Memo 150

Optimal Partitioning of SKA-Low Antenna Elements

D. C. Price
J. Hickish
D. Sinclair
M.E. Jones

September 2013

Optimal partitioning of SKA-low antenna elements

SKA Memorandum 150

D. C. Price[1,E], J. Hickish[1], D. Sinclair[1], M. E. Jones[1]
[1]University of Oxford, Keble Rd, Oxford, United Kingdom OX4 1AH
[E] Author email: danny.price@astro.ox.ac.uk

SKA-low will consist of approximately one quarter of a million antenna elements, dividaed in to a number of stations. In this memo, we investigate the relationship between the number of antennas in each station versus beam-forming and correlation computational requirements, output data rates and station beam quality. We find that the current baseline design is close to computationally optimal, but that larger stations will produce higher quality beams and will decrease output data rates.

1 Introduction

In the current SKA Phase 1 Baseline Design (Dewdney, 2013, henceforth 'SKA-low v2'), the SKA-low telescope will consist of approximately one quarter of a million antenna elements. They are currently specified to be partitioned into 911 beamformed stations of 35 m diameter, each of which consists of 289 antenna elements. In earlier baseline specifications (e.g. Dewdney et al., 2010, henceforth 'SKA-low v1'), each station was to be significantly larger: 11,200 antenna elements per 180 m diameter station, with 50 stations in total. The station size revision was motivated by the findings of the Epoch of Reionization (EoR) working group, whose report (Mellema et al., 2012) recommends stations of size ∼35 m. In contrast, the original aperture array sizing of 11,200 elements was specified in order to deliver low sidelobe levels in the station beam; the motivation for this was to suppress bright sources and achieve high dynamic range imaging (see for example Braun, 2013).

In this memo, we investigate the repercussions of the change in the number of antennas per station upon sensitivity to large-scale sky brightness fluctuations, computational requirements, output data rates, and the quality of each station's beam; discussion of each of these quantities is presented in the four sections below.

2 *uv*-coverage requirements

SKA-low's sensitivity to large-scale sky brightness fluctuations is of paramount importance to EoR science. As such, we must pay careful attention to how the partitioning of antennas into aperture arrays affects large-scale sensitivity.

Within a radio interferometer such as SKA-low, each pair of stations is sensitive to a range of spatial frequencies (i.e. Fourier, or '*uv*' modes). The *uv*-coverage of a station pair is governed by two factors: baseline length and station size. The central *uv*-mode is determined by the baseline length and the range of modes is determined by the station size. This is illustrated in Figure 1. Each 'point' in *uv* space that is measured is in fact an average over an area in the *uv* plane defined by the autocorrelation function of the station aperture field distribution. For a single pointing of the interferometer, this defines the resolution in *uv* space that is available to a power spectrum measurement (in image space, it defines the field of view). For multiple pointings, the *uv* space resolution is improved in proportion to the total sky area surveyed (which may or may not be in a contiguous area). Increasing the sky coverage also decreases the sample variance (the uncertainty in the power spectrum estimate due to

