Precision Analysis of the Imaging Pipeline in the Square Kilometre Array

Anthony Griffin[1,4], Nicolas Pradel[1,2,5], Brody Radford[1,4], David I. Wilson[3,4], Andrew Ensor[1,6]

[1]High Performance Computing Research Laboratory, [2]Institute for Radio Astronomy and Space Research,
[3]Industrial Information and Control Centre, [4]Electrical and Electronic Engineering Department,
[5]Mathematical Sciences Department, [6]Computer Science Department,
School of Engineering, Computer and Mathematical Sciences
Auckland University of Technology, New Zealand

{anthony.griffin, nicolas.pradel, brody.radford, david.wilson, andrew.ensor}@aut.ac.nz

***Abstract*—In this paper we present our end-to-end model of the imaging pipeline in the Square Kilometre Array. Our Sky Generator models the signals that are received by the Central Signal Processor (CSP), our CSP Correlator model then processes those signals to generate visibilities to pass to the Science Data Processor (SDP). Our SDP Imaging model then grids the visibilities and inverse Fourier transforms them to produce a dirty image of the sky. Our modelling allows us to investigate the error that is introduced due to reduced numerical precision, and we then propose techniques to mitigate this error, and thus reduce the required amount of computational hardware.**

***Index Terms*—Radio astronomy, Square Kilometre Array, numerical precision, signal processing, modelling**

I. Introduction

Since the first successful measurements of the diameter of the star Betelgeuse in 1921 by Michelson and Pease [1], the technology of interferometers has made enormous advancements and extraordinarily ambitious projects have been commissioned, with more in the pipeline. Examples of recently commissioned interferometers are the Very Large Telescope Interferometer (VLTI) [2] for optical astronomy, and the Atacama Large (sub-)Millimetre Array (ALMA) [3] for radio astronomy.

Interferometers are more widely used in the radio spectrum than at other wavelengths, mostly because common electronics can more easily handle their lower frequencies. Since the discovery of cosmic radio emission, radio astronomical interferometers have been built to observe it from decametre to micrometre wavelengths, and range in size from a few metres to hundreds of thousands of kilometres with space telescopes HALCA [4] and RadioAstron [5]. However, due to the sparse distribution of the receivers, the reconstruction of the observed part of the celestial sphere is non-trivial and can only be really approximated through signal processing.

The Square Kilometre Array (SKA) [6] is easily the most ambitious radio interferometric project currently under way and will observe metre-to-centimetre wavelengths at unprecedented sensitivity. However, even in its first phase, this sensitivity comes with the cost of a large number of array elements (512 receiver stations), as well as narrow spectral channels (226 Hz), a large bandwidth (2.5 GHz) and short integration time (0.25 s). All of these parameters are part of the science requirements of the SKA and equate to a tremendous computational cost. The overall operations per second necessary to meet the SKA requirements is estimated well above the TOP500 in terms of computing alone, and much worse when the data flow is considered, as all the antennas will produce a combined raw data flow of approximately 159 TB/s.

Processing in the SKA will be done in the Central Signal Processor (CSP) and the Science Data Processor (SDP). The SDP will perform the final stages of image processing, which has traditionally been done in double-precision arithmetic in radio astronomy. However, due to the scope of the SKA, using double-precision would add a significant cost. Indeed, it has been estimated that using single-precision in the SDP would save tens of millions of euro in computing hardware.

In this work, we present our modelling work that aims to prove that single-precision will be sufficient for the gridding and FFT operations in the SDP that are the most compute-intensive parts of the radio astronomy imaging pipeline.

