A Study of the Square Kilometre Array Low-Frequency Aperture Array

David Robert Sinclair
Wolfson College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy in Astrophysics

Trinity 2015

A Study of the Square Kilometre Array Low-Frequency Aperture Array

David Robert Sinclair

Wolfson College

A thesis submitted for the degree of Doctor of Philosophy in Trinity term 2015

This thesis primarily comprises simulations of SKA1-Low, the low-frequency aperture array of the Square Kilometre Array (SKA) radio telescope. Simulations assess how different design specifications would affect the telescope's performance and ability to undertake one of its key science projects: observing the Epoch of Reionisation (EoR). The simulations are based on the previous baseline design of $\sim 260\,000$ dipole antennas in 1024 stations of 256 antennas each.

It is found that errors in the gain and phase of signals at each antenna can significantly affect observations, with the greatest impact at the lowest frequencies. The EoR signal will be detectable at limited redshifts with gain errors of 10% and phase errors of 6°, however to ensure the systematic noise contribution from these errors is always weaker than the EoR signal between $z \sim 10 - 21$ and spherical harmonic multipole coefficients $\ell \sim 10^2 - 10^4$, the gain and phase errors may need to be $\lesssim 5 \times 10^{-6}$ and $\lesssim 3 \times 10^{-4}$ degs, respectively.

Alternative telescope layouts to the SKA1-Low baseline design of 1024 stations containing 256 antennas each are evaluated. Some telescope models which share antennas between stations, allowing stations to be formed in a flexible manner, are found to provide superior thermal noise sensitivity than the current design, but suffer greater detrimental effects from gain and phase errors when imaging.

The implementation of the digital signal processing of a station beamformer is also discussed. Calculations show that five or nine FPGAs with specifications comparable to the Virtex UltraScale+ XCVU13P could be used to beamform a single 256-antenna SKA1-Low station, depending on the utilisation rate of the FPGAs and the number of channels required.

Acknowledgements

This thesis would not exist without the help and support I have had. The most important person, of course, has been my supervisor Prof. Mike Jones, an endless source of knowledge, insightful comments and suggestions. High praise must also go to Ben and Fred for their indispensable assistance and boundless patience. And to Ashling, for always knowing what to do, no matter the hour or day.

Thanks to Danny*, Griffin, Jack, Adam and Charles for advice, encouragement, technical expertise, RadioLab and tapas. To Danielle, my office buddy, for always lending your ear and offering greatly valued wisdom on astrophysics and beyond. To Sarah, for knowing how to analyse actual interferometer data, your incredible work ethic and navigating this D.Phil journey in radio astronomy with me from the off. To Alex*, for your can-do enthusiasm, engineering acumen and for being around forever and yet still being in your first year. To Luke J* for your pretty pictures and eternal optimism on when dinner ('lunch') time should be (after 12:45 being the answer, of course), as well as knowing a thing or two about radio astronomy. To Thibaut*, for the cosmology, the wine and your charm.

Beyond the delights of the Tower and Denys Wilkinson Building, thank you to Hasneen and Liam for being sterling housemates. Thank you to all my wonderful friends in the Korfball Club. Thank you to Luke H, Joseph, Sam, Joana and Paul; without you I'm sure I would never have been in a position to start this D.Phil. And thank you especially to Camilla, you could have so many of the above descriptions attached to your name, but you also get the accolade of having helped solve the first problem I was stuck on in this D.Phil.

Thanks to everyone who has enriched my years here. You have all made it a thoroughly enjoyable experience.

*Bonus thanks to Danny, Alex, Luke J and Thibaut for devoting the time to read drafts of this thesis and provide feedback.

Contents

List of Figures

List of Tables

Abbreviations

2-PAD 2-Polarizations All Digital

AIPS Astronomical Image Processing System

ARTEMIS Advanced Radio Transient Event Monitor and Identification System

ASIC Application Specific Integrated Circuit

ASKAP Australian Square Kilometre Array Pathfinder

ATCA Australia Telescope Compact Array

BRAM Block RAM

C-BASS C-Band All Sky Survey

CASA Common Astronomy Software Applications

CASPER Collaboration for Astronomy Signal Processing and Electronics Research

cMAC Complex Multiply-Accumulate

CMB Cosmic Microwave Background

CPU Central Processing Unit

CUDA Compute Unified Device Architecture

DRAM Distributed RAM

DSP Digital Signal Processing

EoR Epoch of Reionisation

FAST Five-hundred-meter Aperture Spherical radio Telescope

FIR Finite Impulse Response

FFT Fast Fourier Transform

FPGA Field Programmable Gate Array

FRB Fast Radio Burst

GMRT Giant Metrewave Radio Telescope

GPU Graphics Processing Unit

HPBW Half-power beam width

I/O Input / Output

IGM Intergalactic medium

ISM Interstellar medium

LEDA Large-Aperture Experiment to Detect the Dark Ages

LNA Low-noise amplifier

LOFAR The Low-Frequency Array

LWA Long Wavelength Array

MAC Multiply-Accumulate

MERLIN Multi-Element Radio Linked Interferometer Network

MWA Murchison Widefield Array

OSKAR Oxford SKA Radio telescope Simulator

PAPER Precision Array for Probing the Epoch of Reionisation

PFB Polyphase Filter Bank

RAM Random Access Memory

RFI Radio-Frequency Interference

RIME Radio Interferometer Measurement Equation

RIN Residual In-field noise

RMS Root Mean Square

RSN Residual Sidelobe noise

SKA Square Kilometre Array

SKA1-Low The Phase 1 Square Kilometre Array Low-Frequency Aperture Array

SKA1-Mid The Phase 1 Square Kilometre Array Mid-Frequency Dish Array

VEGAS Versatile Green Bank Telescope Astronomical Spectrometer

VHDL VHSIC Hardware Description Language

VHSIC Very High Speed Integrated Circuit

VLA Very Large Array

VLBA Very Long Baseline Array

VLBI Very Long Baseline Interferometry

VLSS Very Large Array Low-Frequency Sky Survey

WSRT Westerbork Synthesis Radio Telescope

Chapter 1

Introduction

In the eight decades since Karl Jansky conducted the first astronomical measurements in the radio band, radio astronomy has become one of the most important means for studying the Universe. The field was responsible for some of the most important discoveries in 20th century astrophysics, such as the Cosmic Microwave Background and pulsars. Radio astronomy intends to continue this tradition with the development of the Square Kilometre Array (SKA).

The SKA is a radio telescope which will be constructed in South Africa and Australia. When completed, it will have orders of magnitude greater sensitivity and survey speed than current telescopes, operating over a wide frequency range of 50 MHz to 13.8 GHz. The capability of the SKA will expand the frontiers of radio astronomy and allow the exploration of an immense range of new science.

This chapter will introduce the principle concepts behind interferometry with aperture arrays and the SKA, along with an overview of historical and contemporary developments in the field.

1.1 History

The field of radio astronomy is regarded as commencing with Jansky's discovery in 1933 that the radio interference his array was detecting had a component with a sidereal period, meaning it was non-terrestrial in origin [1, 2]. Grote Reber followed Jansky's work by constructing the first steerable parabolic radio telescope in his back

garden in 1937, with which he produced the first radio map of the sky [3, 4]. Reber's map correctly showed the strong radio emission from the Milky Way.

This initial success had much for which to thank the field of radio communications. Jansky was studying the effect of atmospheric interference on radio transmissions when he made his discovery, whilst Reber was a radio engineer. This trend of harnessing skills and technology from similar fields continued in the aftermath of World War II, when scientists and engineers who had honed their skills on radar turned their attention to radio astronomy.

