Widefield Science and Technology for the SKA
SKADS Conference 2009
S.A. Torchinsky, A. van Ardenne, T. van den Brink-Havinga, A.J.J. van Es, A.J. Faulkner (eds.)
4-6 November 2009, Château de Limelette, Belgium

OSKAR: Simulating Digital Beamforming for the SKA Aperture Array

F. Dulwich[1], B.J. Mort[1], S. Salvini[1], K. Zarb Adami[2], and M.E. Jones[2]

[1] Oxford e-Research Centre, University of Oxford
[2] Department of Astrophyics, University of Oxford

Abstract. Digital beamforming for the aperture array components of the SKA poses considerable computational challenges. In this paper, we propose a hierarchical scheme aimed at tackling them and introduce OSKAR, a beamforming simulator which implements these ideas and algorithms.

1. Introduction & Motivation

The numerous advantages of aperture arrays, from rapid survey speeds to multiple simultaneous beam pointings, make this detector technology central to the SKA design for the low to mid-frequencies. The work presented here describes beamforming simulations for the mid-frequency SKA station aperture array, investigates the challenges involved, and how the proposed solutions affect the performance and beam quality.

The simulations are based on the frequency-domain, all-digital beamforming design as currently envisaged (Schilizzi et al. 2007): this technology is most likely to scale well with advances in processing hardware.

We first introduce the computational issues involved in beamforming for an SKA-size aperture array, and outline a flexible scheme that could tackle them. Finally, we describe OSKAR, our beamforming simulator, and present examples of how this simulator is used.

2. Computational Challenges

The reference design for the mid-frequency aperture array of an SKA station consists of 65,536 dual-polarisation antenna elements, each sampling the sky at a rate of ~2 GHz to provide 1024 narrow frequency channels spanning several hundred MHz of total bandwidth.

Beams are formed electronically by the linear combination of delayed antenna signals. For a narrowband digital system these delays correspond to complex phase factors, or weights. Beams can then be formed by a simple linear transformation:

$$b_i = \sum_{j=1}^{N} W_{i,j} a_j \tag{1}$$

where b_i is a vector of beams for each direction i, $W_{i,j}$ is the matrix of complex phase weights, and a_j is the vector of channelised, sampled antenna signals. In their simplest form, the phase weights are given by $W = e^{i\psi}$, where ψ is the geometric phase offset required to form a beam in a direction defined by azimuth ϕ and elevation θ, for wavenumber k and an antenna at coordinates (x, y):

$$\psi = k \cos\theta \, (x \sin\phi + y \cos\phi) \tag{2}$$

Equation 1 represents a Fourier relation between directions on the sky, defined by the geometric phase, and a set of antenna positions. The operation count of the Fast Fourier Transform (FFT) makes it initially attractive, but its suitability for beamforming must be considered carefully. In fact, a spatial FFT requires a regular grid of antennas, or a computationally expensive gridding procedure to deal with irregular array configurations. The FFT also lacks flexibility if we consider that forming beams may require some level of anisotropic beam corrections, for example, if on-line calibration or null steering is employed. Furthermore, the FFT computes *all* the independent beams for a given frequency and polarisation, although only a subset of them may be practically used. With this approach, the computational cost of forming beams for one frequency channel would be ~ 20 TFlops, rising to ~ 20 PFlops for all channels. This does not include the extra costs that the FFT scheme is likely to incur due to gridding and beam corrections: the actual cost may be far higher.

With each antenna generating ~ 2 MWords/s for each frequency channel and polarisation, and as many independent beams directions as the number of antennas, the computational challenges are extreme, both in terms of overall bandwidth and number of operations.

An alternative approach to forming the required beams is by direct matrix-vector product, as in Equation 1. This approach, while computationally more intensive, offers a far higher level of flexibility over the FFT: no gridding is required for irregular array configurations; the beamforming weights can be directly modified to include calibration corrections and optional beam shaping; and only the beam directions required for the observation need to be computed at all. This allows the total data rate to be reduced as early as possible. Using this method, the cost of forming a beam for a given channel is ~ 1 TFlop, and this scales linearly with the number of required beams.

