



SDP Memo 62: Numerical Precision

Document Number SDP Memo 62
 Document Type MEMO
 Revision 1
 Author Anthony Griffin and Andrew Ensor
 Release Date 2018-09-06
 Document Classification Unrestricted
 Status Draft

Lead Author	Designation	Affiliation
Anthony Griffin	Imaging Pipeline Lead	New Zealand Alliance — Auckland University of Technology
Signature & Date:	<i>ABGriffin</i>	2018-09-06

Revision	Date of issue	Prepared by	Comments
1	2018-09-06	Anthony Griffin	Initial Release

SDP Memo Disclaimer

The SDP memos are designed to allow the quick recording of investigations and research done by members of the SDP. They are also designed to raise questions about parts of the SDP design or SDP process. The contents of a memo may be the opinion of the author, not the whole of the SDP.

Table of Contents

1 Introduction	5
2 Radio Astronomy Imaging	5
3 Models of the signal processing chain	5
3.1 Sky Generator Model	6
3.2 CSP Correlator Model	6
3.3 SDP Imaging Model	8
4 Simulations	8
4.1 Dirty Image Precision Effects	8
4.2 Minor Cycle Cleaning Precision Effects	11
4.3 Major Cycle Cleaning Precision Effects	14
5 Conclusion	18
List of Figures	19
List of Tables	19
References	20

List of Abbreviations

ASKAP	Australian Square Kilometre Array Pathfinder
BDA	Baseline Dependent Averaging
CCM	CSP Correlator Model
CSP	Central Signal Processor
FFT	Fast Fourier Transform
FOV	Field of View
IFFT	Inverse Fast Fourier Transform
NRMSE	Normalised Root-Mean-Square Error
SDP	Science Data Processor
SGM	Sky Generator Model
SIM	SDP Imaging Model
SNR	Signal-to-Noise Ratio

1 Introduction

This memo is a summary of all the SDP numerical precision work that has been performed at AUT. The goal is to determine whether or not single precision will provide acceptable performance in the Imaging Pipeline of the SKA. We consider not only a fully single precision pipeline, but also a heterogeneous one where some processing is performed in single precision, and the rest is done in double precision. The motivation for this work is that using single precision instead of double would result in significant cost savings for the SDP (in the order of tens of millions of Euro). Significant portions of this work have been published in (Griffin et al., 2017) and (Griffin and Ensor, 2018).

2 Radio Astronomy Imaging

The *Measurement Equation*

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

gives the relationship a measured visibility $V(u, v, w)$ at a position (u, v, w) determined by a baseline has with the sky intensity $I(l, m)$ and the antenna response $A(l, m)$ at an angle (l, m) on the sky. If the field of view is small then the w term can be largely ignored, and the equation simplifies to a Fourier transform relationship between the visibilities and the image on the sky (Thompson et al., 2008).

In the case of the SKA, the field-of-view is several degrees wide, and strategies such as w-projection (Cornwell et al., 2008), w-stacking (Dewdney et al., 2013) and/or w-snapshot (Cornwell et al., 2012) are used, to correct for the w term. Once corrected, and the antenna response found via calibration, the measurement equation can be inverted to give the intensity distribution $I(l, m)$ on the sky as a function of the measured visibilities $V(u, v)$:

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

The inverse Fourier transform given by (2) could be efficiently evaluated by an inverse two-dimensional Fast Fourier Transform (FFT) (Cooley and Tukey, 1965), if the visibilities had been measured at regularly spaced coordinates (u, v) . Unfortunately the baselines result in visibilities along arcs, so to apply the FFT, the visibilities must be approximated on a regularly-spaced grid by convolution with a *gridding kernel*. Several gridding kernels exist but the most commonly used in radio astronomy has been the prolate spheroidal function (Slepian and Pollak, 1961), as it maximizes the energy concentration in the main lobe (Barbosa, 1986).