Ruby Payne-Scott in Sydney at CSIRO conducted the first interferometric observations using a radar system mounted atop a cliff looking out to sea; fringes, the sinusoidal interference pattern which is more typically the product of signals from two antennas, were detected from combining signals that took a direct path to the antenna with those which reflected off the sea [5, 6].

At the same time, Martin Ryle established a group in Cambridge that constructed the first multiple-element radio astronomy interferometer, with two radio dishes.

The existence of the 21-cm hydrogen line, one of the most important sources of radio emission in astronomy, was theorised in 1945 by Hendrick van de Hulst [7]. This prediction was proved right when Ewen and Purcell [8] detected the 21-cm line in 1951, closely followed by detections by Muller and Oort [9] and by Christiansen and Hindman [10, 11].

The 1960s proved perhaps the most successful decade so far, with three achievements being awarded with Nobel prizes. In 1962 Ryle fully described Earth-rotation synthesis mapping [12], whereby the rotation of the Earth changes interferometric baseline directions, and exploited it with the Cambridge One Mile Telescope [13]. The following year, Penzias and Wilson discovered the Cosmic Microwave Background [14] and in 1968 Bell and Hewish discovered pulsars [15].

The following decades saw the construction of several interferometers that significantly increased the sensitivity and resolution of observations. The Cambridge 5-km

telescope in 1972 [16] and the Westerbork Synthesis Radio Telescope (WSRT) in 1973 [17] were followed by the Very Large Array (VLA) in 1980 [18], the Australia Telescope Compact Array (ATCA) [19] in 1988 and the Giant Metrewave Radio Telescope (GMRT) in 1998 [20].

Another milestone was the development of interferometry with continental-scale baselines in 1979 [21], known as Very Large Baseline Interferometry (VLBI).

The modern era has seen the construction of a new generation of telescopes, such as the MeerKAT [22], the Australian Square Kilometre Array Pathfinder (ASKAP) [23], the Murchison Widefield Array (MWA) [24] and the Low-Frequency Telescope (LOFAR) [25]. The next stage will be the construction of the SKA.

1.2 Radio telescopes

Radio telescopes can be either individual or a collection of antennas. Telescopes are typically parabolic dishes, cylindrical paraboloids or aperture arrays.

1.2.1 Parabolic dish

A parabolic dish radio telescope is a two-dimensional (2-D) surface which reflects radio signals to focus them on to a receiver. Reber's classical design mounted the receiver at a centrally-located prime focus. Alternative designs include the offset Gregorian design, which place the receiver offset from the dish centre and change the prime focus position to support the receiver from the edge of the dish, and the Cassegrain design which has a secondary reflector to focus waves to a receiver behind the parabolic dish's centre. Steerable parabolic dishes exist up to 100 m in diameter, namely the Green Bank [26] and Effelsberg [27] telescopes, but the non-steerable Arecibo dish in Puerto Rico has a 305-m diameter [28] and the planned FAST (Five-hundred-meter Aperture Spherical radio Telescope) telescope will have a 500-m non-steerable dish [29].

The antenna receiver chain processes the detected radio signals. A receiver typically comprises a low-noise amplifier (LNA), band-pass filter and, for single dishes, a square-law detector. Alternatively, signals can be digitised and then squared. The LNA increases the amplitude of the radio signal; without one the noise of components on the receiver chain overwhelm the signal. Unfortunately, the LNA also introduces noise to the signal, however cooling the amplifier can ameliorate the noise level. The band-pass filter defines the frequency limits of the receiver, whilst the square law detector converts the voltages in the receiver chain to a power. The power is typically integrated for a period of time to determine an average power level.

For antennas operating at $\gtrsim 1\,\mathrm{GHz}$ a local oscillator is used to heterodyne the signal to a lower frequency. This is important when digital systems are used later in the signal processing chain, as they may not be able to operate at a high enough speed to accurately sample the original frequency of the signal. Heterodyning can also allow parts of a single receiver chain to be used for multiple receiver bands, as each receiver's frequency band is converted to a common, fixed, frequency at which the rest of the receiver chain operates.

1.2.2 Cylindrical paraboloid

A cylindrical paraboloid reflects radio signals in one dimension. Signals are detected at receivers along the cylindrical paraboloid and may be combined by summations or correlation. Beams can be steered by rotating the reflector(s) along one axis and by delaying the times at which signals detected at different positions along the telescope are combined, a process known as beamforming (see below). Cylindrical paraboloid telescopes are often constructed in two orthogonal parts, which allows narrow 2-D beams to be formed, increasing the resolution of observations.

The Mills cross [30] and 4C array [31] used the cylindrical paraboloid design to conduct important surveys in the Southern and Northern hemispheres, respectively [32, 33, 34, 35, 36].

1.2.3 Aperture arrays

An aperture array, also known as a phased array, is a collection of individual antennas. Signals are detected at each antenna and combined to produce beams, which spatially filter signals so as to be sensitive in a given direction. Signals may be combined by summing in a process known as beamforming, or by cross-correlating the signals. Many aperture arrays beamform their signals, as the signal processing cost is lower than correlating straight away. An aperture array may be constructed as multiple localised groups of antennas; these groups can be called 'stations'. The beam(s) from each station can be correlated with those of other groups.

As electromagnetic waves travel at a finite speed, a radio wave will not be detected at the same time simultaneously across an array of antennas (unless the source is at the zenith). The path delay between antennas is calculated from geometry and is illustrated in one dimension in Fig. 1.1.

For a planar array of antennas, spherical polar coordinate considerations lead to a path delay, ΔL, given by

$$\Delta L = d_{\mathrm{x}} \sin(\theta) \cos(\phi) + d_{\mathrm{y}} \sin(\theta) \sin(\phi) \ , \tag{1.1}$$

where d_{x} and d_{y} are an antenna's location in the x and y directions, relative to a reference antenna, and θ and ϕ are the altitude and azimuth observing direction [37]. Eq. (1.1) assumes a planar array. The antenna distribution for a uniform gridded array and coordinate reference frame are illustrated in Fig. 1.2.

Beamforming is necessary to counteract this delay and synchronise the signals detected by each antenna. This process electronically steers the beam of the aperture array depending on the path delay used, and a major appeal of aperture arrays is that beamforming can be done in multiple directions simultaneously, allowing concurrent observations of different parts of the sky. Electronically steering the beam also allows the telescope to be repointed in very short time spans, as it not necessary to physically move any components.