With these extreme computational requirements in mind, we developed a beamforming scheme for maximum flexibility, simple computational structure and fast switching between op-

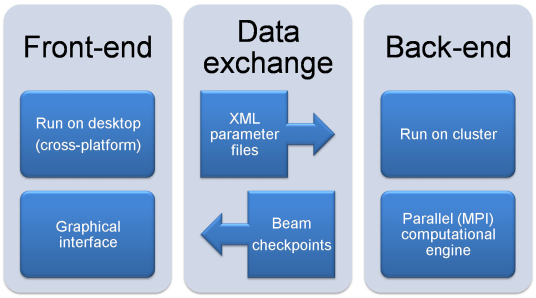

Fig. 1: Overall structure of the OSKAR simulator. The front-end and the back-end are independent components, which exchange data via a file interface.

erational modes. For a given observation, we assume that only a subset of the beams, frequencies or polarisations are needed, and the matrix-vector approach allows us to maintain this flexibility.

In order to maximise parallelism and minimise non-local data transfer as early as possible, we adopt a hierarchical beamforming strategy. In this scheme the aperture array is divided into 'tiles', or sets of contiguous antennas, which each produce tile beams. These tile beams are then combined to form station beams at a second level of processing. Of course, more than two levels could be envisaged.

The beamforming weight coefficients used depend on the geometrical phase terms for the set of pointing directions, as well as any required corrections to the beam. Since the parameters that determine these weights (e.g. beam tracking, and variations in antenna gains) vary much slower than the signal sampling rate, the weights can be computed off the main signal path and at relatively low cost.

For a two-level beamformer, the computational costs of an FFT scheme would require 5 TFlops per beam per frequency channel. Our scheme of explicit multi-level linear combinations of weighted antenna signals would require 1 TFlop per beam per frequency channel. However, costs increase non-linearly. If all the beams and frequencies were required, the FFT scheme would need $\sim$ 5 PFlops, while the matrix-vector scheme would need $\sim$ 250 PFlops.

Overall, the total amount of data that can be generated far exceeds current and, most likely, medium-term future capabilities. It is therefore very important to be able to easily switch between different operational modes to select only the beams and frequencies of interest.

3. The OSKAR Simulator

OSKAR simulates digital beamforming for the aperture array of an SKA-sized station. The simulator has been designed as a flexible framework to investigate different solutions to the computational challenges of the beamformer processing, and the effects on the quality of the beams produced.

The simulator consists of two main components (Figure 1). A front-end graphical interface is used to set up the simulation parameters, including antenna placement, sky model, and beam positions. The front-end uses platform-independent libraries and is designed to run on the user's desktop or workstation. The beamforming simulation itself is performed using a modular, MPI-parallel computational back-end, designed to run on a supercomputing cluster. Modules are provided for sky processing, antenna signal generation, beamforming and complex weights generation, as shown in Figure 2. Data can be check-pointed at any stage of the computational pipeline for use in other simulations if required, and the modular design of OSKAR means that it is easy to experiment with different beamforming algorithms by writing new weights generators. This also means that it is possible to experiment with alternative hardware architectures by intercepting the data stream.

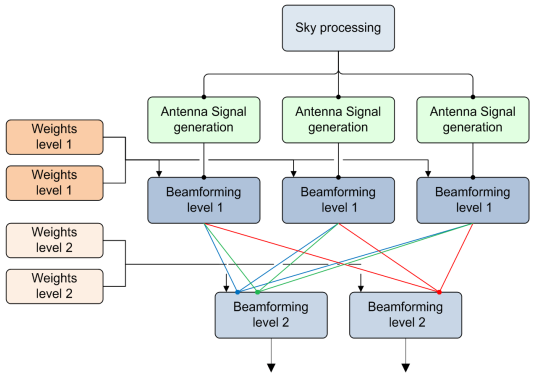

Fig. 2: An example configuration showing the modular structure and communication pattern of the OSKAR back-end, for a two-level system.