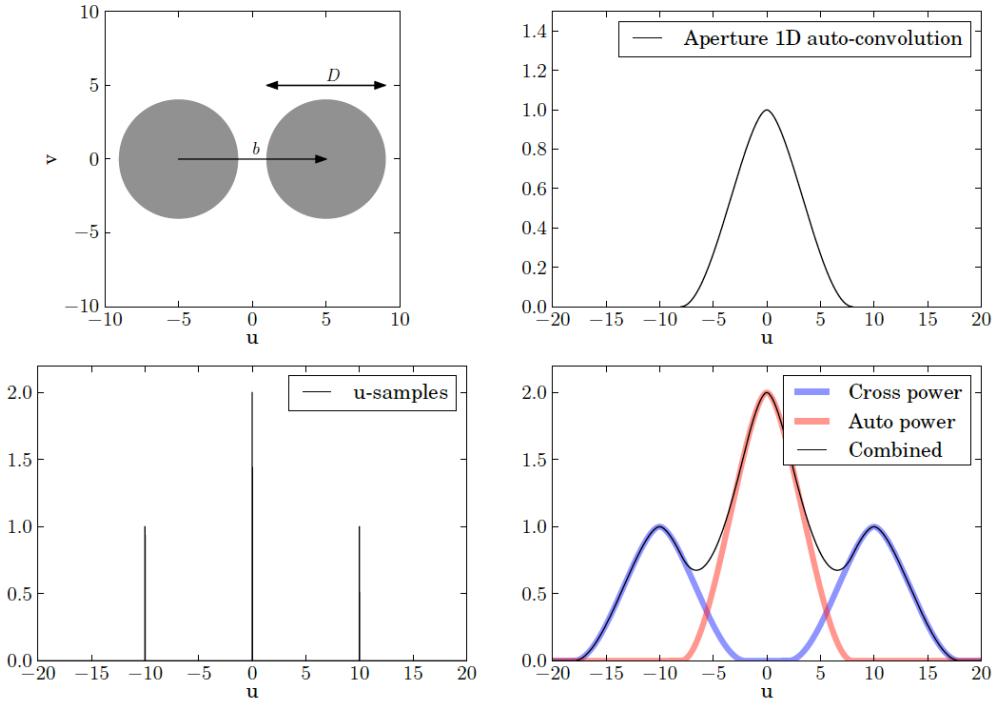

Figure 1: Top left) Cartoon of a simple two-element interferometer formed from circular apertures with diameter D=8, separated by a baseline $|\vec{b}|$=10. Top right) The 1D auto-convolution function of a circular aperture with D=8. Bottom left) The *uv*-coverage of the depicted two-element interferometer. Bottom right) The relative sensitivity of such an interferometer to angular scales in the sky brightness temperature distribution. Figure reproduced from Hickish (2013).

the fact that only a finite number of modes have been sampled). Well-established formalisms exist for combining multi-pointing data in this way for continuum measurements (Myers et al., 2003), and these techniques can be generalised to spectral measurements.

For a power spectrum measurement, or indeed for an imaging experiment that is required to image a defined range of angular scales, the station size should be set such that the *uv*-coverage of a pair of neighbouring stations is sensitive to the smallest *uv*-mode of interest, and the sky coverage should be sufficient to provide the required *uv* space resolution. Such an approach has been successfully used to measure the cosmic microwave background (at GHz frequencies) with interferometeric experiments such as CBI (Pearson et al., 2003) and VSA (Dickinson et al., 2004). In both these experiments, the power spectrum was measured in bins whose central *uv* point was at a radius approximately 1/3 of the antenna diameter, and of course the same information is also available in the image plane. Below this radius the sensitivity drops off rapidly due to the steep tail of the autocorrelation function for circular apertures. For EoR, the largest scale of interest is $\sim 1^\circ$ (Zaroubi et al., 2012); this sets the smallest *uv*-mode to which SKA-low must be sensitive. It follows that the minimum baseline size (and maximum station size) could increased to 85 m (1° FoV at 200 MHz), or even larger, without jeopardizing *uv*-coverage at the largest scale of interest.

The above arguments focus on the shortest baselines, which will be produced by the stations in the densest part of the SKA-low core. For the outer parts of the core, and for the extended baselines, the arguments about small station size do not apply as the stations are only participating in baselines much longer than the station diameter, and hence only measuring angular scales much smaller than the field of view of a single station. Here there is also the potential to increase the station size, and recover the field of view with multi-beaming, if this is computationally advantageous. This configuration,

with large multi-beamed antennas surrounding smaller antennas in the core has also been considered in the context of high-frequency intensity mapping experiments (Bowman, 2012), and is sometimes described as a *foveated* array. It is also similar to the concept used in heterogeneous arrays such as AMI and CARMA where small antennas are used for the short baselines and larger antennas for the longer baselines.

In contrast to these ideas, Mellema et al. (2012) argue for a small station size of 35 m based on a minimum FoV of $\sim 5^\circ$ for EoR power spectrum studies (sample variance limit), and a minimum FoV per beam of $\sim 2^\circ$ for EoR imaging. This minimum beam size for imaging is set by the concern for power spectrum measurements that "although it might be possible to produce a mosaic using multibeaming, this is unlikely to accurately reconstruct modes with wavelengths longer than the size of an individual field". They conclude that stations of 35 m diameter are optimal as they have a FoV of 2.5-10 degrees over 200-50 MHz; as such, SKA-low v2 has adopted this 35 m station size.

Nevertheless, from a basic *uv*-sampling point of view, an SKA-low implementation with stations as large as 85 m will still be sensitive to the power spectrum modes of interest for EoR science. What is unclear – and perhaps the root concern of Mellema et al. (2012) – is whether or not multi-beam systems can reach the calibration requirements for imaging. Recovery of spatial frequency information resolved at the sub-aperture scale does indeed require that the visibility data be well calibrated and that the aperture illumination function (or, equivalently, the station beam pattern) be well understood.

Detailed consideration of these calibration issues is beyond the scope of this memo. However, we show that should these as yet unquantified calibration concerns turn out to be overstated, larger stations have significant advantages. In particular, data output rates from the correlator could be decreased significantly (Section 4), and station beam quality would increase (Section 5). Decreasing data output would have cost benefits for data storage and processing, and increasing beam quality will improve imaging performance, due to increased suppression of bright sources outside the primary beam.

