	€4.7 – 9.4 m/yr	€8.5 – 13.2 m/yr

The differential operating cost per annum of the two options compared to the SKA1 Baseline is therefore in the range €5m - €7m for option A and €4m for option B. There will be additional operating costs associated with a dual site implementation in terms of additional staff however the SOWG note that the proponent sites both intend to operate precursor instruments as facilities into the SKA1 period.

Site Characteristics

- For RFI and EMI characteristics of both core sites (all that is utilised in SKA1), the conclusions of the SSAC report do not provide any useful information in discriminating between options A and B.
- For both options A and B SKA1_Mid is located at the better tropospheric site, however as presently planned the upper operating frequency of 3 GHz means that tropospheric stability is not a significant issue for SKA1.
- All other infrastructure is comparable on the two sites and sufficient for SKA1 requirements.

Implementation, management and other operational factors

- The SOWG noted that, to first order, option A represents a situation in which additional investment for SKA1 by the SKA project is comparable on both sites.
- The two options differ from an implementation standpoint when the build out to SKA2 is considered.
  - Option A has an implementation advantage if a dual-site implementation is adopted for SKA2 with SKA2_Low located in ANZ as discussed earlier.
  - Option B does not preclude the dual-site option for SKA2 considered earlier. The full low-frequency in SKA2 would need to be constructed on the ANZ site leaving the SKA1 low-frequency array infrastructure in RSA. This would incur the additional cost of replicating the SKA1 component, however there may be some modest advantage of building SKA2_Low on a virgin site and a scientific advantage of having two SKA_Low cores each of which has a filling factor (relative density of antenna elements) of approximately unity. There is also the risk of possible loss of staff expertise in changing location of this element from RSA to ANZ.
- Having a geographically extended SKA1 implementation on both sites will help maintain the strong RFI protection (see Section 4); this is an argument for Option A.
- As noted previously adopting either option for SKA1 does not preclude in any way a single site implementation option for SKA2.

Risks

A more detailed discussion of the overall risks associated with a dual-site implementation is presented in Section 4. Here the discussion is focussed on issues directly relevant to the SKA1 implementation.

- A dual-site solution for SKA1 has additional costs – these are reduced to some extent by reuse of existing known infrastructure as reflected in the cost of infrastructure provision already committed at both sites.
- The option to include a survey element requires the availability of PAFs which are part of the AIP; this risk has already been identified. If PAFs are not available, as discussed in the next section, the best scientific option would be to co-locate all dishes on the RSA site. SKA1_Low could still be located on the ANZ site, utilising the existing infrastructure, but without scientific advantage.
- The re-use of the precursors and infrastructure adds very significant risk to the project A mechanism should be found to formally incorporate the precursors into the overall project structure on an appropriate timescale. This offers some mitigation to this risk.
- The differential cost analysis is at best only accurate to first order: the uncertainty in costing adds additional cost risk to the project.
- The added complexity to the design and then construction phase for SKA1 adds programmatic risk to the project.

Conclusion

A majority of the SOWG did not express a clear preference between these two options.

3 Analysis and down-selection of SKA1 options

The SOWG performed a systematic analysis of the range of combinatorial options possible for a dual-site SKA1 implementation. These are summarised in the following table (colour encoding indicates the order in which options were downselected).

	Options							
Component	**1**	**2**	**3 (A)**	**4**	**5 (B)**	**6**	**7**	**8**
SKA1_Low	**ANZ**	**SA**	**ANZ**	**ANZ**	**SA**	**SA**	**SA**	**ANZ**
SKA1_Mid	**SA**	**ANZ**	**SA**	**SA**	**SA**	**ANZ**	**ANZ**	**ANZ**
SKA1_AIP_Survey			**ANZ**	**SA**	**ANZ**	**ANZ**	**SA**	**SA**

The SOWG down-selection analysis can be summarised as follows:

1. Only options 1 and 2 are possible if PAFs are not available
2. Options 7 and 8 do not permit reuse of precursor instruments and are discarded
3. Options 4 and 6 do not make optimal reuse of precursors but do co-locate all dishes
    a. Analysis of survey speed (see Section 2) demonstrates that best survey speed is achieved by incorporating ASKAP with 60 SKA dishes equipped with PAFs.
    b. Option 6 can only reach required A/T if all 250 dishes are used with low T_{sys} feeds.
    c. Option 6 does not permit incorporation of MeerKAT which provides significant A/T (25% of SKA1 requirement).
    d. Discard option 6.
    e. Option 4 maximises A/T at a single site, does not allow incorporation of ASKAP and requires PAFs on more than 60 SKA antennas to get to survey speed of options 3 or 5.
    f. Discard option 4.
4. Neither options 1 or 2 provide a scientific advantage for SKA1 over a single site solution in which the precursor at that site is incorporated into SKA1.
    a. Option 2 does not make good reuse of MeerKAT with high A/T, but offers potential for higher survey speed.
    b. Option 2 locates high power demand at RSA site and low power demand at ANZ site; sufficient power to also co-locate processing at RSA site.
    c. With co-location of processing at RSA site data transport capability does not provide an additional constraint
    d. Overall preference for option 1 over option 2 on grounds of optimal reuse of MeerKAT.
5. Strong preference for options 3 and 5 – take forward to further detailed analysis.
6. Option 1 remains possible, but offers no advantage on scientific grounds
    a. Additional science capability will be a strong motivation for wider community to appreciate and accept a dual-site implementation option.

4 Risks and risk mitigation

From a project management perspective, the dual site options being considered in this report bring with it a number of risks.

- Higher costs because of duplication of some core infrastructure and reduced flexibility of sharing infrastructure between elements.
  - It will not be possible to fully mitigate this risk; detailed design will reduce impact but cannot eliminate it.
  - As noted above current cost modelling indicates dominance of infrastructure costs within the cores, however it will still be necessary to provide bulk infrastructure on two sites which while dependent on the detailed design is likely to represent a significant cost risk
- Greater complexity, and duplication of effort, required to manage and maintain different configurations at two locations (one in Southern Africa, another in Australia/New Zealand). An increase in system engineering and project management activities as well as dealing with an increased number of interfaces.
  - Mitigated to some extent by splitting on technology grounds but cannot be eliminated.
  - Mitigated through greater project partner inclusion
  - Additional cost incurred during detailed design phase (PEP) for project staff
- Increased complexity and effort for commissioning and verification resulting from geographical split.
  - Mitigated to some extent by splitting on technology grounds but cannot be eliminated and will require additional project staff
- Increased complexity in managing a project in two countries with different legal, taxation and social characteristics.
  - Mitigation requires additional staff in project office.
- Greater operational costs resulting from the long term requirement to maintain two sites. Increase in maintenance costs, since less ability to provide flexible engineering support between technology elements.
  - Fractional increase in operational costs may be small compared to overall operations budget if the latter is dominated by power costs which are split. The SSAC report notes the projected lower cost of power on the RSA site.
  - Some operational aspects need not be duplicated (science support, SDP support etc.)
  - Mitigated by split on technology grounds with specialist engineers on different sites for different technologies, but cannot eliminate all additional maintenance costs
- Incorporating ASKAP and/or MeerKAT increases the number of interfaces and introduces new standards which may not be compatible with SKA needs. Incorporation of operating radio telescopes (or ones still under construction) into the SKA could also impact the availability of the precursors and lead to delays.
  - Mitigated but not eliminated by early convergence between SKA project and the precursors
  - Becomes increasingly less important during transition from phase I to phase II

    - Bringing science delivery of precursors under SKA project provides potential to maximise science return beyond that of independent precursor projects which are not formally coordinated.
- Dual site implementation for SKA1 could be seen as “Splitting the SKA” with consequential loss of identity, and possibly viewed as the build-out of the precursors as two national telescopes (ASKAP++ and MeerKAT++) rather than the SKA. Substantial risk to lose the SKA science vision and make the project unattractive to other partners.
    - Mitigation is to provide early convergence of the precursors and the SKA project and enforce the authority of the SKA project at an appropriate time (see Section 5).

The SSAC Report lists a number of risks which could be relevant when considering dual-site implementation options. The SOWG considered the implications of each risk factor on dual-site implementations: the following is a summary of those factors which directly effect consideration of a dual-site implementation.