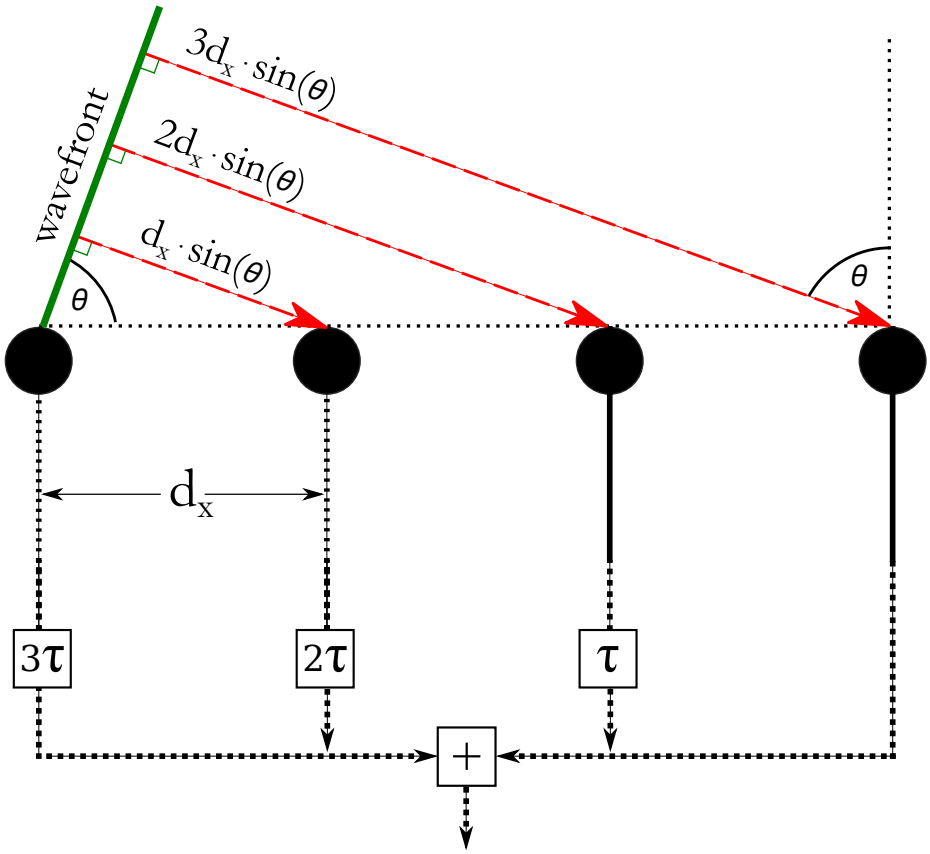

Figure 1.1: Path delays to successive antennas in a linear array. The path delays are compensated by time delays (or equivalent) at each antenna of $\tau = d_x \sin(\theta)/c$, where d_x is the distance from the reference antenna to each antenna, θ is the azimuth angle of the observation and c is the speed of light.

To ensure the combined signals from each antenna correspond to the same point on the wave it is thus necessary to delay the signals received by each antenna by a length dependent on Eq. (1.1). By delaying the signals with reference to the desired observing direction, the observer effectively forms a manoeuvrable beam on the sky, electronically pointing the aperture array. If these delays are not accurately calculated the signals may deconstructively interfere with one another, lowering the signal-to-noise ratio. Indeed, in most cases radio-frequency interference (RFI) will not arrive from the same direction as the beam direction, the effect of RFI is mitigated [38].

Beamformers can be broadly divided into two categories: time-delay and phase-shift. A time-delay beamformer delays signals in time. In analogue beamformers this is achieved by physically varying lengths of transmission lines for the signals to pass

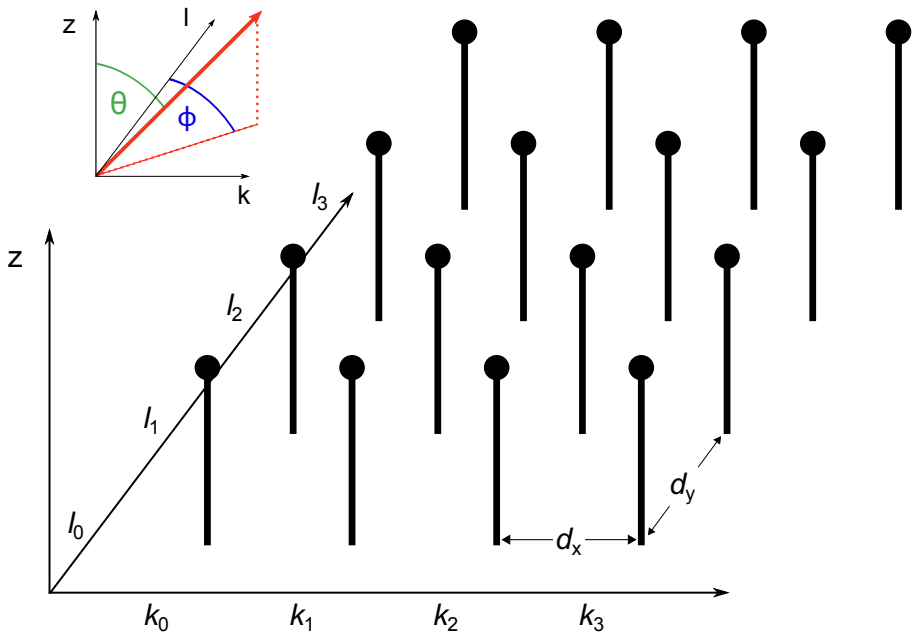

Figure 1.2: An example layout of a 4×4 gridded antenna array. Antenna (k_0, l_0) can be taken as the reference antenna; delays to the other antennas are calculated from its location. The delay's value depends upon the distances between antennas, d_x and d_y and the observing direction (θ, ϕ).

through, but in modern systems digital delays are more likely to be utilised. Although conceptually straightforward, time-delay beamformers have their drawbacks. A time-delay system is relatively expensive to construct and typically utilises switchable signal path lengths with circuit boards or cables. The switchable path lengths must be replicated for each beam required. Digitally-sampling signals allows delays to be implemented more easily, but the delays are discrete, limited to multiples of the sampling interval. Interpolation must be used to delay a signal by a finer resolution, requiring extra signal processing and degrading signal quality.

Phase-shift beamformers channelise the broadband signal detected by antennas into narrow frequency channels, such as to effectively give quasi-monochromatic sine waves, to which phase-shifts can be applied to synchronise the signals. Separating signals into narrow frequency channels facilitates scientific analysis of observations and also allows narrowband RFI spikes to be easily excised [39].

1.2.3.1 Dense and Sparse Arrays

The Nyquist-Shannon sampling theory states that a band-limited signal must be sampled at a rate at least twice the highest frequency in the signal to accurately describe the signal [40]. Alternatively, this can be phrased as sampling a signal at least every half wavelength. This is a fundamental principle in signal processing and determines the sampling rate at which the analogue-to-digital converters must operate.

It is also relevant to the inter-antenna spacing in an aperture array. To fully sample signals in the spatial domain, antennas must be spaced within half a wavelength of one another. If the aperture array covers a wide frequency ratio, such as the 50-$-$ 350 MHz planned for the SKA's aperture array, this may not be possible. The size of the antenna detectors necessary to receive the lowest frequency signals may physically prevent antennas being positioned within half a wavelength at their highest frequency.

An aperture array with antennas within half a wavelength separation of one another is described as 'dense', otherwise it is 'sparse'. For a given number of antennas, sparse arrays have stronger sidelobes, known as grating lobes, in their primary beam patterns than dense arrays; however, as a consequence of the larger station diameters, sparse arrays will also have narrower main lobes, giving them smaller fields of view. Sparse arrays perform better in terms of sensitivity and survey speed, whilst dense arrays perform better in terms of temperature sensitivity (due to their larger beam solid angles) [41, 42]. The frequency at which an aperture array changes from the dense to the sparse regime is called the critical, or transition, frequency.

1.3 Interferometers

Radio interferometry is a radio-astronomical technique whereby signals detected at multiple antennas are combined by cross-correlation to synthesise an image. Interferometry allows imaging of much greater resolution than would be practical with a