II. Radio Astronomy Imaging

A radio interferometer is made up of an array of $N \geq 2$ antennas. The maximum angular resolution of a single dish with a diameter D observing a signal of a wavelength λ is determined by its diffraction limit and is approximately $\frac{\lambda}{D}$. On the other hand, the angular resolution of the interferometer is limited by the distance between the furthest two antennas in the array $B_{\max}$ in the same way, that is $\frac{\lambda}{B_{\max}}$. Achieving sub-arcsecond resolution in single-dish antennas requires large diameters that are prohibitive (*e.g.*, a primary reflector larger than 2 km for $\lambda = 1$ cm observation). For interferometers, achieving sub-arcsecond resolution is just a matter of having antennas as far away as possible from each other.

The basic measurement device (the Michelson interferometer) is composed of just two elements. Each possible two-element combination of antennas forms a two-element measuring device. During an observation, each antenna tracks a position on the celestial sphere normally within the field of view of the observation. Each antenna receives the signal

ISBN 978-0-9928626-7-1

$$V(u,v,w) = \frac{1}{A_0}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\frac{A(l,m)I(l,m)e^{-2j\pi(ul+vm+w(\sqrt{1-l^2-m^2}-1))}}{\sqrt{1-l^2-m^2}}dldm \quad (1)$$

as time-series voltages, which are then synchronised in a correlator, cross-multiplied, and accumulated to improve the signal-to-noise ratio, as the signal of interest is generally much lower than the ambient noise. These accumulated values are called visibilities. Each visibility is an algebraic complex function and a measurement of the intensity distribution of the sky as seen by the element array. The $N(N-1)/2$ visibility readings are done simultaneously by the whole N-element antenna array and each reading differs in the vectorial geometrical set-up of the antennas and the radio sources. The geometrical set-up is described by the baseline which is the vector distance between the two antennas. Earth rotation changes the directions of the baseline (assuming a frame of reference stationary with the celestial sphere), and this is taken advantage of by taking subsequent visibility readings. As shown in [7], there is a Fourier relationship between visibility as a function of baseline vector and the intensity distribution, *i.e.*, the image of the sky.

Many variables in interferometry are expressed against a reference direction known as the phase reference centre or phase centre. The phase centre is fixed to the sky and it is common practice that it is set to point in the same direction as the delay centre. The general coordinate system used in interferometry is the (u,v,w) coordinate system. It is a right-handed Cartesian system where axes U and V are on a plane normal to the phase centre and the W-axis in the direction of the phase centre. The U-axis is in the East-West direction while the V-axis is in the North-South direction. One of the main uses of the coordinate system is to measure the baseline against the phase centre. The baseline components expressed over the UVW-axes are defined by (u,v,w). The components are normally given in units of number of wavelengths, λ. The measurement equation of the interferometer defines the relationship between the interferometer-measured quantities $V(u,v,w)$ and the intensity distribution of the sky $I(l,m)$ multiplied by the antenna response $A(l,m)$. It is defined in the most general case as (1) (see the top of this page).

The visibilities obtained from the correlator are samples of the measurement equation for the given (u,v,w) coordinate of the baseline vector. In order to obtain the intensity distribution of the sky, it is necessary to Fourier invert (1). Whilst possible in the generalised form given above, it is conceptually difficult and computationally expensive. If the visibility equation (1) can be reduced to the form of a two-dimensional Fourier transform (necessitating the removal of the w term), the inversion can be carried out using a Fast Fourier Transform FFT [8], which is much more computationally tractable. For most radio interferometry, this w term can be ignored as it is negligible if the field of view is small. However, in the case of SKA, the field of view is several degrees wide, so some strategies such as w-projection [9], w-stacking [10] and/or w-snapshot [11] are used, but are outside of the scope of this work. Once the w term is corrected, (1) can be reduced to

$$V(u,v,w) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}A(l,m)I(l,m)e^{-2j\pi(ul+vm)}dldm, \quad (2)$$

which leads to the inverse Fourier function

$$\frac{1}{A_0}A(l,m)I(l,m) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}V(u,v)e^{2j\pi(ul+vm)}dldm. \quad (3)$$

The FFT is the most efficient algorithm to solve this inversion, but requires a regularly-spaced sampled function. However, the value of the visibilities are non-regularly sampled continuous (u,v) coordinates. To use the FFT, the values of the visibilities need to be evaluated on a regularly-spaced grid by a so-called gridding kernel. Several gridding kernels exist but the most commonly used is the prolate spheroidal function [12].