3.1. Simulator Implementation

The antenna signal modules produce a stream of complex-valued dual-polarisation signals for each antenna using the sky model. The weights generators produce sets of complex beamforming weights to define the direction of interest for each beam. Beam steering is achieved by updating the weight coefficients, but as these vary much slower than the signal sample rate, they can be computed away from the main signal path. This scheme allows the computationally-intensive beamforming modules to be very simple, as they only perform multiply-accumulate operations between the antenna signals and complex weights. For reasons of flexibility, we have adopted the matrix-vector approach in the OSKAR simulator, but other beamforming algorithms could be easily added to OSKAR, including the use of the FFT.

Parallelism is achieved by splitting the station into tiles, consisting of order 10^2 antennas, and by adopting a hierarchical scheme whereby beams are produced by successively splitting those from the previous level. The hierarchical approach has the advantage of allowing the beamformer to discard unwanted data as early as possible, thus reducing the data rate through the system.

OSKAR has two modes of operation: 'beam-pattern' mode and 'end-to-end beamforming simulation' mode.

The 'beam-pattern' mode is used to produce a single snapshot of the directional beam response for the given array configuration. There is no time dependence in this mode, so the beamforming weights are held constant for the direction of interest while a test source of unit amplitude is swept around the model sky to evaluate the station response at each position. This mode is suitable for investigating the effects of beam shaping (such as applying apodisation, or by forcing null steering), and to investigate using different array configurations and/or changing the antenna types.

The other mode of operation is the 'end-to-end beamforming simulation' mode, which uses a point-source sky model to generate antenna signals for each element of the aperture array. This mode simulates a realistic, full-size SKA station, so it is possible to model beam tracking, multiple beams per level, and to incorporate noise and other time-variable effects. This mode is well-suited for investigating the effects of using a different processing hierarchy, and experimenting with different beamforming algorithms while studying their performance.

4. Examples

This section describes some simple examples obtained from the OSKAR simulator using model isotropic antennas. Visualisations were performed using the OSKAR front-end.

4.1. Simple two-level beam pattern

The simplest example consists of a regular square station containing 16^2 close-packed tiles, each of which is a dense array of 16^2 antennas separated by $\lambda/2$, where $\lambda = 0.30$m is the wavelength of the radiation. The station thus contains 65,536 antennas. Each tile produces a beam pointing straight up at the zenith, and the main lobe of the tile beam is shown in Figure 3a: as expected, this is the Fourier transform of a square aperture.

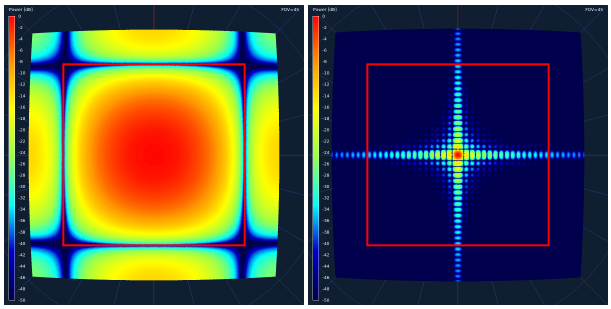

Fig. 3: Beam patterns shown using a logarithmic scale. **Left (a.):** Main lobe of the tile beam at the zenith, 15 degrees across. The red box indicates the position of the first null. **Right (b.):** The corresponding station beam, produced at the tile beam centre. Parts (a.) and (b.) are shown to the same scale.

These tile beams can be combined to form a station beam, also at the zenith in the first instance (Figure 3b). No grating lobes are visible in the station beam pattern, even though the tile centres are separated by distances greater than half a wavelength. In this case, the tile separation is the same as the tile width: the former determines the grating lobe separation in the station beam, while the latter determines the separation of the tile beam nulls. The station beam is modulated by the tile beam pattern, so the grating lobes coincide with the null positions, and are entirely suppressed.