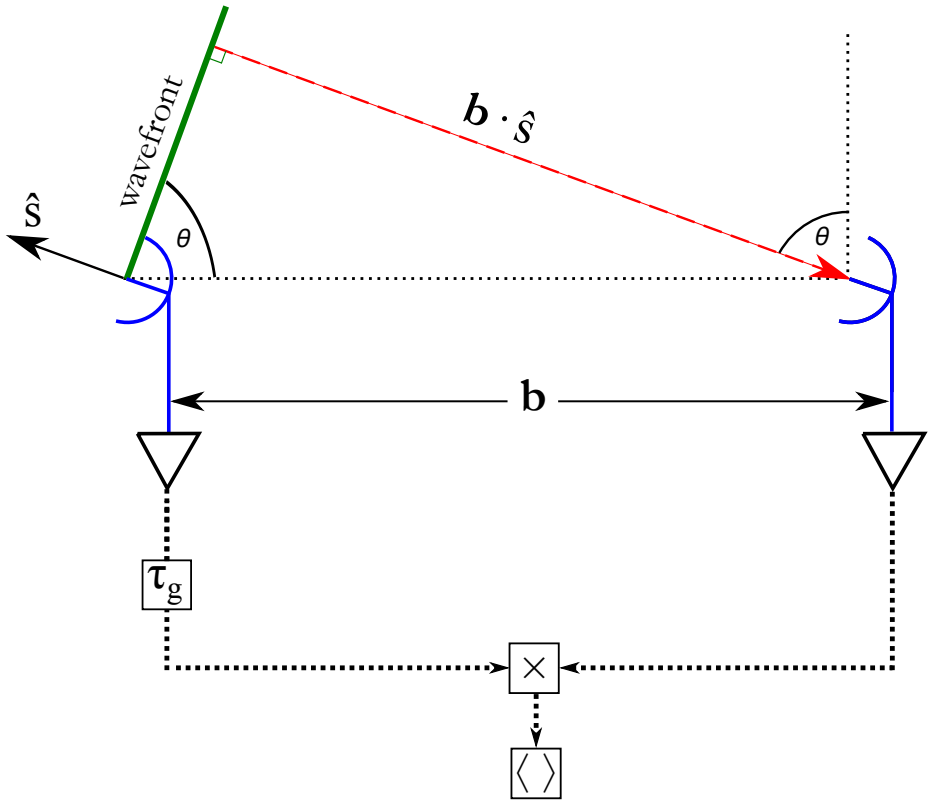

Figure 1.3: A two-element interferometer, separated by a baseline **b**, observing in direction $\hat{s}$. Signals received at the antennas are amplified, filtered and a time delay of τ_g is applied to one antenna. The signals from the antennas are then correlated and integrated.

single dish or single aperture array station, as the resolution is determined by the distance between antennas, rather than the size of any individual dish or station. Furthermore, interferometry can help overcome RFI when it does not correlate in time between antennas, just as with beamforming. Likewise, noise specific to each antenna's individual receiver chain does not correlate, meaning that receiver stability is not as crucial as with single dishes or stations.

Barring exceptions such as the sea-cliff interferometer, radio-astronomy interferometers consist of two or more receiving antennas that combine their signals to reduce noise and increase sensitivity and angular resolution.

When spatially separated antennas observe in the same direction in the sky, the radio signals are received at different points in time due to the different path length

the signal must travel to each of the antennas and the finite speed of light. The corresponding time difference is known as the geometrical delay.

1.3.1 Determining the visibilities

As shown in Fig. 1.3, the time length of the geometrical delay, $\tau_{\rm s}$, is given by

$$\tau_{\rm s} = \frac{\mathbf{b} \cdot \hat{s}}{c}\,, \tag{1.2}$$

where $\mathbf{b}$ is the baseline length between the two antennas, $\hat{s}$ is a unit vector in the direction of observation and c is the speed of light in the medium.

At the antennas the signals are converted to voltages, which have amplitudes and phases in the quasi-monochromatic approximation. For the voltages to be combined in a coherent manner, a process known as correlation, the geometrical delay must be compensated for before the signals are correlated. The correlation operation is defined as

$$r_{\rm cor}(\tau_{\rm s}) = \lim_{T\to\infty} \frac{1}{2T} \int_{-T}^{T} v_1(t) v_2^*(t-\tau_{\rm s})\,{\rm d}t\,, \tag{1.3}$$

where $r_{\rm cor}$ is the output of the correlator, v_1 and v_2 are the two input voltages [38].

The path length of the geometrical delay in terms of the wavelength, λ, is

$$\tau_\lambda = \frac{\mathbf{b} \cdot \hat{s}}{\lambda}\,, \tag{1.4}$$

hence the phase difference, $\Delta\phi$, between the signal at the two antennas is

$$\Delta\phi = \exp\left(-2\pi i \frac{\mathbf{b} \cdot \hat{s}}{\lambda}\right)\,, \tag{1.5}$$

where the factor of -1 is a sign convention indicating which antenna is the reference antenna and λ is the wavelength.

The response of the interferometer for two spatially separated antennas observing at frequency ν is defined by the visibility function of the antenna pair V,

$$V(\nu) = I_\nu \exp\left(-2\pi i \frac{\mathbf{b} \cdot \hat{s}}{\lambda}\right)\,, \tag{1.6}$$

where I_ν is the sky intensity of a single point source at frequency ν and frequency is related to wavelength by $c = \lambda\nu$.

The baseline between the two antennas can be written in terms of orthogonal units (u, v, w), where u is taken as the East-West direction and v the North-South direction in the plane perpendicular to the source direction, and w is the source direction. Likewise $\hat{s}$ can be written in terms of direction cosines on the sky with $\left(l, m, \sqrt{1-l^2-m^2}\right)$, where $l = \sin(\theta_{\rm EW})$ and $m = \sin(\theta_{\rm NS})$; $\theta_{\rm EW}$ and $\theta_{\rm NS}$ are the angles away from the zenith in the East-West and North-South directions, respectively. The third term, $\sqrt{1-l^2-m^2}$, can be used as $\hat{s}$ is a unit vector.

Using these terms Eq. (1.4) becomes

$$\frac{\mathbf{b}\cdot\hat{s}}{\lambda} = ul + vm + w\sqrt{1-l^2-m^2}\,, \tag{1.7}$$

hence extending Eq. (1.6) to cover the whole sky yields

$$\begin{aligned} V_\nu(u,v,w) &= \int\int A(l,m) I_\nu(l,m)\, \mathrm{e}^{-2\pi i\left(ul+vm+w\sqrt{1-l^2-m^2}\right)}\, \mathrm{d}\Omega\,, \\ &= \int\int \frac{A(l,m) I_\nu(l,m)}{\sqrt{1-l^2-m^2}}\, \mathrm{e}^{-2\pi i\left(ul+vm+w\sqrt{1-l^2-m^2}\right)}\, \mathrm{d}l\, \mathrm{d}m\,, \end{aligned} \tag{1.8}$$

as

$$\mathrm{d}\Omega = \frac{\mathrm{d}l\, \mathrm{d}m}{\sqrt{1-l^2-m^2}}\,. \tag{1.9}$$

$I_\nu(l,m)$ is the source flux density of the sky and $A(l,m)$ is the antenna response pattern of each antenna. $A(l,m)I_\nu(l,m)$ gives the response of a single antenna to a sky intensity distribution.

If w is small, Eq. (1.8) simplifies to

$$V_\nu(u,v) = \int\int A(l,m) I_\nu(l,m)\, \mathrm{e}^{-2\pi i(ul+vm)}\, \mathrm{d}l\, \mathrm{d}m\,. \tag{1.10}$$

Equation (1.10) is a Fourier transform between the (l,m) coordinates and (u,v) coordinates, hence it can be inverted to give

$$A(l,m) I_\nu(l,m) = \int\int V_\nu(u,v)\, \mathrm{e}^{-2\pi i(ul+vm)}\, \mathrm{d}u\, \mathrm{d}v\,. \tag{1.11}$$

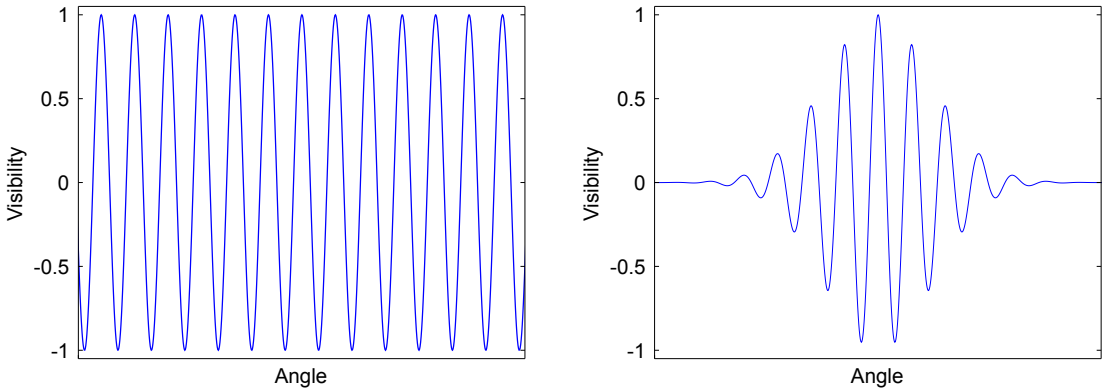

(a) Fringe pattern with no tapering from beam pattern

(b) Fringe pattern with Gaussian tapering from fringe pattern

Figure 1.4: Fringe pattern of a two-element interferometer with and without tapering from a Gaussian beam pattern.