1. Ionospheric Turbulence: nothing specifically related to a dual-site implementation.
2. RFI environment:
    - The SSAC report concludes that insufficient information was provided to discriminate between the RFI environments in the core
    - The GSM suppression technique that has been implemented in South Africa is untested by the SPO and remains a risk.
    - The SOWG noted the commentary of the Board on the RFI consideration on remote stations – the issue of RFI environment on remote stations did not play a role in the SOWG’s consideration of implementation options
3. Radio Quiet Zone protection:
    - The SOWG noted that it may be difficult to maintain full protection of the Radio Quiet Zone on a given site if the scale of the SKA built there is significantly reduced.
4. Long Term RFI Environment: The SSAC Report states that "We find that legal enforcement may be a more severe problem for remote sites in the RSA partner countries, mitigated, however, by the higher tolerance of RFI for signals from the remote sites."
    - The SOWG did not feel this should have a high weight in consideration of dual-site options.
5. Array Science Performance: the SSAC report notes the advantage of RSA on long baseline imaging performance for Phase II
    - the SOWG resolved to include new information on configuration if relevant
6. Physical Characteristics of the Sites: none.
7. Tropospheric Turbulence: The average altitude of the sites in Southern Africa is about 1000m and 350m for Australia/New Zealand. This means that there is a preference in siting the high frequency array(s) in Southern Africa. This applies in particular to frequencies above 3 GHz, and therefore for the dish-array.
    - This factor was taken into account in the SOWG dual-site options analysis

The SOWG noted the conclusions of the SSAC that overall the SSAC report concludes there is significant advantage to Australia/New Zealand in the “other selection factors” and a significant advantage to Southern Africa in the “Implementation plans and cost factors”. Taking together the other selection factors and the implementation plan and cost factors the SOWG noted that there is potential for some risk mitigation for the project across of these factors for a dual-site implementation option.

5 Management requirements for implementation options

The programmatic risks and perception risks to the outside community of adopting a dual-site implementation option have been noted above.

A critical requirement is to maintain the “big-vision” of the SKA project. Introducing the concept of the SKA observatory in which the science mission is implemented using technology appropriately located enables this vision to be maintained. It is essential that if a dual-site implementation is adopted it does not become a split of the SKA into two, competing, instrument programmes. Therefore string central management and operation of a single SKA observatory is required. This will also permit, in the long term, further expansion of the observatory possibly (if well motivated) using further geographical sites.

There is clear advantage for SKA1 of incorporating into the SKA1 design the precursor instruments and infrastructure. This however has the potential to introduce very high levels of risk in the project if the current situation of independent precursor programmes continues indefinitely. There are two clear critical milestones in the development of the SKA project incorporating the precursors:

1. The decision, or not, to follow a dual-site implementation model (now)
2. Obtaining funding to proceed with SKA1 construction (2014 or afterwards)

To de-risk the adoption of a dual-site implementation, or indeed a single site implementation incorporating the precursor on that site, the future management and interaction between precursors and the SKA project must be addressed at this time.

The SOWG noted the following possible model for incorporation of the precursors into the overall project structure: this discussion represents the view of those SOWG members not directly associated with the precursors.

A good model and timeline for incorporating the precursors is:

- Agree at this stage that the precursor projects should, in the same way as other independently funded SKA Work Package contributions, report progress formally to the Board of the SKA organisation against the agreed delivery of the precursors and infrastructure.
- The SKA Organisation comes to an agreement with the hosts on the responsibility for the delivery to the scientific community of the planned science programmes with the precursors. The option of these science programmes becoming part of the early SKA science programme in due course should be considered.
- Development of the precursors of their science delivery beyond the currently agreed programme comes under the auspices of the SKA Organisation.
- As the SPO staffing level increases the SKA Director General and precursors develop plans and timescales to begin integration of the precursor projects with increased direct involvement of the SPO in the precursor projects.
- The timeline will require that by the time funding for SKA1 construction is agreed the precursors will come fully within the control of the SKA Organisation / Director General.

The additional benefit to the external perception of the project of achieving convergence of the precursors and the central project the SOWG believes to be a major advantage to the project in the widest sense. The SOWG noted that following this route might represent an extension of the mandate of the SKA Organisation as currently defined in the Business Case.