III. Models of the signal processing chain

We developed a novel suite of models that comprise a full end-to-end model that takes a parametrised description of an area of sky with semi-realistic stellar sources, and after simulating an observation produces a corresponding intensity map $I(l,m)$. These are the Sky Generator Model (SGM), CSP Correlator Model (CCM), and the SDP Imaging Model (SIM), depicted in Figs. 1 and 2. The SGM and CCM were largely implemented in Simulink for its natural ability to succinctly model digital signal processing systems, while SIM was written in Matlab. We consider a single polarisation and a single coarse channel, as X and Y polarisations and adjacent frequency channels are considered to be essentially orthogonal. While parts of the complete radio astronomy signal processing pipeline have been modelled by others, we believe ours to be the first operational end-to-end model in existence.

A. Sky Generator Model

The SGM models the locations of 40 SKA Low antennas (a substantial subset of the full array of $\sim$ 500) and their geometrical relationship with synthetic astronomical sources. It generates the astronomical signals that would be received by the SKA Low array with its particular antenna loci, source direction (declination and right ascension) and intensities, and time of observation. The model also generates the appropriate time-varying geometric time delays that are central to the interferometer's operation and used by the CCM delay correction modules to synchronise the signals received by each antenna. This capability allows us to verify the overall model as it provides sufficiently realistic signals for processing by the CCM and SIM to reconstruct the source image. We considered using real data from an existing, different, radio telescope array, but determined it was important to model the layout

ISBN 978-0-9928626-7-1 © EURASIP 2017

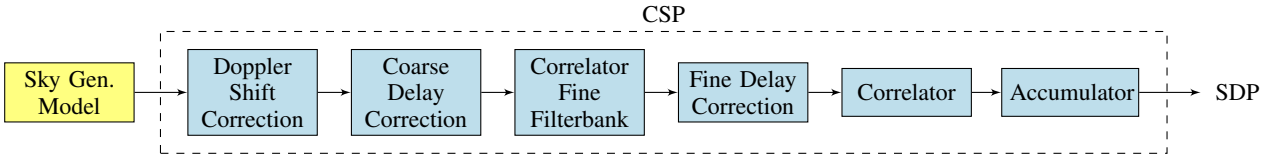

Fig. 1. The CSP correlator model.

of the SKA as it will be constructed and to be able to perform simulations with arbitrary source configurations.

An important function of SGM is to produce signals that are delayed a fraction of the sample time. These are known as sub-sample, or fractional, delays and the generation, and subsequent correction by the observatory is key to the correct operation of the telescope. In the elementary case of a pure sinusoid, this may be achieved analytically, however SGM's synthesis mechanism follows [13] which proposed approximating a band-limited noisy signal with a chorus of sine waves with closely spaced random frequencies and phases. A simple Gaussian random number generator is not used as one must be able to accurately compute fractional delays.

B. CSP Correlator Model

The CCM executes the CSP functions in Fig. 1, transforming unprocessed and incoherent radio wave samples into visibilities. This section only describes the imaging pipeline of the model, and does not include elements such as the radio frequency identification and mitigation components that do not affect the precision of the model.