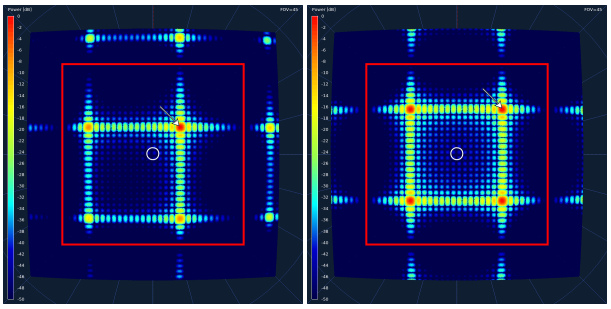

Fig. 4: Beam patterns showing the effects of forming station beams away from the tile beam centre, which is indicated by the white circle. The red boxes show the location of the first null in the tile beam. Note that the logarithmic colour scales are relative to the peaks in each image. **Left (a.):** Station beam produced at (45°, 87°) in azimuth and elevation, three degrees from the tile beam centre. **Right (b.):** Station beam produced at (45°, 85°) in azimuth and elevation, five degrees from the tile beam centre. The white arrows show the location of the main lobe of the station beam.

The pattern becomes more interesting when forming a station beam away from the tile beam centre. Figure 4a shows a station beam three degrees off the zenith. The main lobe of the station beam is at the expected location, but the grating lobes, previously suppressed, have reappeared: the station beam is now offset with respect to the tile beam centre. The station beam response has also dropped to ~ 75 per cent of its previous value because of the tapering on the tile beam at this position. This effect becomes more significant when the station and tile beams are further apart. At five degrees separation (Figure 4b), the main lobe has ~ 42 per cent of the central amplitude. At seven degrees, near the edge of the tile beam, the response in the direction of interest has dropped to ~ 15 per cent, while the most prominent grating lobe (about 10 degrees away) has a response of ~ 72 per cent of the tile beam amplitude.

This shows that, while the tile beam is 15 degrees across, only the central few degrees can be used to form station beams effectively. Station beams produced very close to the edge of the main lobe of tile beam are likely to introduce major errors due to the presence of the grating lobes.

4.2. More complex two-level beam pattern

Figure 5 shows a more complex beam pattern. This example used a smaller station containing 256 square tiles arranged in a circle, separated by 0.07m. Each tile contains 4^2 antennas separated by 0.15m, and the signals from each antenna were combined with Gaussian-distributed random phase errors (σ = 20 degrees). The tile beam in this example is at 60 degrees elevation, and the station beam is at the tile beam centre. The

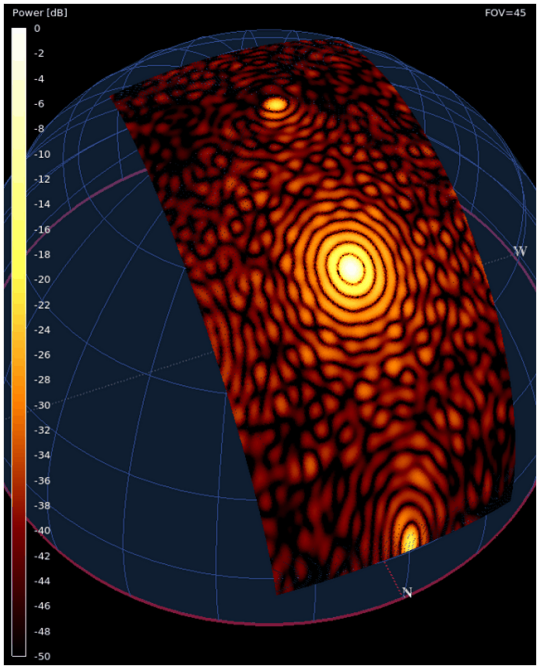

Fig. 5: Complex beam pattern as described in Section 4.2.

circular symmetry of the station is evident from the circular sidelobes, while the phase errors show up as noise in the beam pattern, causing a reduction in the station beam sensitivity by 5 per cent. Low-level grating lobes are also visible because the tiles are not close-packed.