3 Models of the signal processing chain

Our novel suite of models comprise an end-to-end system, taking a parametrised description of a region of sky containing semi-realistic stellar sources and—after simulating the observation processing—producing the corresponding intensity map $I(l, m)$ as described in (2). They are:

- the Sky Generator Model (SGM),
- the CSP Correlator Model (CCM),
- the SDP Imaging Model (SIM),

as depicted in Figure 1. The SGM and CCM were implemented in Simulink—due its ability to succinctly model digital signal processing systems—and the SIM was written in Matlab, as discussed later. We simulate only a single polarisation and single coarse channel, as we consider X and Y polarisations and adjacent frequency channels to be essentially orthogonal. Parts of the complete radio astronomy signal processing chain have been modelled previously, but we believe that ours is the first full end-to-end model, and the first to accurately model the intended SKA signal processing chain.




Figure 1: The three models making up the end-to-end model of the Imaging Pipeline in the SKA.

3.1 Sky Generator Model

The Sky Generator Model (SGM) is the name we gave to the model that generates the signals that the correlator in the CSP receives. The SGM models the locations of a subset of the SKA Low antennas 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 approximately 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 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 (Paxson, 1997) 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.

3.2 CSP Correlator Model

The CSP Correlator Model (CCM) executes the CSP functions in Figure 2, transforming unprocessed and incoherent radio wave samples into visibilities. This section only describes the parts of the CSP model that are relevant to 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.




Figure 2: The CSP correlator 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:

$$\hat{s}_i = s_i e^{-j\omega_c d_i(t)} \quad (3)$$

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 $1.08\mu\text{s}$ (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 (Harris and Haines, 2011) and (Price, 2016). The fine filterbank performs the second and final step of frequency analysis, refining each of the 384×926 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 $c \times \frac{T_c}{2} = 162$ 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

$$\hat{s}_i = s_i e^{-j d_i(t+\Delta t)} \quad (4)$$

- **Correlator.**

The cross-correlation of two discrete functions x and y is

$$x_n \star y_n = \sum_{m=-\infty}^{\infty} x_m^* \cdot y_{m+n} \quad (5)$$

and from the Correlation Theorem we know that:

$$\mathcal{F}\{x \star y\} = \mathcal{F}\{x\}^* \cdot \mathcal{F}\{y\} \quad (6)$$

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.

3.3 SDP Imaging Model

In order to ensure that we implemented the final part of the Imaging Pipeline correctly, we based the SDP Imaging Model on (SIM) ASKAPsoft (CSIRO Australia, b), the software running on the Australian Square Kilometre Array Pathfinder (ASKAP) (CSIRO Australia, a). ASKAPsoft was written in C++, but we used Matlab for the SIM as it allowed for deeper analysis and meshed better with our other models.

As discussed earlier, the visibilities generated by the CCM must be gridded onto an evenly-spaced grid using prolate spheroidal functions. This grid is then passed through an IFFT to generate an image as expressed in (2). Through a process called deconvolution, this “dirty image” is then cleaned to produce the final image of the sky.

In addition being able to vary the precision used throughout the stages of our full model, we have a mode that mimics the precision in ASKAPsoft, allowing us to reproduce the bit-accurate output of ASKAPsoft, enabling us to be very confident in our Matlab implementation’s accuracy. It must be noted here that the generation of the prolate spheroidal function (Slepian and Pollak, 1961) is always done in double precision as it very sensitive to numerical precision. However this is not a big constraint, as it is only calculated intermittently.

4 Simulations

4.1 Dirty Image Precision Effects

Our initial investigations focused on the degradation in the dirty image due to reduced numerical precision. The system model used here is shown in Figure 3. The current SKA design specifies that the input to the CCM will be quantised as 8 (real) + 8 (imaginary) bits (SKA1 LFAA to CSP Interface Control Document) and the output of the CCM will be quantised to 32 (real) + 32 (imaginary) bits (.). The standard deviation of the input is assumed to be 12. Note that the LFAA tile beamformers, station beamformers, and coarse channelization are not included in the model, but the CCM does include the fine channelization filter bank as specified in the CSP_Low design. Although we have labelled the block between the CCM and the SIM as a 32-bit quantiser, here it chooses between single- and double-precision floating point.