- **Doppler Shift Correction** As the Earth rotates, Doppler shifts are introduced, modulating the frequency at each channel by a small, yet non-trivial, amount. The amount of Doppler shift is different for each antenna; some will be moving toward the source while others are moving away as the source passes through zenith.
  The original incoming signal s_i is Doppler-corrected to $\hat{s}_i$, by applying a complex time-varying phase-shift particular to each antenna:
  $$\hat{s}_i = s_i \, e^{-j\omega_c \, d_i(t)} \tag{4}$$
  where ω_c is the angular frequency of the centre of the coarse channel, and $d_i(t)$ is the time-varying delay for signal i.
- **Coarse Delay Correction** For the Low CSP stage, the incoming data are sampled at 1.08μs (926 kHz). The Coarse Delay Correction module aligns all antenna feeds to within ± 0.5 samples referenced to the centre of the array. This is achieved by applying the integer component of the delay, z^{-n}, appropriate to each antenna signal. In practice these delays are provided by the telescope manager given that they are known *a priori* for a particular observation, although in this model the Sky Generator provides these parameters.
- **Correlator/Fine Filterbank** The filterbank is a channeliser, performing frequency analysis with a polyphase filterbank architecture as described in [14] and [15]. The fine filterbank performs the second and final step of frequency analysis, refining each of the 384 $\times$ 926 kHz coarse channels into 4,096 $\times$ 226 Hz fine channels.
- **Fine Delay Correction** The fine delay correction removes the potential signal misalignment of up to ± 0.5 sample after the coarse delay is removed. This delay equates to $c \times \frac{T_s}{2} = 162$ m, many times the antenna dish diameter. This delay is removed by applying a complex phase-shift in the Fourier domain to the channelised data
  $$\hat{s}_i = s_i \, e^{-j \, d_i(t+\Delta t)} \tag{5}$$
- **Correlator**.
  The cross-correlation of two discrete functions x and y is
  $$x_n \star y_n = \sum_{m=-\infty}^{\infty} x_m^* \cdot y_{m+n} \tag{6}$$
  and from the Correlation Theorem we know that:
  $$\mathcal{F}\{x \star y\} = \mathcal{F}\{x\}^* \cdot \mathcal{F}\{y\} \tag{7}$$
  The Filterbank performs the Fourier analysis, and the Correlator performs the complex element-wise vector multiplication $\mathcal{F}\{x\}^* \cdot \mathcal{F}\{y\}$ for all baselines, all pairs of antennas, all coarse channels, and all fine channels, emitting complex visibilities for each baseline. These visibilities are a measure of the common astronomical signal detected by each baseline.
- **Accumulator**. The Accumulator integrates visibilities up to the dump time, 0.25 s. This is a straightforward yet crucial step, as the astronomical signals are far below the noise and thus must be accumulated for a relatively long time to increase the signal-to-noise ratio; approximately 57 times per dump. Accumulated visibilities are sent to the SDP.

C. SDP Imaging Model

Our SDP Imaging Model (SIM) was written in Matlab, and is based on ASKAPsoft [16], which is the C++ software that the Australian Square Kilometre Array Pathfinder [17] runs. The incoming visibilities are gridded onto a regular grid using prolate spheroidal functions, and then converted to the image domain using the IFFT, and thus implementing (3). This 'dirty image' is then deconvolved to produce a cleaned, and then restored image. In this work we only consider the dirty image.

Although we have flexibility in the precision used in each stage of our model, we also have a mode that follows the precision used in ASKAPsoft, allowing us to reproduce the exact output of ASKAPsoft, giving us confidence in the accuracy of our Matlab implementation.

ISBN 978-0-9928626-7-1 © EURASIP 2017

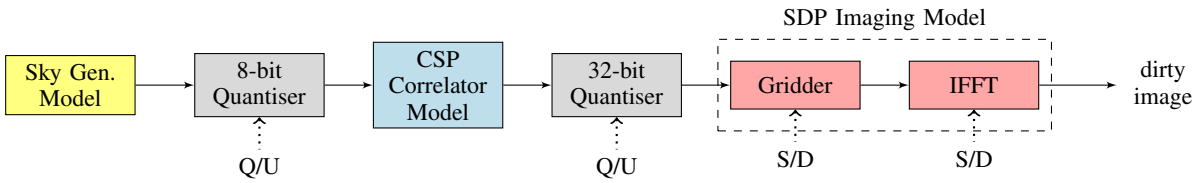

Fig. 2. Full system model, with precision options. Q/U and S/D denote quantised or unquantised, and single- or double-precision, respectively.