Appendices

Assumed description for SKA1 and SKA2

	SKA1_low	SKA1_mid	SKA2_low	SKA2_mid_dish	SKA2_AIP_AA	AIP_PAF	Comments
Collector type	Sparse AA [1]	15m dish [1]	Sparse AA [1]	15m dish [1]	Dense AA [1]	15m dish+PAF [1]	Offset feed dishes
No. of collectors	280 [3][9]	250 [1]	280 [3][10]	2,500 [11]	280 [3]	2000 [15]	
Frequency range GHz	0.07 – 0.45 [1]	0.45 – 3.0 [1]	0.07 – 0.45 [2]	0.45 – 10 [11]	0.4 – 1.4 [2]	0.45 – 3.0 [13]	50MHz goal
Max bandwidth GHz	0.38 [1]	1.5 [8]	0.38 [2]	*Depends on feed*	1.0 [8]	0.3	
Dish feeds: 1. GHz 2. GHz 3. GHz		0.45 – 0.9 [1] 0.8 – 1.6 [1] 1.5 – 3.0 [1]		To be decided		0.45 – 0.9 [13] 0.8 – 1.6 [13] 1.5 – 3.0 [13]	
Effective FoV deg^2		*1GHz:* 1.0 [1]	200 [4]	*1GHz:* 1.0 [1]		*0.5GHz:* 144 deg^2 [13] *1GHz:* 36 deg^2 [13] *2GHz:* 9 deg^2 [13]	15m dish FoV
No. of beams	160 [1]	1		1		36	
Sensitivity: /element m^2K^{-1} total sensitivity m^2K^{-1}	*131 MHz:* 7.2 [8] *131MHz:* 1,515 [1] *300 MHz:* 889 [1]	*1-2GHz:* 4.0 [8] *1-2GHz:* 1,031 [1] *0.45-1GHz:* 773 [1]	*>90MHz:* 14.3 [8] *>90MHz:* 4,000 [2]	4.0 [8] 10,000 [2]	*<1.2GHz:* 36 [8] *<1.2GHz:* 10,000 [2] *1.4GHz:* 5,000 [2]	*1-2GHz:* 3.5 [14]	Sensitivity of AA on boresight

		SKA1_low	SKA1_mid	SKA2_low	SKA2_mid_dish	SKA2_AIP_AA	AIP_PAF	Comments
Configuration:								
<1.0 km dia	%	50% [1]	50% [1]	30% [1]	20% [1]	30% [1]		
1-5 km dia	%	20% [1]	20% [1]	36% [1]	30% [1]	36% [1]		
SKA1/2: 5-100/180 km	%	30% [1]	30% [1]	34% [1]	30% [1]	34% [1]		
>180km	%	-- [1]	-- [1]	-- [1]	20% [1]	-- [1]		
Correlator O/P max data rate	GB/s	40 [3]	*Imaging:* 50 [3] *Pulsar search:* 364 [3]					Data rate to return to post-processing
Power max:								
per element	kW	12 [6]	4 [12]	31 [5]	4 [12]	96 [5]	2 [12]	
Correlator	kW	480 [7]	37 [7]					
Total	kW	3,730 [8]	1,037 [8]					

References & notes:

[1] Delta System CoDR, High level system description: 06-wp2-005.030.010-td-002-a-highlevelsysdescription

[2] System CoDR, High level system description: 04-WP2-005.030.010-TD-001-A_HighLevelSysDescr

[3] Software and computing CoDR analysis of requirements from DRM D2A_WP2-050.020.010-RR-001-E_Analysis_of_DRM_Requirements

[4] AA Concept descriptions: AA-CoDR 02-WP2-010.020.010-TD-001-E-AAConceptDescription

[5] AA Deployment and Operation: AA-CoDR 04-WP2-010.020.010-TR-001-F_AAImplementation

[6] AA Deployment and Operation: AA-CoDR, scaled from 50 to 280 arrays *(maybe a bit low)* 04-WP2-010.020.010-TR-001-F_AAImplementation

[7] Memo 139: A Software Correlator for SKA1

[8] Calculated

[9] For imaging quality the number of AA-low stations has been calculated to be 280 of ~80m diameter in the Software and computing CoDR.