4.3. End-to-end simulation mode

Multiple simultaneous beam snapshots can be formed in full beamforming ('end-to-end') simulation mode. Figure 6 shows a mosaic of beams formed around one source at low elevation, where each pixel is a single station beam formed by splitting a grid of tile beams separated by three degrees. As expected, the source appears elongated in the vertical direction, and the elliptical sidelobe pattern arises from the circular symmetry of the station. The power in the beam pixels to the right of the source is due to the tile beam sidelobe. The patchwork pattern arises as a result of amplitude scaling by the tile beams, which are shown in Figure 7.

5. Conclusions & Future Work

At the outset, all-digital beamforming for an SKA station aperture array poses considerable computational difficulties. However, the scalable, flexible framework that we have proposed would allow the SKA key observational goals to be achieved. This is a practical alternative to the overwhelming challenge of generating all beams at all frequencies for each observation.

We hope to continue the development of OSKAR to extend its functionality. Current plans include: alternative weights generation algorithms, such as those providing null-steering capabilities; a full treatment of polarization; simulation of ionospheric corruptions; and techniques for station level calibration. The next stage of SKA design may well require full interferometer simulations, and the beams produced by OSKAR could be well suited for these studies.

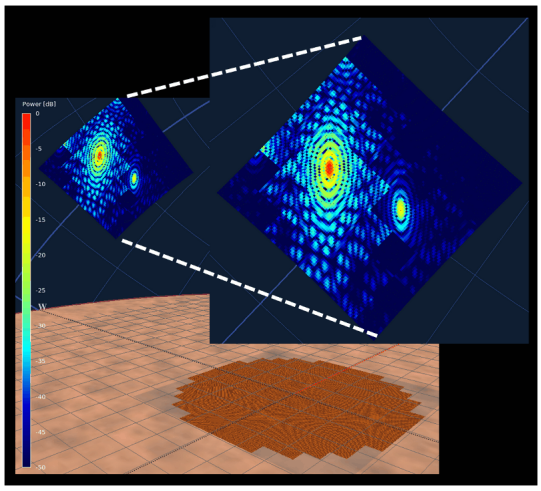

Fig. 6: Snapshot of a 20 by 20 degree station beam mosaic from an end-to-end simulation. The patchwork pattern is due to amplitude scaling by the tile beams, as shown in Figure 7.

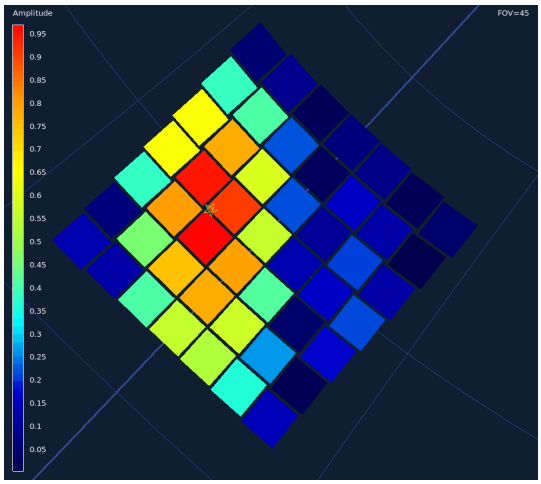

Fig. 7: The tile beam mosaic for the patch of sky shown in Figure 6.

OSKAR was developed under a BSD open-source licence. The first version has been released and is available with a documented API from `http://www.oerc.ox.ac.uk/research/oskar`.

Acknowledgements. We gratefully acknowledge the funding by EPSRC for the development of OSKAR. The project, carried out under the overall umbrella of SKADS activities, was completed in April 2009.

References

Faulkner, A. J. et al. 2007, "Design of an Aperture Phased Array System for the SKA," `www.skads-eu.org/p/`

memos.php
Schilizzi, R. T. et al. 2007, "Preliminary Specifications for the Square Kilometre Array," www.skatelescope.org/pages/memos