3 Computation rates

To ensure that science goals are reached at minimum cost, the first quantity that one might optimize for SKA-low is its computational requirements. Here we consider only the computational costs of beamforming and correlation, not post-correlation calibration and imaging (however, we note that these post-correlation computing costs tend to increase with the size of the field to be instananeously imaged). As the cost of the correlator will scale with the square of the number of stations, having very small stations will require an unfeasibly large correlator. On the other hand, large stations require more beams to reach FoV requirements, which increases the cost of the beamformer. In this section, we present a simple model of the SKA-low digital signal processing (DSP) systems, and show that there is a computationally optimal station size. Our approach is similar to that presented by Weinreb & D'Addario (2001), in the very early phases of SKA design. An expanded discussion of DSP and station size for SKA-low is given in Chapter 3 of Hickish (2013).

In order to create a simple computational cost model for SKA-low, we make some assumptions about the DSP systems. Firstly, we assume that floating-point arithmetic is used throughout. This is so that we may count floating-point operations per second (FLOPs). Secondly, we assume a particular order of operations, and that these operations are all-digital: Fourier transform ('F'), beamforming ('B'), then correlation ('X'); this is shown in a simplified block diagram in Figure 2, and is consistent with the SKA station beamformer concept description (Jones et al., 2011). We assume that a constant instantaneous FoV is built up through multi-beaming, and that every beam that is formed is correlated. There will thus be N_B identical correlator engines running in parallel: one for each beam.

With these assumptions, we are able to define a simple "computational cost" function, CF, in terms

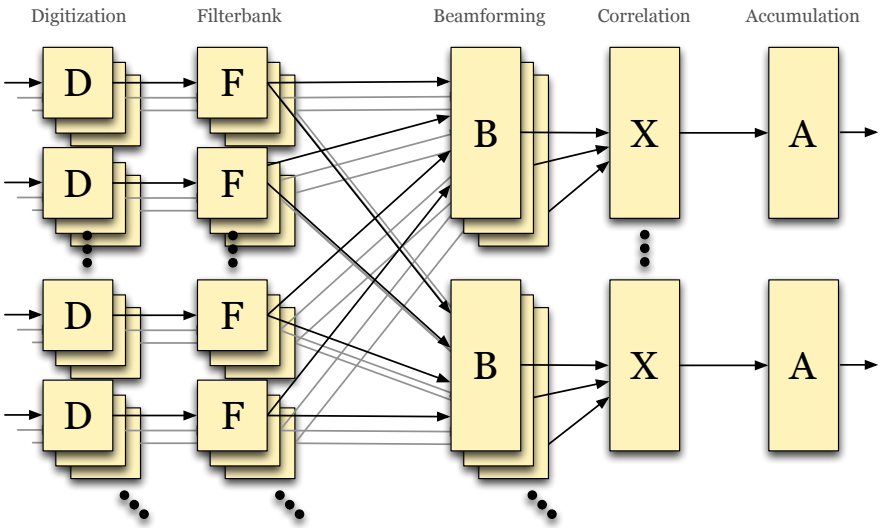

Figure 2: Block diagram of the simplified model of the SKA-low digital signal processing system that is considered in this memo. Within each station, every antenna element is first digitized (D), then channelized using a polyphasse filterbank (F). The channelized data for each are then beamformed (B), forming N_B beams per station. Each of these N_B beams are then correlated (X) with corresponding beams from other stations, and then accumulated (A).

of FLOPs. The total computational cost is the sum of the cost of the F, B and X engines:

$$CF = CF_F + CF_B + CF_X.$$

For now, we ignore the F engine, as channelization does do not affect the number of FLOPs required for the B and X engines. First, let's look at the B-engine. The computational cost can be approximated as

$$CF_B = N_A N_B N_S N_P;$$

the parameters in this equation — and the equations to follow — are defined in Table 1. Note that this equation assumes independent antenna weightings for each beam (i.e. partial digital beamforming); fewer computations are required in hierarchical (multi-stage) beamformers (Colegate et al., 2012; Jiwani et al., 2012).