If w is not small then the issue of non-coplanar baselines must be resolved. There are several techniques, which are discussed by Cornwell et al. [43].

If the antenna response of each antenna is considered to be a delta function, then the correlation of the two antennas is a cosine function, as shown in Fig. 1.4a. This is known as a fringe pattern. A closer approximation to the response of an antenna is a Gaussian; this puts a Gaussian envelope on the cosine response, shown in Fig. 1.4b.

The period of the fringe pattern is dependent upon the observation frequency and the baseline length between the two antennas.

The field of view, θ_{FoV}, is dependent upon the width of the main lobe of the antenna beam pattern, and is approximately given by

$$\theta_{\mathrm{FoV}} \approx \frac{\lambda}{d}\,, \tag{1.12}$$

where d is the size of the antenna.

1.3.2 The synthesised beam

Further baselines can also improve the point-source response of the telescope. Summing the fringe patterns from multiple baselines of different lengths yields a synthesised beam which tends towards a narrow beam with weak sidelobes, as illustrated in

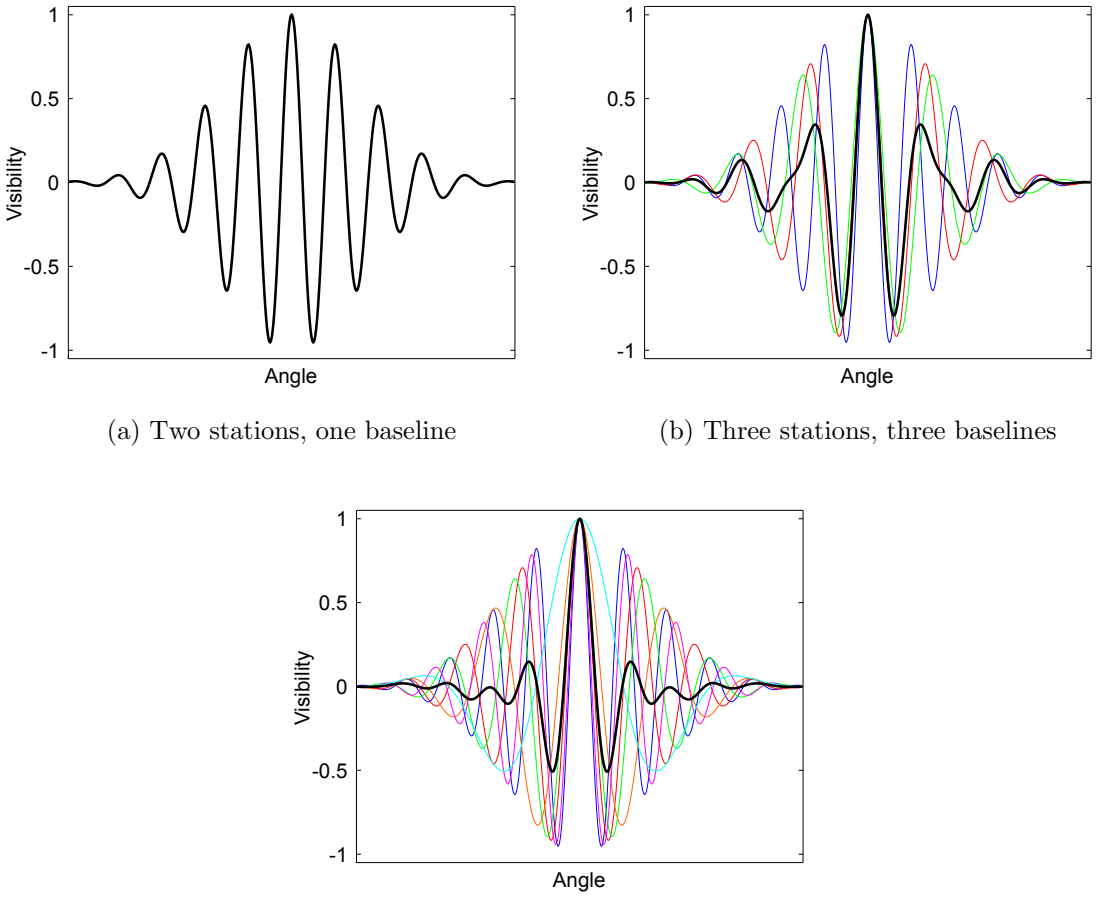

(a) Two stations, one baseline

(b) Three stations, three baselines

(c) Four stations, six baselines

Figure 1.5: The synthesised beam of an interferometer with baselines in one dimension. The narrow coloured lines show the response of individual baselines of different lengths. The thick black lines show the mean response of the individual baseline responses; this is the synthesised beam. Increasing the number of baselines converges the synthesised beam towards a narrow central lobe with smaller sidelobes.

Fig. 1.5. The actual response of the synthesised beam is a convolution of the antenna response pattern with the uv-plane sampling function.

1.3.3 The uv-plane

Figures 1.4 and 1.5 only consider baselines in one dimension. If the baselines are spread out over a 2-D plane (i.e. the u and v directions) a 2-D synthesised beams can be achieved.

As described in Eq. (1.10), visibilities are recorded as a function of (u, v). The

length and direction of each baseline is marked on a set of coordinates known as the uv-plane, the Fourier transform of which gives the point-source response of the interferometer. An example uv-plane is shown in Figs. 1.6a and 1.6b for a three-element interferometer. For each pair of antennas, two points are marked on a uv-plane: once for the baseline from antenna i to antenna j and once for the baseline in the reverse direction. The number of baselines, $N_{\rm b}$, is given by

$$N_{\rm b} = N_{\rm stat}(N_{\rm stat} - 1)\,, \tag{1.13}$$

where $N_{\rm stat}$ is the number of dishes or stations. However the number of unique baselines, $N_{\rm bu}$, is

$$N_{\rm bu} = \frac{N_{\rm stat}(N_{\rm stat} - 1)}{2}\,, \tag{1.14}$$

as the visibility on one baseline is the complex-conjugate of that of its reverse baseline (as the sky is real and the Fourier Transform of a real function is Hermitian conjugate).

The synthesised beam of an interferometer is given by the Fourier transform of the uv-plane sampling function. The greater the fraction of the uv-plane that is filled in, the greater the amount of information that can be determined about the sky intensity distribution.

One method of filling in the uv-plane to a greater extent is to have more dishes/stations with unique spacings between each antenna pair (so as not to duplicate existing baselines). Another approach is to change the spacings between existing antennas. Signals can be correlated for a period of time and combined with the correlated signals from the same antennas located at different physical positions observing at a later time. Some telescopes, such as the VLA, have their antennas mounted on train tracks to allow them to be moved.