TABLE I
PRECISION COMBINATIONS STUDIED

legend label	P
8-bit	QUDD
32-bit	UQDD
gridding	UUSD
IFFT	UUDS

D. Full System Model

Our full system model is shown in Fig. 2. The current SKA design specifies that the input to the CCM will be quantised to 8-bit integers with a standard deviation of 12, and that the output of the CCM will be quantised to 32 bits.

The precision options in Fig. 2 allow us to refer to a combination of precision options as a four-letter code. For example the highest-precision case is UUDD, which means:

- U no quantisation between the SGM and CCM
- U no quantisation between the CCM and Gridder
- D the Gridder runs at double-precision
- D the IFFT runs at double-precision,

and the lowest-precision case is QQSS, which means:

- Q 8-bit quantisation between the SGM and CCM
- Q 32-bit quantisation between the CCM and Gridder
- S the Gridder runs at single-precision
- S the IFFT runs at single-precision.

This allows us to study the isolated precision effects of different blocks in the dirty images produced by various precision combinations. Table I shows the combinations and the labels we use to refer to them.

Note that for this work, the SGM operates in double precision (as it is approximating a natural phenomenon). The CCM also operates in double precision, although parallel investigations have developed CCM model variants that also run at varying precision levels.

IV. Results and Discussion

We used our Sky Generator Model to generate the signals that would be received by 40 receivers over a five minute observation period. The receiver locations were chosen as a subset of the 512 planned for SKA Low. These (possibly quantised) sky signals were then processed by the CSP Correlator Model as 780 baselines, resulting in 250,000 visibilities. These were possibly further quantised, before being processed by the SDP Imaging Model at various precisions to produce dirty images.

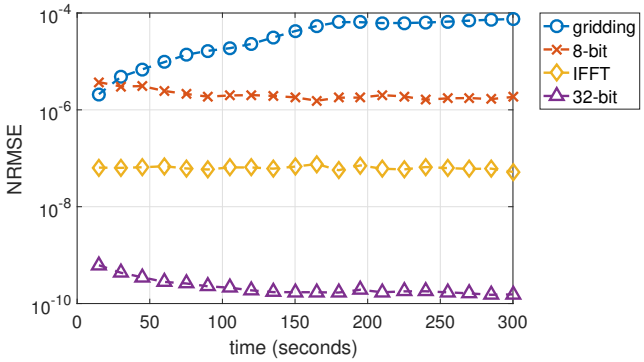

Fig. 3. Isolated precision errors for 40 receivers, and a 256^2 grid.

We calculated the error due to reduced precision in the following manner. Let $\boldsymbol{X}_P$ be the dirty image obtained with precision combination P, then we measure the normalised root-mean-square error (NRMSE) between the precision combination P and a reference P_{ref} as

$$\mathrm{E}(P, P_{\text{ref}}) = \frac{||\boldsymbol{X}_P - \boldsymbol{X}_{P_{\text{ref}}}||_2}{||\boldsymbol{X}_{P_{\text{ref}}}||_2} \tag{8}$$

where $||\cdot||_2$ denotes the ℓ_2 norm. See Table I for the precision combinations used. In this work, P_{ref} is always the most precise computation, UUDD.

Fig. 3 presents the isolated errors due to reduced precision, plotted over a five minute interval. All the errors reduce or remain flat except for that due to gridding. Indeed, the gridding is the main source of error, and approaches 1×10^{-4}. The error due to the 8-bit quantisation is just over 1×10^{-6}, while the IFFT error is less than 1×10^{-7}. The error due to the 32-bit quantisation is orders of magnitude down, well below 1×10^{-9}, and will be ignored in the rest of the results.