The precision options in Figure 3 allow us to refer to a combination of precision options as a four-letter code. For example the highest-precision case is UDD, 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,




Figure 3: Dirty Image System Model with precision options. Q/U and S/D denote quantised or unquantised, and single- or double-precision, respectively.

Table 1: Precision combinations for dirty image effects

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

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 1 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.

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.

We calculated the error due to reduced precision in the following manner. Let 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

$$E(P, P_{\text{ref}}) = \frac{\|X_P - X_{P_{\text{ref}}}\|_2}{\|X_{P_{\text{ref}}}\|_2} \quad (7)$$

where $\|\cdot\|_2$ denotes the ℓ_2 norm. See Table 1 for the precision combinations used. In this work, P_{ref} is always the most precise computation, UUDD.

Figure 4 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




Figure 4: Isolated precision errors for 40 receivers, and a 256^2 grid.

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 Figure 4 on larger grids. These results are shown in Figure 5, 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.

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 gridded 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 Figure 6. 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 gridded, where we pre-summed three 0.25 s visibilities before they were gridded. This result is shown in Figure 6 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) (Cotton, 1999). 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




Figure 5: Isolated precision errors for 40 receivers with various grid sizes.




Figure 6: Isolated precision errors for 40 receivers with 1024^2 grids and various integration times.

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.

4.2 Minor Cycle Cleaning Precision Effects

The next part of this work looked at how errors introduced through reduced precision affected deconvolution, specifically the Högbom CLEAN algorithm (Högbom, 1974). The system model for this is shown in Figure 7

As discussed above, in this work we only investigate the effect of precision choices in the SIM, and the SGM and CCM operate in double-precision. The precision options for the SIM in Figure 1 allow us to refer to a precision combination with a three-letter code. For example, DDD is the highest-precision case, where all the blocks are run in double-precision; and the lowest-precision case is SSS, where all the blocks are run in single-precision. If just the gridded, IFFT, or deconvolution are run in single-precision the codes are SDD, DSD, and




Figure 7: System model, with precision options, where S and D denote single- or double-precision, respectively.




Figure 8: Zoomed in view of the seven sources observation, the white box shows the area used for dynamic range calculations.

DDS, respectively. There is a step in the deconvolution called the convolution correction and this is always performed in double precision, however this is not very computationally intensive.

The first simulation we did in this section was to model seven sources arranged in a spiral, as shown in Figure 8. The Sky Generator Model was used to generate the signals from these sources that would be received by 20 receivers over a five-minute observation period. The receiver locations are a subset of the 512 that have been chosen for SKA Low. Differing values of thermal noise were added at each receiver so that the signal-to-noise ratio (SNR) varied from 0 dB to -60 dB. These sky signals were then processed by the CSP Correlator Model as 190 baselines, over 41 fine frequency channels, resulting in over 9 million visibilities. The SDP Imaging Model then gridded these visibilities, passed them through an IFFT, and performed deconvolution using 1000 cycles of the Högbom CLEAN, as implemented by ASKAPsoft. Figure 8(a) shows the dirty image, which is the output of the IFFT, and Figure 8(b) shows the image after deconvolution where the seven sources can be seen clearly. Note that these are 200×200 pixel zoomed views of the 1024×1024 pixels image of the sky.

A common figure of merit in astronomical images is that of dynamic range, which is defined as the peak intensity divided by the root mean squared energy in a patch of the image where there are no sources (Thompson et al., 2008). The choice of patch can affect the dynamic range measured, especially in complex images. In order to be sure we did this fairly, we measured the dynamic range using a range of patches. The starting position of the patch is shown in Figure 8(c), and the patch was then moved in 23 steps towards the top of the image until it just touched the edge of the cleaned inner quarter—a 512×512 area.