Continuing on, as we have assumed a fixed FoV requirement, we can write $N_B = \Omega_R/\Omega_B$. From the antenna theorem, we have

$$\Omega_B A_e = \lambda^2,$$

so can relate FoV to collecting area. For an aperture array, the physical area of an array scales with $A \propto N_A$. Using the average antenna spacing from SKA-low v2, we find that for $A = \alpha N_A$ a value $\alpha = 3.33$. It follows that the number of beams required is given by

$$N_B = \frac{\Omega_R}{\Omega_B} = \frac{\alpha \Omega_R}{\lambda^2} N_A \equiv \beta N_A,$$

Figure 3a shows N_B as a function of N_A, for a number of different wavelengths. Note N_B is ceiling rounded so that it is always an integer. We can compute the average of β over the frequency range,

from λ_a=6 m (50 MHz) to λ_b=0.8 m (∼350 MHz), as

$$\begin{aligned}\bar{\beta} &= \frac{1}{\lambda_b - \lambda_a}\int_{\lambda_a}^{\lambda_b}\frac{\alpha\Omega_R}{\lambda^2}N_A d\lambda \\ &= \frac{\alpha\Omega_R}{\lambda_a\lambda_b} \approx 0.7\Omega_R.\end{aligned}$$

The beamformer cost function is then

$$CF_B = \beta N_A N_T N_P.$$

We can similarly define a computational cost function for the correlator. To form all correlations (including autocorrelations), we have

$$\begin{aligned}CF_X &= N_B(N_S^2N_P^2 + N_SN_P)^2/2 \\ &= \frac{\beta}{2}\left(\frac{N_T^2N_P^2}{N_A} + N_TN_P\right).\end{aligned}$$

Setting $N_P = 2$, the total cost function is then

$$CF = 2\beta N_T(N_A + \frac{N_T}{N_A} + \frac{1}{2}N_T),$$

This is plotted in Figure 3b, using fiducial values from Table 1. The cost is minimized for

$$\frac{\partial}{\partial N_A}(CF) = 2\beta N_T N_C(1 - \frac{N_T}{N_A^2}) = 0,$$

so the minimum occurs when

$$N_A = N_S = \sqrt{N_T}.$$

that is, when the number of antennas per station is equal to the number of stations. For 250,000 antenna elements, the optimum is given for $N_A = N_S = 500$. As such, the SKA-low v2 station size of N_A=289 is significantly more computationally optimal than the SKA-low v1 size of $N_A = 11,200$.

However, there are a number of factors which will alter the position of this minimum. The first is that as we have taken N_B to be an integer, the minimum is a function of frequency (as is apparent in Figure 3b). As such, fewer beams are required at longer wavelengths, so larger stations are favoured. A second factor is the precise implementation of the beamformer. A hierarchical beamformer that uses a first-stage analogue beamformer would require fewer computations (Jiwani et al., 2012), which also encourages larger stations.

In the above derivation, we assumed floating-point arithmetic. In fixed-point technologies (such as FPGAs), one must also consider how adders and multipliers are implemented. For example, multiple low-bitwidth complex multiplies can be conducted by clever use of a single high-bitwidth multiplier plus some additional logic (de Souza et al., 2007; Hickish, 2013). The relative cost of a low-bitwidth correlator may therefore be lower than that of a high-bitwidth beamformer; this would promote larger correlators (and as such, smaller stations). Nevertheless, significant deviation from this computed minima would likely suggest a computationally inefficient SKA-low implementation.

4 Data rates

Our second consideration is data transfer requirements. Of particular importance is the correlator output data rate, as these data must be either reduced in real-time to form secondary data products, or

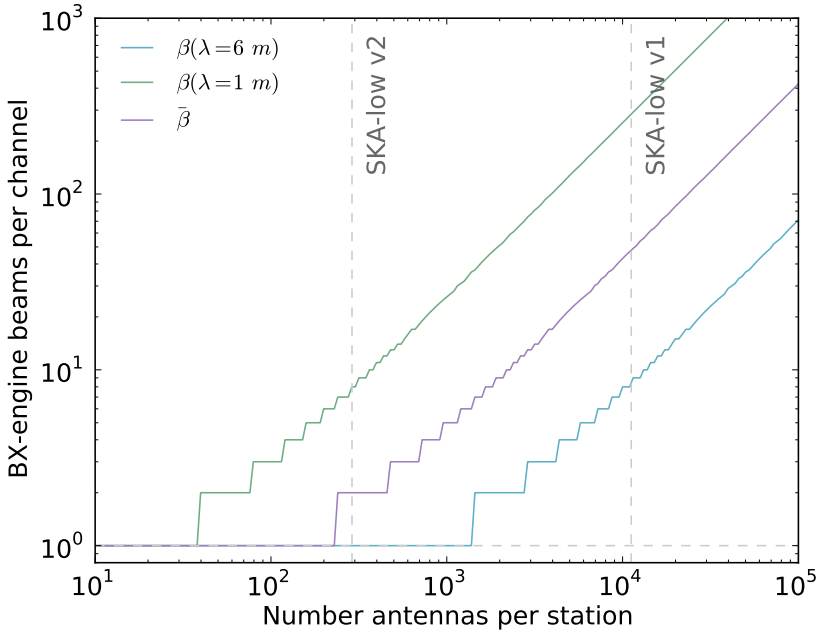

(a) The number of beams required to satisfy field of view requirement, Ω_R, rounded up to the nearest integer.