A common source of error in single-precision processing is that a very small number is added to a very large number, resulting in no change to the large number due to the limitations of single-precision representation. We hypothesised that this was the cause of the high gridding error, as many visibilities could be added up together on the same grid point. To test this hypothesis, we redid the results in Fig. 3 on larger grids. These results are shown in Fig. 4, where it is clear that the increasing grid size reduces the gridding error significantly. It is intuitively satisfying that the 8-bit quantisation and the IFFT errors are unaffected by the grid size. This is also the case for the 32-bit quantisation, which remains below 1×10^{-9} for all grid sizes, but for clarity is not shown.

ISBN 978-0-9928626-7-1 © EURASIP 2017

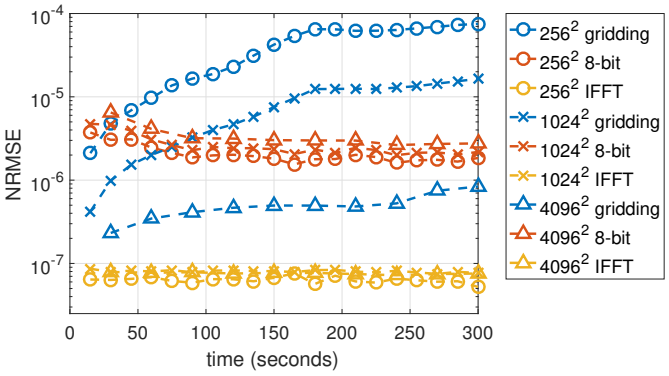

Fig. 4. Isolated precision errors for 40 receivers with various grid sizes.

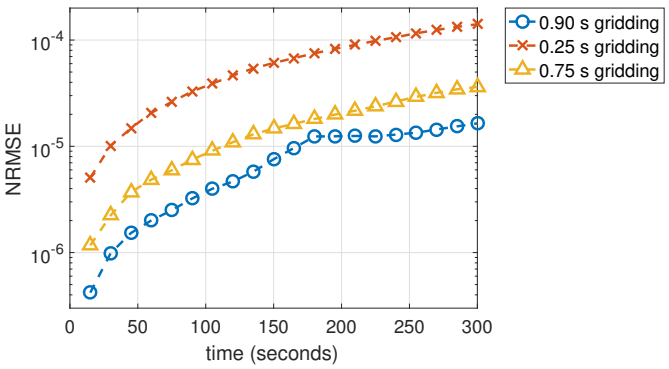

Fig. 5. Isolated precision errors for 40 receivers with 1024^2 grids and various integration times.

During the course of this work, there was a change to the SKA requirement governing the integration time in the CSP, from 0.90 seconds down to 0.25 seconds. The effect of this requirement change is that the visibilities arrive at the gridder about 3.6 times faster. We used our model to investigate the effect of this change on the error due to reduced precision. The results are shown in Fig. 5. It is clear that the decrease in integration time from 0.90 s to 0.25 s results in an order-of-magnitude increase in the gridding error, due to more visibilities being added to a grid point. The IFFT and 8-bit quantisation errors are unaffected by the decrease in integration time, and for clarity are not shown here.

In order to mitigate this increase in gridding error, we implemented a simple extension to our gridder, where we pre-summed three 0.25 s visibilities before they were gridded. This result is shown in Fig. 5 as the 0.75 s curves. It is clear that this simple addition decreases the error significantly, paving the way for more sophisticated schemes such as baseline-dependent averaging (BDA) [18]. BDA involves averaging the visibilities so that shorter baselines have longer integration times, and is usually used to reduce the time spent on the gridding process, but the results here suggest that it could also be used to decrease the single-precision error significantly.

In the SKA, the integration time in the output of the CCM is fixed and cannot be varied according to baseline length. However, as shown above we can simply accumulate them further at the input to the SDP before they are gridded, and this could easily be done according to baseline length.