The measured values are shown in Figure 9, where it is clear the higher values of SNR




Figure 9: Dynamic Range measured over a series of patches that moved further and further away from the centre of the image, for various values of signal-to-noise ratio (SNR).

give better dynamic range, which is intuitively satisfying. In fact, only the images with an SNR ≥ -30 dB show the sources clearly, for lower values of SNR the sources cannot be resolved. Figure 9 also illustrates that the dynamic range rises slightly as the patch moves further away from the sources, which is to be expected.

By keeping the patch fixed, we can investigate the effects of precision choices in the SIM. Figure 10 illustrates this, and it is evident that the use of single-precision throughout the SIM makes no practical difference to measured dynamic range.

Although it is not obvious in Figure 10, from closer examination of these results it was clear that the greatest effect from using single precision came from using single-precision in the gridded. This was also a major finding of the work in Section 4.1. Thus from this point on, we only compared results using SDD against DDD.

We also investigated the effect of distance between two sources. We positioned a weaker source near the centre of the field of view, and then a stronger source at the same distance from the centre but diagonally opposite the weak source. We then moved the stronger source further and further out along the diagonal line until it was at the edge of the cleaned quarter. Let us now define γ as the flux ratio,

$$\gamma = 10 \log_{10} \frac{I_W}{I_S}, \quad (8)$$

where I_W is the intensity of the weak source, and I_S is the intensity of the strong source. Thus if $\gamma = -6$ dB, then $I_W \approx I_S/4$. Now let γ_T be the true flux ratio, and γ_M be the *measured* flux ratio, the ratio that we calculate from our cleaned image. Obviously, in the ideal case, $\gamma_M = \gamma_T$. Figure 11 shows the results of our simulations which were performed with an SNR of 0 dB. For the case with $\gamma_T = -6$ dB, $\gamma_M \approx -6$ dB at all positions of the strong source, but this is the only case where this is true. When $\gamma_T = -9$ dB, γ_M falls away as the strong source moves further away, and the weak source is not even detected when the strong source is at the 15-th position. This trend increases as γ_T decreases, with γ_M being less accurate and the weak source not being detected with the strong source closer and closer in. In terms of precision, the only case where using single-precision results in a significant difference from double-precision processing is when $\gamma_T = -18$ dB. But it must be noted that the SDD case always detects the weak source when the DDD case does, it is only that γ_M differs. These results are consistent with (Wijnholds et al., 2016) in that as the strong source moves further away from the phase centre, the phase error will increase leading to reduced accuracy in estimation of the weak source.




Figure 10: Dynamic Range over a fixed patch for various values of SNR and precision. See Figure 7 for more details on the precision modes.




Figure 11: Measured Flux Ratio γ_M for various source positions, values of True Flux Ratio γ_T and precision modes. See Figure 7 for more details on the precision modes. Note that when a curve ends, it means that the weak source was no longer detected.

4.3 Major Cycle Cleaning Precision Effects

The major cycles of ASKAPsoft were then added to the SIM, as shown in Figure 12. In our simulations we chose to mirror the precision of gridding and degriding and the IFFT and FFT, so that the precision combinations and their labels are explained in Table 2. The precision combination SDS is the precision combination used in ASKAPsoft. Similar to Section 4.2, the convolution correction is always performed in double precision.

Our first simulation in this section was run with ASKAPsoft's 1934 test data, which contains over 600,000 visibilities (233 x 15 x 180), representing one point source near the centre of the field-of-view (FOV). 500 minor cycles of Högbom cleaning were performed with a loop gain of

Table 2: Precision combinations for major cycle effects

legend label	gridding and degriding	IFFT and FFT	deconvolution
SSS	single	single	single
SDS	single	double	single
DDD	double	double	double