[10] For SKA2 it is assumed that the number of SKA-low stations is constant and the diameter increases to ~180m.
[11] Scaled from SKA1 to the sensitivity and frequency range of SKA2
[12] Per AA station, per dish; incremental cost of adding PAF to dish. (MvA)
[13] Each PAF can cover about 2.5:1 with good noise performance but FOV per band is better optimised with a 2:1 design. Costs area already dominated by digital electronics of the beam-former and this can be shared over multiple PAF antenna arrays if required. (RB)
[14] The SPF value of 4.0 m^2/K implies T_{sys}/η=44 K. Current PAF performance is already T_{sys}/η=55 K, while next generation performance of T_{sys}/η=50 K is a conservative projection. (RB)
[15] The underlying assumption is that PAFs are deployed in all dishes < 180 km. (RB)

Summary of on-site infrastructure and value

Australia / ASKAP

Item	Total value	Location	Status	% SKA1	Comments	SKA1 value
ASKAP						
ASKAP Antennas	$10,260,000	MRO	Existing	100%	36 x 12m 3-axis	$10,260,000
PAF receivers and beamformers	$11,160,000	MRO	15% done	100%	18 192-element chequerboard	$11,160,000
Correlator	$3,500,000	MRO	10% done	10%	SKA1 Testing use only	$350,000
ASKAP Telescope total:	**$24,920,000**					**$21,770,000**
Infrastructure:						
Network hardware	$1,900,000	MRO, MSF	Existing	100%	40Gbps MRO-Perth, 1Tbps MRO	$1,900,000
Antenna foundations	$3,060,000	MRO	Existing	100%	ASKAP antennas	$3,060,000
Main site acces road	$1,980,000	MRO	Existing	100%	9 km main road	$1,980,000
3 spiral arm roads (existing)	$945,000	MRO/Bool.	Existing	100%	3 x 15 km tracks along arms	$945,000
Power & fibre cabling to antennas	$4,995,000	MRO	Existing	100%	ASKAP 6km site	$4,995,000
On-site access tracks	$567,000	MRO	Existing	100%	ASKAP 6km site	$567,000
Control Building	$10,000,000	MRO	Existing	100%	$6M in RFI shielding	$10,000,000
Accommodation at Boolardy	$5,250,000	Boolardy	Existing	100%	35 people	$5,250,000
MRO airstip	$150,000	MRO	Existing	100%	RFDS standard	$150,000
Boolardy airstip	$80,000	Boolardy	Existing	100%	RFDS standard	$80,000
MSF (Geraldton)	$4,600,000	Geraldton	Existing	100%	Includes architect fees	$4,600,000
Perth Head office accommodation	$6,600,000	Perth WA	Existing	100%	1650 m sq	$6,600,000
Fibre link MRO - Pawsey	$30,000,000	MRO-Perth	Existing	100%	24 fibre pair	$30,000,000
MRO Power Plant	$15,500,000	*MRO*	*Funded*	*100%*	*1MW capacity*	*$15,500,000*
MRO Power plant - EIF Phase II	*$11,000,000*	*MRO*	*Funded*	*100%*	*Additional 1 MW (Install 2014)*	*$11,000,000*
Pawsey HPC for SKA Science	*$80,000,000*	*Perth, WA*	*Funded*	*25%*	*2 Petaflop (install 2013)*	*$20,000,000*
International connectivity link	$16,000,000	Aust-RoW	Existing	100%	10Gbps	$16,000,000
MWA fibre and power cabling	$2,915,000	MRO	Existing	100%	25 Receiver "drop points"	$2,915,000
Infrastructure Total:	**$195,542,000**					**$135,542,000**
Grand Total:	**$220,462,000**				**Total SKA Phase 1 value**	**$157,312,000**

Key points to note

- Value of total existing/funded ANZ infrastructure for SKA1 **$136M** ($90M already built, all funded)
- Value including ASKAP **$168M** ($104M already built, all funded)
- Funding for additional 18 PAFs being sought.

Summary

ASKAP	$21,770,000
Network/Computing	$71,855,000
Buildings	$26,450,000
Power	$30,455,000
Roads/Airstrip	$3,722,000
Foundations	$3,060,000
S/C	$31,200,000