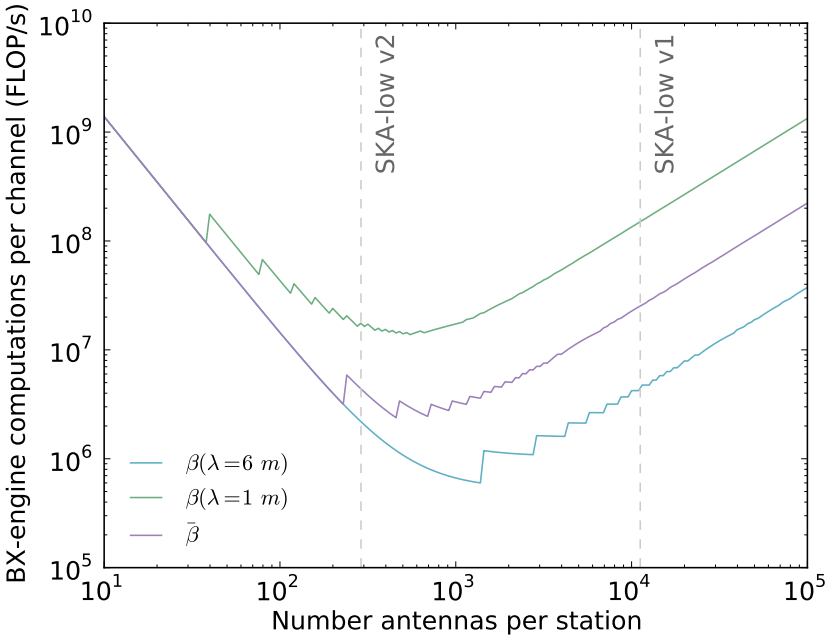

(b) The minimum computational cost function for a BX array.

Figure 3: Number of beams required (top), and corresponding number of computations required for the combined BX-engine. The number of beams is related to number of antennas per station by β, a function of frequency and required FoV.

Table 1: Parameters used in cost function.

Variable	Description	Value
N_A	Number antennas per station	1 - 250,000
λ	Wavelength	0.8-6 m
d_s	Diameter of station (depends on N_A)	<85 m
N_T	Total number of antennas	250,000
N_P	Number of polarizations	2
N_C	Number of channels	2.5×10^5
N_{bi}	Number of output bits from B-engine	64
N_{bo}	Number of output bits from X-engine	64
G_E	Directivity of a single element	8 dBi
Ω_R	Required field of view	25 deg^2
$\lambda_0/2$	average spacing of antennas	1.5 m
N_B	Number of beams	Ω_R/Ω_B
N_S	Number of stations	N_{TOT}/N_{APS}
Ω_B	Field of view of a single beam	$\lambda^2/(4\pi d_S^2)$
A_S	Collecting area of a station	$N_{APS} \times A_{EL}$
A_E	Collecting area of a single element	$G_E\lambda^2/4\pi$

archived and stored. In a similar manner to the computational cost function above, we may compute data rates throughout the system as a function of antennas per station.

4.1 Correlator input data rate

The input data rate to the correlator depends directly upon the total output data rate of the station beamformers. This data rate is given by

$$\begin{aligned} DR_I &= \Delta\nu N_{bi} N_B N_P N_S \\ &= \Delta\nu\beta N_{bi} N_T. \end{aligned}$$

That is, the correlator input data rate does not depend upon how antennas are partitioned into stations.

4.2 Correlator output data rate

In constrast, the correlator output data rate is indeed a function of antennas per station. For N_B parallel correlators (one per beam), each with an intrinsic output data rate of DR_X, the output data rate DR_O is given by

$$\begin{aligned} DR_O &= N_B(DR_X) \\ &= \frac{1}{2} N_B N_S^2 N_P^2 N_{bo} N_C/\tau \end{aligned}$$

Here, the minimum integration time τ is set by the FoV and baseline length. We use the formula

$$\Delta\theta\tau \ll \frac{\theta_S P}{2\pi},$$

where $P \approx 86164$ seconds is the period of earth's rotation per sidereal day, θ_s is the synthesized beam's resolution, and $\Delta\theta$ is angular distance from the phase centre of a patch of sky to be imaged (we set

this to the size of the station beam). We then set a threshold m=5, and compute τ as

$$\tau = \frac{\theta_S P}{\Delta\theta 2\pi} m$$

In addition to this limit, there are additional dump requirements on this borne out of ionospheric and instrumental gain fluctuations. Setting an absolute maximum integration time of 60 s (approximate ionospheric coherence time at 74 MHz, Kassim et al., 2007), the maximum acceptable integration times for SKA-low are shown in Figure 4a, once again using the fiducial values of Table 1. The corresponding data output rates are then shown in Figure 4b.