A similar benefit can be yielded through utilising the rotation of the Earth. As the Earth rotates on its axis the directions of baselines are rotated relative to the celestial sphere, filling in the uv-plane. An example is shown in Fig. 1.6 with a three-element interferometer. After 12 hours, the rotation allows an entire circle to be filled in on

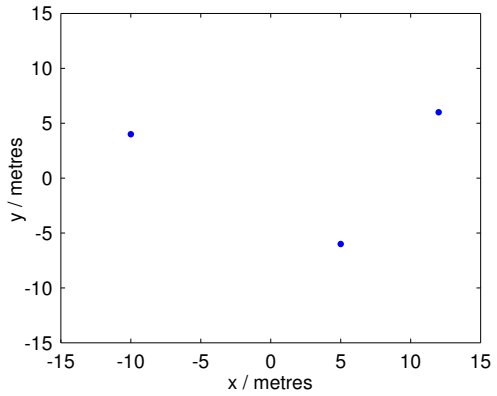

(a) Telescope layout at beginning of observation

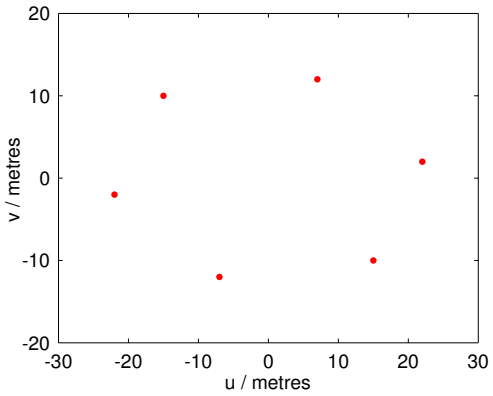

(b) uv-plane at beginning of observation

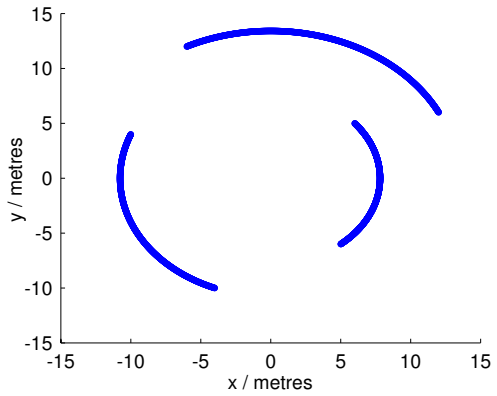

(c) Integrated telescope layout over 6 hour observation (90° rotation)

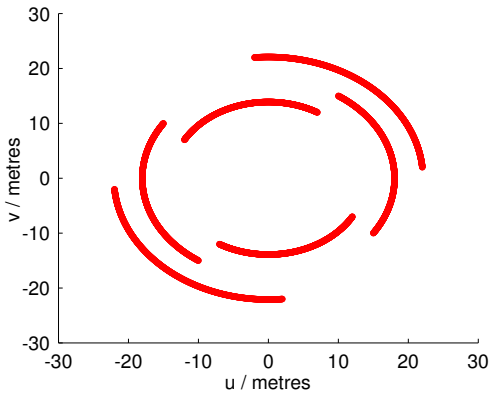

(d) Integrated uv-plane over 6 hour observation (90° rotation)

Figure 1.6: The telescope layout and uv-plane of a three-element interferometer positioned at the North Pole, as seen from the zenith. The Earth rotates over time, changing the projection of the station positions and baselines relative to the celestial sphere. An interferometer positioned away from one of the poles would have its uv-plane tracks foreshortened in the v direction.

the uv-plane. Longer integrations do not fill in the uv-plane further, as baselines are repeated (although the signal-to-noise ratio of observations will be improved).

The layouts of interferometers are generally designed to maximise the benefits of Earth rotation synthesis mapping. The VLA, for instance, has three linear arms spaced at 120° intervals, allowing complete circles on the uv-plane to be filled in 4 hours.

1.4 The Radio Interferometer Measurement Equation

The behaviour of an interferometer can be modelled using an equation known as the Radio Interferometer Measurement Equation (RIME), which expresses the interferometer response in a compact and useful form. It was originally published by Hamaker et al. [44] using Muller matrix formalism and simplified into Jones matrix formalism four years later [45]. The RIME is discussed extensively in a series of papers by Smirnov [46, 47, 48, 49].

The complex, polarised visibility, $\langle V \rangle$, for a baseline can be calculated as

$$\langle V_{p,q} \rangle = \sum_s \mathbf{K}_{p,s}\mathbf{E}_{p,s}\mathbf{G}_{p,s}\mathbf{R}_{p,s}\langle \mathbf{B}_S \rangle \mathbf{R}^H_{q,s}\mathbf{G}^H_{p,s}\mathbf{E}_{q,s}\mathbf{K}^H_{q,s}\,, \tag{1.15}$$

where p and q are the stations forming the baseline, s indexes the sources in the sky and the H superscripts indicates a Hermitian transpose [50].

$\mathbf{B}_s$ is the brightness matrix, defined by

$$\mathbf{B}_s = \begin{bmatrix} I+Q & U+iV \\ U-iV & I-Q \end{bmatrix}\,, \tag{1.16}$$

where I, Q, U and V are the Stokes parameters,

$$I \equiv \langle \mathcal{E}_x^2 + \mathcal{E}_y^2 \rangle \tag{1.17}$$

$$Q \equiv \langle \mathcal{E}_x^2 - \mathcal{E}_y^2 \rangle \tag{1.18}$$

$$U \equiv \langle 2\mathcal{E}_x\mathcal{E}_y \cos(\Delta\phi) \rangle \tag{1.19}$$

$$V \equiv \langle 2\mathcal{E}_x\mathcal{E}_y \sin(\Delta\phi) \rangle\,, \tag{1.20}$$

where $\mathcal{E}_x$ and $\mathcal{E}_y$ are the amplitudes of the components of the electric field and $\Delta\phi$ the phase difference between them.

$\mathbf{R}$ is the parallatic angle rotation matrix

$$\mathbf{R} = \begin{bmatrix} \cos(\psi) & -\sin(\psi) \\ \sin(\psi) & \cos(\psi) \end{bmatrix}\,, \tag{1.21}$$

which rotates signals from a source into the antennas' reference frame. The parallactic angle, ψ, is calculated by

$$\psi = \arctan\left(\frac{\cos(\varphi)\sin(H)}{\sin(\varphi)\cos(\delta) - \cos(\varphi)\sin(\delta)\cos(H)}\right), \tag{1.22}$$

where φ is the telescope's latitude, H the source's local hour angle and δ the source's declination.

G, the element factor (also known as the antenna's field pattern), is the response pattern of an antenna, given by

$$\mathbf{G} = \begin{bmatrix} g_\theta^X & g_\phi^X \\ g_\theta^Y & g_\phi^Y \end{bmatrix}, \tag{1.23}$$

where g_θ^X, g_θ^Y are the responses of the antenna in the θ (altitude) direction and g_ϕ^X, g_ϕ^Y the responses in the ϕ (azimuth) direction. The X and Y superscripts denote components of the antenna with orthogonal polarization responses.

E is the array factor, which is the response pattern of the station and is dependent upon the antenna spacing.

$$\mathbf{E} = e\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \tag{1.24}$$

where

$$e = \sum_a w_a \mathrm{e}^{-i\mathbf{k}\cdot\mathbf{r_a}}, \tag{1.25}$$

and w the beamforming weight, $\mathbf{k}$ the wave vector and $\mathbf{r_a}$ the position of antenna a, given by

$$\mathbf{r}_i = (x_i, y_i, z_i). \tag{1.26}$$

K is the interferometer phase representing the phase delay from the geometrical path inference between stations,

$$\mathbf{K} = k_\mathrm{d}\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \tag{1.27}$$

where

$$k_\mathrm{d} = \mathrm{e}^{-2\pi i\left(ul+vm+w\sqrt{1-l^2-m^2}\right)}. \tag{1.28}$$

The OSKAR-2 (Oxford SKA Radio telescope simulator 2) interferometer simulator[1] uses the RIME to model the visibility data a given telescope would produce for a given observation. These visibilities can be turned into an image using a radio astronomy data post-processing software such as CASA (Common Astronomy Software Applications)[2] or AIPS (Astronomical Image Processing System)[3].