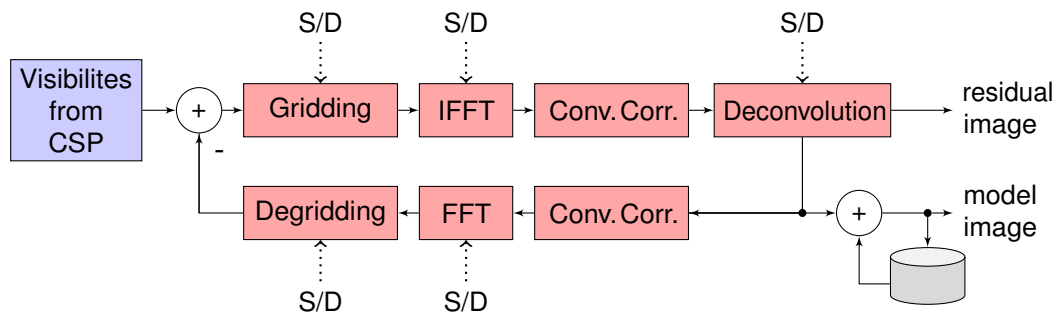


Figure 12: Full system model, with precision options, where S and D denote single- or double-precision, respectively.




Figure 13: Dynamic Range measured over a series of patches that moved further and further away from the centre of the FOV, with 500 minor cycles in each major cycle, and DDD precision.

0.1, along with 5 major cycles. The Dynamic Range was measured over a 50 x 50 pixel patch that was directly above the point source and then moved towards the top of the FOV.

Figure 13 shows the Dynamic Range of the dirty image and the restored image after each major cycle, for DDD precision. It is intuitively satisfying that the Dynamic Range increases with each major cycle for most patches. Similarly, Figure 14 shows the difference of using 100 minor cycles instead of 500 in each major cycle. Again, it is intuitively satisfying that the Dynamic Range increases with more minor cycles for most patches.

Figures 15, 16 and 17 shows the Dynamic Range of the dirty image, the restored image after one major cycle and the restored image after five major cycles, respectively. It is clear that there is very little difference in DR when comparing the different precision combinations.




Figure 14: Dynamic Range measured over a series of patches that moved further and further away from the centre of the FOV, with 5 major cycles, and the given number of minor cycles in each major cycle, and DDD precision.

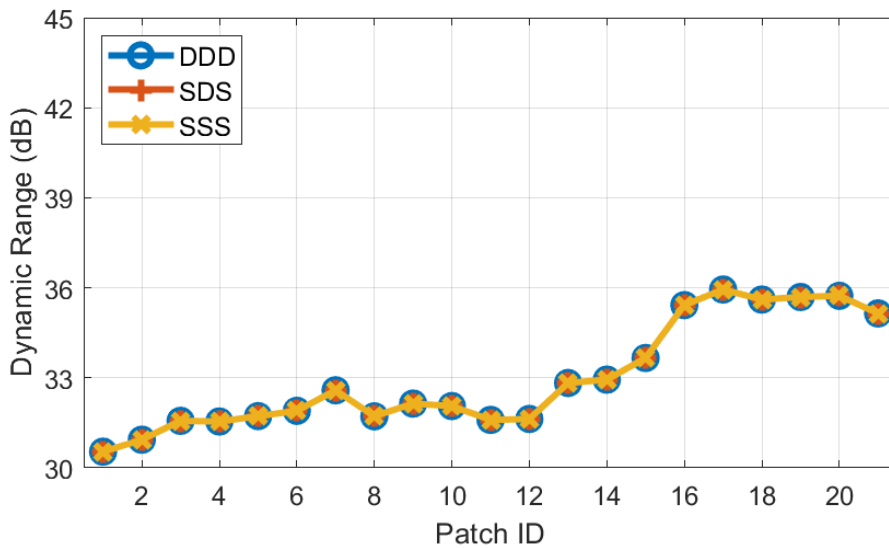


Figure 15: Dynamic Range measured over a series of patches that moved further and further away from the centre of the image, measured on the dirty image, for different precision combinations.




Figure 16: Dynamic Range measured over a series of patches that moved further and further away from the centre of the image, measured on the restored image, with 500 minor cycles, 1 major cycle, for different precision combinations.