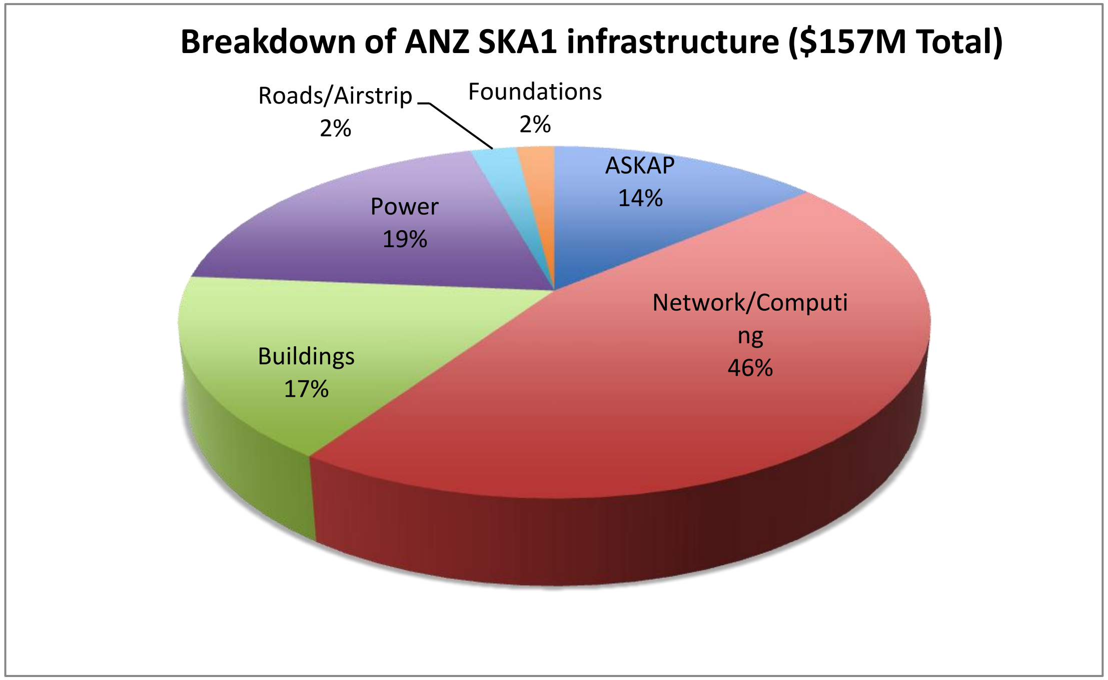

SA / MeerKAT

Item	Description	2012 RSA Cost	2007 Euro Cost
1	**Kat 7 infrastructure**	**R 211,000,000.00**	**€ 20,488,100**
1.1.	Upgrade to Provincial Road to site (NC Govt)	R 130,000,000.00	€ 12,623,000
1.2.	Construction of KAT 7 road	R 17,000,000.00	€ 1,650,700
1.3.	Dish Assembly Shed	R 10,000,000.00	€ 971,000
1.4.	Site Complex accommodation	R 1,500,000.00	€ 145,650
1.5.	Accommodation - portacabins	R 700,000.00	€ 67,970
1.6.	New Klerefontein water, sewer, electrical reticulation	R 1,000,000.00	€ 97,100
1.7.	Kat 7 fibre ducting and electrical reticulation	R 8,000,000.00	€ 776,800
1.8.	Land (Meys Dam and Losberg) - capital cost	R 16,000,000.00	€ 1,553,600
1.9.	Buy-out of Usufruct Agreements	R 14,000,000.00	€ 1,359,400
1.10.	KAT 7 antenna foundations	R 5,000,000.00	€ 485,500
1.11.	RFI shielded containers	R 1,600,000.00	€ 155,360
1.12.	KAT 7 diesel generators	R 1,200,000.00	€ 116,520
1.13	Klerefontein Support Base (Offices, workshops, accommodation)	R 5,000,000.00	€ 485,500