1.5 Gain and phase calibration

The visibilities produced by an interferometer differ from the true visibilities a sky produces for a given spatial sampling due to errors in the signal processing. These errors can arise from a range of inaccuracies, including antenna locations and alignment, quantisation errors, instabilities in amplifiers and other components, atmospheric distortions, and phase-shift differences over frequency channels. The combined effect is an uncertainty on the amplitude gain and phase of the antenna signals, creating inaccuracies in the visibilities, which cause additional noise in images. These are generally uncorrelated between signal processing paths, leading to visibilities having randomised errors [51].

To an extent, telescopes can calibrate for phase and gain errors. Strong sky sources, for instance, can be utilised for calibration. The difference in visibility phases relative to the beam centre, for a given observation, determines the phase offsets. Sources of well-known flux are used to calibrate the gains of signals. During long observations of weak sources, telescopes will periodically observe nearby strong sources to update the calibration offsets, as these can vary with time.

Interferometers can also be calibrated by exploring 'closure relationships', where the baselines of multiple dishes/stations can be combined to remove errors in the phase and gain (see, for example, [38, 52]). For a three-element interferometer the

[1] http://www.oerc.ox.ac.uk/~ska/oskar2/
[2] http://casa.nrao.edu
[3] http://www.aips.nrao.edu/index.shtml

measured phase on each element is given by

$$\begin{aligned}\phi_{12} &= \psi_{12} + \alpha_1 - \alpha_2\,,\\ \phi_{23} &= \psi_{23} + \alpha_2 - \alpha_3\,,\\ \phi_{31} &= \psi_{31} + \alpha_3 - \alpha_1\,,\end{aligned} \tag{1.29}$$

where ϕ is the measured phase, ψ the true phase, α the phase error and the subscripts indicate the individual elements, $1-3$, and their baselines. The sum of Eq. (1.29) cancels the phase errors out

$$\phi_{123} \equiv \phi_{12} + \phi_{23} + \phi_{13} = \psi_{12} + \psi_{23} + \psi_{13}\,. \tag{1.30}$$

By combining ϕ_{123}, the closure phase, from different antenna groups, the phase error on each element can be determined.

Similarly, four elements can be combined to calculate the 'closure amplitude'. The correlator output for a baseline, r_{12}, can be described by

$$r_{12} = G_{12}V_{12}\,, \tag{1.31}$$

where V_{12} is the visibility and G_{12} a gain factor which can be considered as the product of the individual element gains, $G_{12} = g_1 g_2^*$. The individual element gains can be cancelled out by expressing the correlator output of four baselines as

$$\frac{|r_{12}||r_{34}|}{|r_{13}||r_{24}|} = \frac{|V_{12}||V_{34}|}{|V_{13}||V_{24}|}\,. \tag{1.32}$$

The accuracy of the closure phase and closure amplitude techniques increases with the number of dishes/stations [52], however some residual errors will always remain.

1.6 The Square Kilometre Array

The SKA is a next-generation radio interferometer currently being designed with the aim of creating the world's largest radio telescope. The $\sim$€650 million[4] international project is planned to be constructed in the Western Australian desert and South

[4]Budgeted phase 1 construction cost

African Karoo desert. The first stage, 'phase 1', is specified to comprise one hundred and thirty-three 15-m dishes, sixty-four 13.5-m dishes and 131 072 log-periodic dipole antennas. Phase 1 is specified to be 10% of the size of the phase 2 telescope.

The dishes will be located in South Africa with single pixel receivers and observe from $0.35 - 13.8\,\mathrm{GHz}$. This telescope will be called SKA1-Mid.

The log-periodic antennas, to be placed in Australia, will be grouped into five hundred and twelve 35-m stations of 256 antennas each, covering a band-pass of $50 - 350\,\mathrm{MHz}$. This aperture array will be known as SKA1-Low. SKA1-Low will have a collecting area of $\sim 0.5\,\mathrm{km}^2$, approximately three times more collecting area than the UTR-2 (Ukrainian T-shape Radio Telescope, second modification), the current largest radio telescope by collecting area [53].

Log-periodic antennas allow a large frequency range to be covered with one antenna design. They can also be manufactured with rigid pieces of wire, which reduces the amount of metal required and hence the cost.

First conceived in the early 1990s [20, 54, 55], the SKA is now an international project with twenty countries working towards its realisation. It is headquartered at the Jodrell Bank Observatory, south of Manchester, UK, adjacent to the Lovell telescope. Construction is scheduled to begin in 2018.

1.6.1 Site selection

Beyond the political and economical factors that are inevitably entangled with any major scientific project such as the SKA, there are numerous astronomical reasons why the Western Australian desert and South African Karoo desert have been selected to host the telescope.

The first is that these sites are in the Southern Hemisphere, which is commonly regarded as the more astronomically 'interesting' hemisphere: the Galactic Centre of the Milky Way is visible from the Southern Hemisphere, for instance.

However, the relatively low RFI level at both sites is the most important factor. Man-made radio signals for communications, from electronic components and electrical infrastructure, are a significant hindrance to radio astronomical observations, not least because these signals are generally orders of magnitude more powerful than astronomical signals. The signals detected by radio telescopes are split into narrow frequency channels; if a channel measures an extremely strong signal that data is generally regarded as contaminated by RFI and is discarded.

Both SKA sites are in sparsely populated areas away from most sources of RFI, although some types of RFI, such as from satellites, are inescapable anywhere on the planet. Both the Australian and South African government have passed laws to help maintain their respective site's radio quiet environments over the lifetime of the SKA [56, 57].

Another important factor when assessing the suitability of a site is the atmosphere above the telescope. Both SKA sites had extensive investigations conducted to measure their atmospheric behaviour [58].

One atmospheric effect is scintillation, in which wavefronts are corrupted by ionised particles. Whilst commonly understood as an issue for optical observations, scintillation can have significant implications for low-frequency radio observations.

Free electrons in the ionosphere can absorb low-frequency astronomical signals. This absorption is dependent on the electron density; lower densities cause more absorption and it is typically only noticeable below 25 MHz [59]. However at 100s of MHz, refraction by the free electrons in the ionosphere is a significant potential problem. This causes variations in the apparent position and brightness of low-frequency sources. At the other end of the SKA's frequency band, absorption by water vapour and other molecules in the troposphere can also hinder observations. The desert sites should mitigate the presence of water vapour. Calibrating out ionospheric effects will be a significant problem for SKA1-Low, however ionospheric effects will not be directly addressed in this thesis.

1.6.2 Rebaselining

The specifications for the SKA were originally given in March 2013 in the *Baseline design* [60], which defined SKA1-Low as comprising 262 144 (2^{18}) antennas, grouped into 1024 stations of 256 antennas each. The same baseline design also specified SKA1-Mid as comprising 256 dish antennas (sixty-four 13.5-m dish antennas of MeerKAT plus 190 new 15-m SKA dish antennas). A third array, SKA-survey, was also to be constructed with made of thirty-six 12-m dish antennas from ASKAP and sixty 15-m SKA dish antennas, all with phased-array feeds.

In April 2015, following a review of the science requirements and estimated costs of this design, the project produced a new specification, *Requirements specification, revision 6* [61], which reduced the number of SKA1-Low antennas by a factor of two, reduced the number of new 15-m dish antennas in SKA1-Mid to 133, and eliminated SKA-survey entirely.