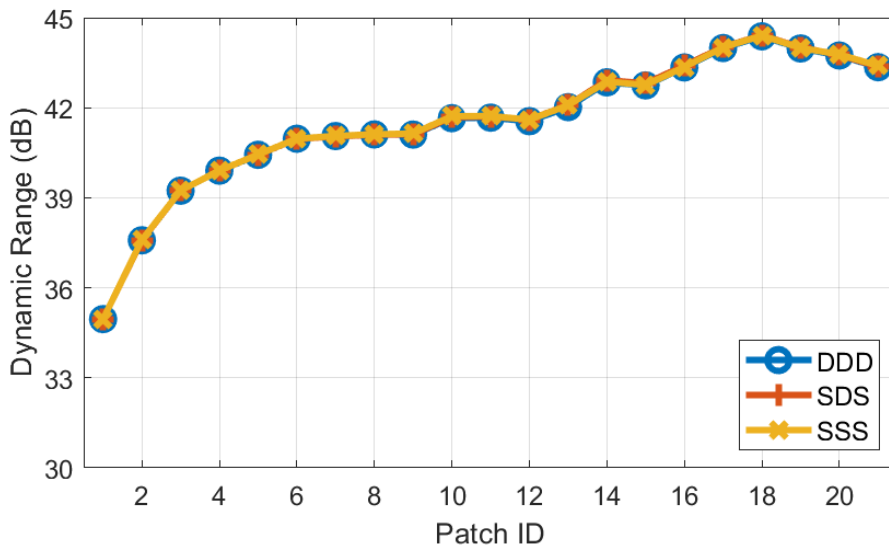


Figure 17: Dynamic Range measured over a series of patches that moved further and further away from the centre of the image, measured on the restored image, with 500 minor cycles, 5 major cycles, for different precision combinations.

5 Conclusion

Our end-to-end models of the Imaging Pipeline in the SKA simulate the signals received from sources in the sky and produce a deconvolved image of that region of the sky in agreement with ASKAPsoft. Although reduced precision increases the NRMSE in the dirty image, this is not significant when measuring the Dynamic Range in the cleaned and restored images. These results are encouraging, and indicate that single-precision processing may be suitable for a large part of the imaging pipeline, resulting in significant savings in computing hardware.

List of Figures

1	The three models making up the end-to-end model of the Imaging Pipeline in the SKA.	6
2	The CSP correlator model.	7
3	Dirty Image System Model with precision options. Q/U and S/D denote quantised or unquantised, and single- or double-precision, respectively.	9
4	Isolated precision errors for 40 receivers, and a 256^2 grid.	10
5	Isolated precision errors for 40 receivers with various grid sizes.	11
6	Isolated precision errors for 40 receivers with 1024^2 grids and various integration times.	11
7	System model, with precision options, where S and D denote single- or double-precision, respectively.	12
8	Zoomed in view of the seven sources observation, the white box shows the area used for dynamic range calculations.	12
9	Dynamic Range measured over a series of patches that moved further and further away from the centre of the image, for various values of signal-to-noise ratio (SNR).	13
10	Dynamic Range over a fixed patch for various values of SNR and precision. See Figure 7 for more details on the precision modes.	14
11	Measured Flux Ratio γ_M for various source positions, values of True Flux Ratio γ_T and precision modes. See Figure 7 for more details on the precision modes. Note that when a curve ends, it means that the weak source was no longer detected.	14
12	Full system model, with precision options, where S and D denote single- or double-precision, respectively.	15
13	Dynamic Range measured over a series of patches that moved further and further away from the centre of the FOV, with 500 minor cycles in each major cycle, and DDD precision.	15
14	Dynamic Range measured over a series of patches that moved further and further away from the centre of the FOV, with 5 major cycles, and the given number of minor cycles in each major cycle, and DDD precision.	16
15	Dynamic Range measured over a series of patches that moved further and further away from the centre of the image, measured on the dirty image, for different precision combinations.	16
16	Dynamic Range measured over a series of patches that moved further and further away from the centre of the image, measured on the restored image, with 500 minor cycles, 1 major cycle, for different precision combinations.	17
17	Dynamic Range measured over a series of patches that moved further and further away from the centre of the image, measured on the restored image, with 500 minor cycles, 5 major cycles, for different precision combinations.	17