From Figure 4b, it is clear that larger stations are preferable as their data rate is lower, and that SKA-low v2 will output an order of magnitude more data than SKA-low v1 would. In summary, increasing the size of SKA-low stations will decrease the data output rate of its DSP systems.

5 Beam quality

We have now discussed that somewhat larger stations than the SKA-low v2 baselines can maintain EoR science capability, and have lower data rates. We now briefly touch upon how a station's beam quality is affected by the number of constituent antennas. There are many ways in which beam quality can be quantified, but for brevity, we consider only the average sidelobe level. A more extensive discussion of beam quality issues for SKA-low will be given in a later publication (Sinclair et al., in prep).

Having a well-understood and well-behaved beam pattern is a requirement for high dynamic range synthesis imaging and calibration (e.g. Smirnov, 2011; Braun, 2013). As such, the sidelobe level in SKA-low stations must be suppressed and/or accurately modelled. Figure 5 shows the average sidelobe levels found in simulations of quasi-random arrays. These simulations were done using the OSKAR2 aperture array simulation software (Dulwich et al., prep). For each data point, 100 quasi-random arrays of a fixed radius and filling factor (i.e. inter-antenna spacing) were generated and their beam characteristics were analysed to form statistical averages. The dashed blue line corresponds to a $y = c(N_A)^{\alpha}$ relationship, where α=-0.94 and $c = 0.80$.

A slope of $\alpha = -1$ is also expected from the following argument. At large off-axis angles the phase delays to each antenna can be considered to be essentially random, so the voltages from each antenna add incoherently, and the beamformed voltageis reduced by a factor $\sqrt{N_A}$ compared to boresight. The power pattern is proportional to voltage^2, hence the typical beam response at a given large angle scales as N_A^{-1}.

5.1 Beam size and the ionosphere

For calibration, it is advantageous if the beam size is smaller than an isoplanatic patch of the ionosphere; if the beam is larger, more complex methods must be applied to calibrate out the ionospheric effects (Cotton et al., 2004). Isoplanatic patches are of order 10 km in size, and the ionosphere is at a height of approximately 400 km, so the maximum FoV before ionospheric calibration becomes more complex is

$$\theta_I \approx 10/400\,\text{rad} = 1.4\,\text{deg}.$$

This would require very large stations at low frequencies (> 200 m at 50 MHz), and indeed this factor was one of the motivations for the large station size in SKA-low v1. SKA-low station sizes similar to those in SKA-low v2 will deliver beam sizes larger than an isoplanatic patch. This will require additional computational complexity in the post-correlation calibration and imaging, to allow for direction-dependent calibration of the ionospheric phases. The cost of this will need to be traded off against the real-time computational cost of the smaller stations.

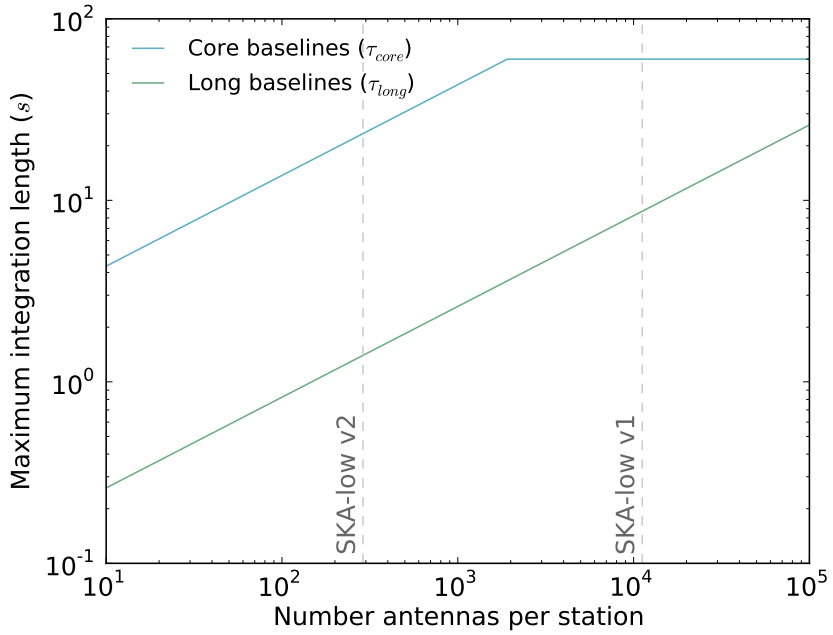

(a) Required minimum correlator dump times (i.e. maximum integration times), based on the specifications of SKA-low v2, for long baselines (top), and short (bottom).