V. Conclusion

In this paper we have presented our precision modelling work of the Imaging Pipeline in the SKA, along with some initial results. These initial results are very encouraging, and indicate that the IFFT may be performed in single-precision as the error in doing so is less than that due to the 8-bit quantisation. Single-precision gridding introduces a more significant error, but our results strongly suggest that this error may be reduced by performing averaging of the visibilities in the SDP before gridding them.

References

[1] A. A. Michelson and F. G. Pease, "Measurement of the diameter of Alpha Orionis with the interferometer." *Astrophysical Journal*, vol. 53, pp. 249–259, 1921.

[2] R. G. Petrov, F. Malbet, A. Richichi *et al.*, "AMBER: the near-infrared focal instrument for the very large telescope interferometer," in *Astronomical Telescopes and Instrumentation*. International Society for Optics and Photonics, 2000, pp. 68–79.

[3] A. Wootten, "Atacama large millimeter array (ALMA)," in *Astronomical Telescopes and Instrumentation*. International Society for Optics and Photonics, 2003, pp. 110–118.

[4] H. Hirabayashi *et al.*, "The VLBI space observatory programme and the radio-astronomical satellite HALCA," *Publications of the Astronomical Society of Japan*, vol. 52, no. 6, p. 955, 2000.

[5] N. Kardashev, "RadioAstron-a radio telescope much greater than the Earth," *Experimental Astronomy*, vol. 7, no. 4, pp. 329–343, 1997.

[6] P. J. Hall, *The Square Kilometre Array: An International Engineering Perspective*. Dordrecht: Springer Netherlands, 2005, pp. 5–16.

[7] A. Thompson, J. Moran, and G. Swenson, *Interferometry and Synthesis in Radio Astronomy*. Wiley, 2008.

[8] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex Fourier series," *Mathematics of Computation*, vol. 19, pp. 297–301, 1965.

[9] T. J. Cornwell, K. Golap, and S. Bhatnagar, "The noncoplanar baselines effect in radio interferometry: The w-projection algorithm," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 5, pp. 647–657, 2008.

[10] P. Dewdney, W. Turner, R. Millenaar, R. McCool, J. Lazio, and T. Cornwell, "SKA1 system baseline design," *Document number SKA-TEL-SKO-DD-001 Revision*, vol. 1, no. 1, 2013.

[11] T. Cornwell, M. Voronkov, and B. Humphreys, "Wide field imaging for the square kilometre array," in *SPIE Optical Engineering+ Applications*. International Society for Optics and Photonics, 2012, pp. 85 000L–85 000L.

[12] D. Slepian and H. O. Pollak, "Prolate spheroidal wave functions, Fourier analysis and uncertainty I," *Bell System Technical Journal*, vol. 40, no. 1, pp. 43–63, 1961.

[13] V. Paxson, "Fast, Approximate Synthesis of Fractional Gaussian Noise for Generating Self-Similar Network Traffic," *Computer Communication*, vol. 27, no. 5, pp. 5–18, 1997.

[14] C. Harris and K. Haines, "A mathematical review of polyphase filterbank implementations for radio astronomy," *Publications of the Astronomical Society of Australia*, vol. 28, no. 4, pp. 317–322, 001 2011.

[15] D. C. Price, "Spectrometers and Polyphase Filterbanks in Radio Astronomy," no. 1, pp. 1–21, 2016.

[16] CSIRO Australia, "ASKAP science data processor software," http://doi.org/10.4225/08/583f6c9316254 [Online; accessed 12-February-2017].

[17] ——, "ASKAP home," http://www.atnf.csiro.au/projects/askap/index.html [Online; accessed 12-February-2017].

[18] W. D. Cotton, "Special Problems in Imaging," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, and R. A. Perley, Eds., vol. 180, 1999, p. 357.

ISBN 978-0-9928626-7-1 © EURASIP 2017