List of Tables

1	Precision combinations for dirty image effects	9
2	Precision combinations for major cycle effects	14

References

- Barbosa, L. C.: A maximum-energy-concentration spectral window, *IBM Journal of Research and Development*, 30, 321–325, doi:10.1147/rd.303.0321, 1986.
- Cooley, J. W. and Tukey, J. W.: An algorithm for the machine calculation of complex Fourier series, *Mathematics of Computation*, 19, 297–301, 1965.
- Cornwell, T., Voronkov, M., and Humphreys, B.: Wide field imaging for the Square Kilometre Array, in: *SPIE Optical Engineering+ Applications*, pp. 85 000L–85 000L, International Society for Optics and Photonics, 2012.
- Cornwell, T. J., Golap, K., and Bhatnagar, S.: The noncoplanar baselines effect in radio interferometry: The W-projection algorithm, *IEEE Journal of Selected Topics in Signal Processing*, 2, 647–657, 2008.
- Cotton, W. D.: Special Problems in Imaging, in: *Synthesis Imaging in Radio Astronomy II*, edited by Taylor, G. B., Carilli, C. L., and Perley, R. A., vol. 180 of *Astronomical Society of the Pacific Conference Series*, p. 357, 1999.
- CSIRO Australia: ASKAP Home, <http://www.atnf.csiro.au/projects/askap/index.html> [Online; accessed 14-March-2018], a.
- CSIRO Australia: ASKAP Science Data Processor software, <http://doi.org/10.4225/08/583f6c9316254> [Online; accessed 14-March-2018], b.
- Dewdney, P., Turner, W., Millenaar, R., McCool, R., Lazio, J., and Cornwell, T.: SKA1 system baseline design, Document number SKA-TEL-SKO-DD-001 Revision, 1, 2013.
- Griffin, A. and Ensor, A.: End-to-end Modelling of the Imaging Pipeline in Radio Astronomy, in: *IEEE Sensor Array And Multichannel Signal Processing Workshop (SAM)*, 2018.
- Griffin, A., Pradel, N., Radford, B., Wilson, D. I., and Ensor, A.: Precision analysis of the imaging pipeline in the square kilometre array, in: *European Signal Processing Conference (EUSIPCO)*, pp. 2418–2422, doi:10.23919/EUSIPCO.2017.8081644, 2017.
- Harris, C. and Haines, K.: A Mathematical Review of Polyphase Filterbank Implementations for Radio Astronomy, *Publications of the Astronomical Society of Australia*, 28, 317–322, doi:10.1071/AS11032, 2011.
- Högbom, J. A.: Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines, *Astronomy and Astrophysics Supplement Series*, 15, 417–426, <http://adsabs.harvard.edu/abs/1974A%26AS...15..417H>, 1974.
- Paxson, V.: Fast, Approximate Synthesis of Fractional Gaussian Noise for Generating Self-Similar Network Traffic, *Computer Communication*, 27, 5–18, doi:10.1145/269790.269792, 1997.
- Price, D. C.: Spectrometers and Polyphase Filterbanks in Radio Astronomy, pp. 1–21, 2016.
- Slepian, D. and Pollak, H. O.: Prolate Spheroidal Wave Functions, Fourier Analysis and Uncertainty — I, *Bell System Technical Journal*, 40, 43–63, doi:10.1002/j.1538-7305.1961.tb03976.x, 1961.

Thompson, A., Moran, J., and Swenson, G.: Interferometry and Synthesis in Radio Astronomy, Wiley, 2008.

Wijnholds, S. J., Salvini, S., Dodson, R., Mort, B., Dulwich, F., Willis, A., and Still, J.: SDP Memo: Feasibility analysis of baseline-dependent averaging, Tech. Rep. SKA-TEL-SDP-0000017, SKA Science Data Processor Consortium, <http://ska-sdp.org/publications/released-sdp-memos>, 2016.