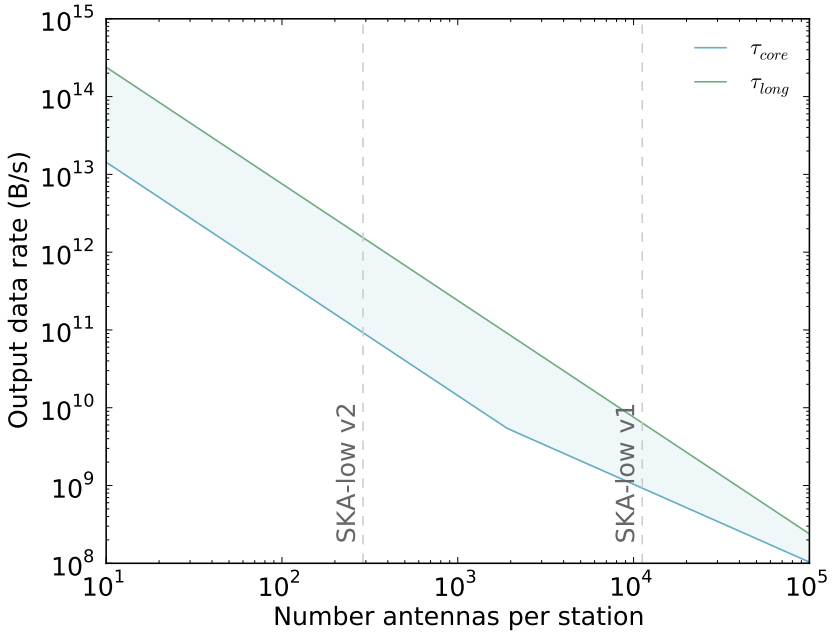

(b) Output data rate of SKA-low correlator, as a function of number of antennas per station. As the total data rate depends on the partition between long baselines and core baselines, the shaded area represents the possible range of output data rates.

Figure 4: Correlator maximum integration times (top) and corresponding output data rates (bottom).

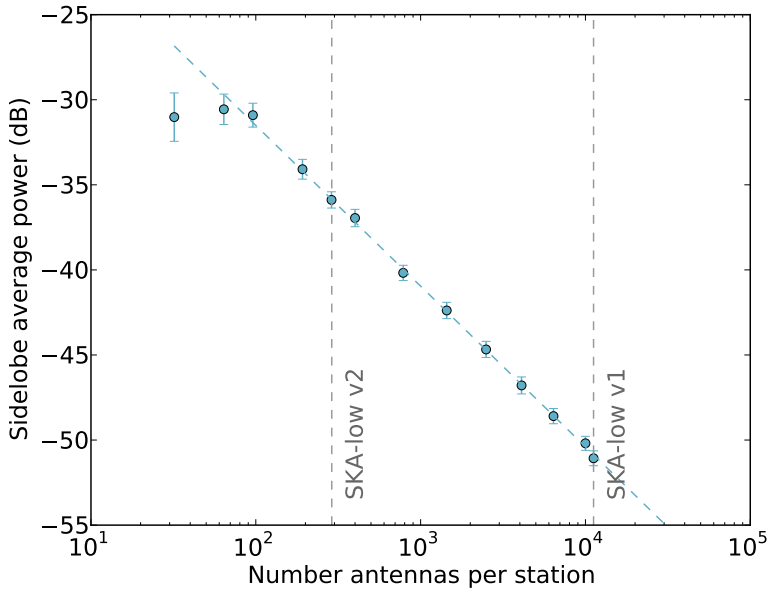

Figure 5: Average sidelobe level in decibels for a quasi-random aperture array station, shown as a function of antennas per station. As station size increases, sidelobe levels decrease. Data points correspond to average values of OSKAR simulations of quasi-randomly generated arrays; the dashed blue line shows a $y = c(N_A)^{\alpha}$ relationship, where α=-0.94 and $c = 0.80$.

6 Discussion and conclusions

Combining all these findings, we conclude that there is a compelling case to repartition the antennas within SKA-low into larger stations than the SKA-low v2 prescription. Importantly, SKA-low stations can be increased to as large as 85 m while still maintaining sensitivity to the longest wavelength mode of the EoR power spectrum. We have shown that larger stations may be computationally advantageous (depending on implementation), and that larger stations will produce lower data output rates. In addition, average sidelobe levels of quasi-random arrays decrease as stations increase in size, which is advantageous for high dynamic range imaging.