2	**MeerKAT infrastructure (Phase I)**	**R 1,221,000,000.00**	**€ 118,559,100**
2.1.	Roads, Landing strip, construction camps, electrical and fibre ducting reticulation (Bid SKA SA SSLE 004/2011)	R 105,000,000.00	€ 10,195,500
2.2.	MeerKAT buildings (Bid SKA SA SSLE 005/2011)		€ 0
2.2.1.	Extensions to Dish Assembly Shed and Pedestal Integration shed for MeerKAT	R 17,000,000.00	€ 1,650,700
2.2.2.	Karoo Array Processor building	R 40,000,000.00	€ 3,884,000
2.2.3.	Power Facility	R 30,000,000.00	€ 2,913,000
2.2.4.	Klerefontein workshops / stores	R 5,500,000.00	€ 534,050
3.3.	Rotary UPS, MV Switchgear, Transformers (Bid SKA SSLE 009/2011)	R 50,000,000.00	€ 4,855,000
3.4.	Building Management System (Bid SKA SSLE 007/2011)	R 4,500,000.00	€ 436,950
3.5.	RFI Shielding for KAPB (Bid SKA SSLE 010/2011)	R 11,000,000.00	€ 1,068,100
3.6.	Antenna Foundations (Bid SKA SSLE 006/2011)	R 45,000,000.00	€ 4,369,500
3.7.	Perimeter Fencing (MeerKAT Site Complex; Klerefontein offices)	R 10,000,000.00	€ 971,000
3.8.	33kV power line	R 46,000,000.00	€ 4,466,600
3.9.	Upgrade to Karoo substation (meerKAT - Civil works and transformers)	R 15,000,000.00	€ 1,456,500
3.10.	10Gbps optic fibre link (long-haul fibre)	R 57,000,000.00	€ 5,534,700
3.11	MeerKAT optic fibre network (internal)	R 15,000,000.00	€ 1,456,500
3.12	64 MeerKAT Dishes	R 680,000,000.00	€ 66,028,000
3.13	Receivers for dishes (excl. NRE)	R 90,000,000.00	€ 8,739,000
	KAT-7 + MeerKAT Total:		**€ 139,047,200**

SKA1 processing power requirements

The following assumptions were made about the required processing performance and power requirements of the post correlator/beamformer processing system. These numbers are justified in an associated paper.

Visibility data rate:	350GB/s
UV processor max. raw performance:	35 PFLOP/s
UV processor actual performance:	8.75 PFLOP/s
Imaging processor performance:	875 TFLOP/s
Total buffer storage:	700TB
Typical power:	<3.0 MW
Maximum power:	4.7 MW

The construction would consist of 3,500 2U servers with GPUs fitted with a regular HPC backend. The UV processor could be constructed in 175 racks of 20 servers each.

Supporting Documents

- Terms of reference
- Infrastructure description RSA/MeerKAT
- Infrastructure description ANZ/ASKAP
- Estimates of computational power requirements for SKA1

Operation of the SOWG

Schedule of meetings: records of meetings

The SOWG met on five occasions:

- A 1.5-hr kickoff meeting/telecon on Thursday 19th April in Manchester
- A 3-hr videocon on Tuesday 1 May
- A three-day face-to-face meeting from 7-9 May in ASTRON.
- A 1.5-hr videcon on Wednesday 16th May
- A 1.5-hr Videocon on Thursday 17th May

All members of the Working Group were able to attend the first two meeting, Di Li was not able to attend the face-to-face: Russ Taylor and Joe Lazio joined for the full meeting via videocon. Di Li and Russ Taylor were unable to join the videcon on 16th May: all members were present on the videocom on 17th May.

Additional work in preparation for the meetings was done by email discussion and circulation.

Detailed minutes for all of these meetings were taken and are available as additional documents to this report.

Members of the working group

Members of the working group are:

- Professor Paul Alexander (Chair)
- Professor Phil Diamond
- Professor Mike Garrett
- Professor Justin Jonas
- Professor Russ Taylor
- Dr Joseph Lazio
- Dr Michiel van Haarlem
- Dr Melanie Johnston-Hollitt
- Dr Luigina Feretti
- Dr Di Li
- Dr S Berry (Secretary)

Abbreviations used throughout this document:

- SOWG Site Options Working Group
- SSAC Site Selection Advisory Committee
- SA Southern Africa site option
- ANZ Australia / New Zealand site option
- AA Aperture Array
- PAF Phased Array Feed
- SWG Science Working Group of the SKA project
- PEP Project Execution Plan
- AIP Advanced Instrumentation Programme

Main input documents and information sources

A document repository was established for the SOWG as part of the SKA Wiki and the SOWG were given access to the full range of SKA documentation including the input documents to the site selection process. The SOWG used as its main input however a subset of documents which were particularly pertinent to its work. These were as follows:

- SSAC Report
- Commentary of the Board on the SSAC Report
- the remarks on the SSAC report sent to the Board by the host candidates
- the SSG and SSAC presentations to the Board
- the responses to the Request for Information from the host candidates and other information made available by the sites to the SSAC
- Memo 125: defining the scope of SKA1
- DRM for SKA1: quantifying the requirements for SKA1

- DRM for SKA2 and SKA Science Case (Carilli and Rawlings)
- “On a split site solution for the SKA” (Schilizzi)
- Minutes of the SSEC pertaining to the split- or dual-site considerations