For example, consider a modest change to the baseline design: that the SKA-low station size is increased to 1024 antennas per station, for a total of 256 stations. This would result in station sizes of ∼65 m. The advantages of this implementation are as follows:

- For a comparable computational cost of the current baseline design, one could implement a multi-beam system (i.e. 2 beams or more). The overall FoV would be maintained (or increased), and the calibration advantages of a multi-beam system could be realized.
- The output data rate would decrease by a factor of ∼6.
- The average sidelobe level of each station would be reduced by ∼6 dB.

Further, we note that multiple, small independent FoVs are inherently more flexible and may be beneficial for non-EoR science cases.

In this memo, we have only discussed low-level characteristics of the SKA-low array. Ultimately, it is the scientific capability of SKA-low that must be optimized. Factors such as how well the array can be calibrated may place more stringent engineering requirements than those presented here. As such,

we suggest further investigation into the optimum partitioning of antenna elements; in particular, end-to-end simulations that include ionspheric effects will be vital to validate the EoR science capability of proposed SKA-low implementations.

References

Bowman, J, e. a. (2012). The first billion years. *Keck Institute for Space Sciences Study Report*.

Braun, R. (2013). Understanding synthesis imaging dynamic range. *Astronomy and Astrophysics*, 551, 91.

Colegate, T. M., Hall, P. J., & Gunst, A. W. (2012). Memo 140: Cost-effective aperture arrays for SKA Phase 1: single or dual-band? *SKA Memo Series*.

Cotton, W. D., Condon, J. J., Perley, R. A., et al. (2004). Beyond the isoplanatic patch in the VLA Low-frequency Sky Survey. *Ground-based Telescopes. Edited by Oschmann*, 5489, 180–189.

de Souza, L., Bunton, J. D., Campbell-Wilson, D., Cappallo, R. J., & Kincaid, B. (2007). A Radio Astronomy Correlator Optimized for the Xilinx Virtex-4 SX FPGA. In *Field Programmable Logic and Applications, 2007. FPL 2007. International Conference on* (pp. 62–67).

Dewdney, P. (2013). SKA1 System Baseline Design. *SKA Project Documents*, (pp. 1–98).

Dewdney, P., Vaate, J.-G. b. d., Cloete, K., et al. (2010). Memo 130: SKA Phase 1: Preliminary System Description. *SKA Memos*, (pp. 1–25).

Dickinson, C., Battye, R. A., Carreira, P., et al. (2004). High-sensitivity measurements of the cosmic microwave background power spectrum with the extended Very Small Array. *MNRAS*, 353, 732–746.

Dulwich, F., Mort, B., & Salvini, S. (2013, in prep). Using oskar to simulate data from radio interferometers. *MNRAS*.

Hickish, J. (2013). *PhD Thesis*. PhD thesis, University of Oxford.

Jiwani, A., Colegate, T., Razavi-Ghods, N., et al. (2012). Square Kilometre Array station configuration using two-stage beamforming. *arXiv.org*.

Jones, M. E., Salvini, S., Zarb Adami, K., et al. (2011). SKA Station Beamformer Concept Description. *SKA Project Documents*.

Kassim, N. E., Lazio, T. J. W., Erickson, W. C., et al. (2007). The 74MHz System on the Very Large Array. *arXiv.org*, (pp. 3088).

Mellema, G., Koopmans, L., Abdalla, F., et al. (2012). Reionization and the Cosmic Dawn with the Square Kilometre Array. *arXiv.org*.

Myers, S. T., Contaldi, C. R., Bond, J. R., et al. (2003). A Fast Gridded Method for the Estimation of the Power Spectrum of the Cosmic Microwave Background from Interferometer Data with Application to the Cosmic Background Imager. *ApJ*, 591, 575–598.

Pearson, T. J., Mason, B. S., Readhead, A. C. S., et al. (2003). The Anisotropy of the Microwave Background to l = 3500: Mosaic Observations with the Cosmic Background Imager. *ApJ*, 591, 556–574.

Smirnov, O. M. (2011). Revisiting the radio interferometer measurement equation. II. Calibration and direction-dependent effects. *arXiv*, astro-ph.IM.

Weinreb, S. & D'Addario, L. (2001). Memo 1: Cost Equation for the SKA. *SKA Memo Series*.

Zaroubi, S., Bruyn, A. G. d., Harker, G., et al. (2012). Imaging neutral hydrogen on large-scales during the Epoch of Reionization with LOFAR. *arXiv*, astro-ph.CO.