As most of the work in this thesis had been completed prior to the release of *Requirements specification, revision 6*, we have kept to the original baseline for all the results presented here. We do not expect there to be any significant qualitative change in the conclusions reached about the performance of SKA1-Low besides a factor-of-two decrease in instantaneous sensitivity.

1.7 Science with the SKA

The SKA is intended to be a general-purpose observatory, however it has been motivated by 'key science projects', which the design has been tailored towards. They were originally defined by the SKA International Science Advisory Committee and outlined in *Science with the Square Kilometre Array* [62]. These science cases have recently been updated in the decade since the original Science Case was published, most recently at the *Advancing Astrophysics with the Square Kilometre Array* conference [63]. Seven 'Science Working Groups' are now defined for the SKA's key science projects [64]. The key science projects are briefly described below.

1.7.1 Cosmic reionisation and the first galaxies

At high redshifts ($z \sim 10 - 20$) neutral hydrogen was reionised by early stars and galaxies. Mapping this era, the Epoch of Reionisation (EoR) will provide insight into the early universe and evolution of large-scale structure. The frequency coverage of SKA1-Low has been set to allow observations from the beginning of the period, the Cosmic Dawn, until the end of Reionisation, allowing a detailed study of the formation of the first stars and galaxies [65].

1.7.2 Cosmology

The SKA will allow studies of the equation-of-station of dark energy and the dark matter power spectrum. A large-scale continuum survey could provide insight into weak lensing and constrain its properties. The SKA will be in a strong position to compliment the cosmological observations of other telescopes such as Euclid [66].

Cosmology with the SKA will primarily focus on a range of surveys with SKA1-Mid, such as a HI galaxy redshift survey detecting $\sim 10^7$ galaxies out to $z \sim 0.7$, radio continuum surveys out to $z \sim 6$ of $\sim 10^8$ galaxies and weak lensing surveys. These surveys will examine the expansion history and geometry of the Universe and the growth of large-scale structure [67].

1.7.3 Galaxy Evolution

As hydrogen is the most abundant atom in galaxies, hydrogen's 21-cm radiation can be used to measure galactic phenomena. HI tracks large-scale galactic dynamics and yields insights into the interstellar medium (ISM), such as the relationship between star formation and the state of the ISM [68]. The SKA will be able to make all-hemisphere surveys of HI emission to a redshift, z, of ~ 1.5.

Continuum surveys will provide data on the star formation history of the Universe, the role of black holes in galaxy evolution and increase the number of radio-detected gravitational lenses to $\sim 10^5$ [69].

1.7.4 Cosmic Magnetism

The SKA will exploit Faraday rotation, polarized synchrotron emission and the Zeeman effect to study magnetic fields, which remain a relatively unstudied aspect of astrophysics. Measurements of the intensity of synchrotron emission can be related to the strength of the magnetic field it was produced in and the orientation of the magnetic field can be ascertained from the polarization of the synchrotron emission.

Faraday rotation is the rotation of the plane of polarization of an electromagnetic wave along its line-of-sight when passing through a plasma and magnetic field. The rotation of the polarization angle is dependent upon the magnetic field strength and the electron density in the plasma. The magnetic field strength can be inferred from measuring the frequency dependence of the polarization angle of radio waves.

1.7.5 Pulsars and Gravity

The unparalleled sensitivity and survey speeds of the SKA increase the likelihood of detecting a neutron star - black hole binary system, allowing predictions from General Relativity to be tested in strong-field cases. The timing of pulses from a neutron star orbiting a black hole would provide accurate measurements on the effects of strong-field gravity.

A large-scale pulsar survey with the SKA's sensitivity is expected to detect a large number of millisecond pulsars (pulsars with rotation periods of order a millisecond), allowing a pulsar-timing array to be established. Direction dependent deviations in the pulses' arrival times could herald the detection of gravitational waves.

1.7.6 The Cradle of Life

This science goal encompasses studying terrestrial planet formation and the evolution of proto-planetary disks as well as detecting spectral lines from amino acids and other complex carbon molecules. The SKA will be able to detect centimetre wavelength emission from large dust grains and pebbles, with the aim of understanding how the

pebbles coalesce to form boulder-sized objects, an important but unsolved question in planet formation [70].

The SKA could also detect leaked radio frequency emissions from extra-terrestrial life around nearby stars. A survey of the nearest thousand stars to detect leakage of airport-radar strength signals could be conducted within 1000 hours of observations [71]. The sensitivity of the SKA also enables the detection of a focussed 1-MW beam, like the Arecibo planetary radar, emitted up to thousands of light years away [72].

1.7.7 Time domain

The survey speed and sensitivity of the SKA will facilitate synoptic surveys, studying the variation of the sky and sources on scales of minutes to days. It should also help detect more Fast Radio Bursts (FRBs) and Tidal Disruption Events (TDEs), which are thought to occur when massive objects pass closely to central galactic black holes. The SKA is predicted to detect $1-1000$ FRBs and TDEs per week, vastly increasing the current number of measurements. This will help astronomers to understand the nature and implications of these events [73].

1.7.8 SKA1-Low key science

The remit of SKA1-Low is focused on the EoR and pulsar searching. The EoR's 21-cm radiation is redshifted to the low-frequency regime of SKA1-Low. The large field of view offered by SKA1-Low make it amenable to large-scale pulsar surveys.

1.8 Current low-frequency telescopes

The scientific interest in exploring the low-frequency radio spectrum, coupled with advances in digital signal processing technologies, has seen a surge in the commissioning of low-frequency telescopes in recent years. Many of these telescopes are now recognised as SKA 'precursors' or 'pathfinders' (the precursors are located on

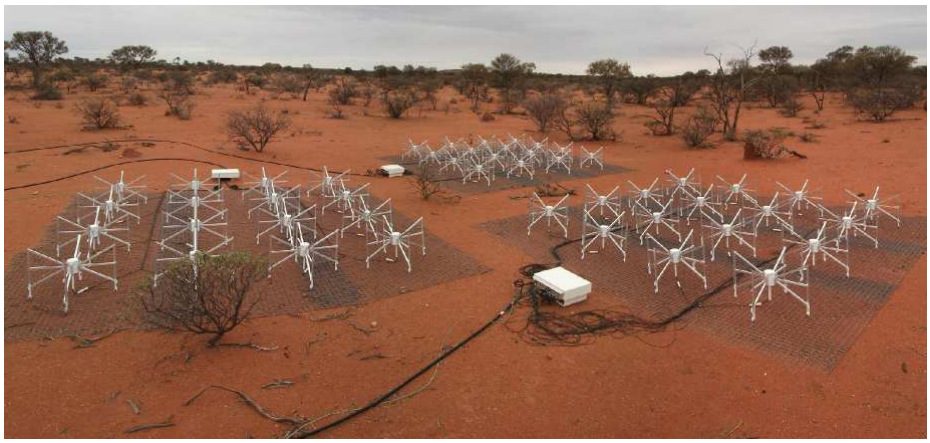

Figure 1.7: Three Murchison Widefield Array (MWA) tiles. Each tile is composed of a 4×4-grid of bow tie antennas on top of 5×5-m ground screens. The ground screens reduce the ground noise detected by the antennas and help maintain constant electrical properties beneath the antennas, which may vary with soil moisture levels. The antenna signals are beamformed in the white boxes adjacent to the tiles. (Photograph by Paul Bourke and Jonathan Kaispel. Supported by WASP (UWA), iVEC, ICRAR and CSIRO.)

one of the SKA sites), as the knowledge and experience gained from their construction, operation and observations will be able to directly contribute to the SKA's own low-frequency efforts. Some of these telescopes are outlined below.

1.8.1 Murchison Widefield Array (MWA)

The MWA comprises 2048 dual-polarization dipole antennas covering $80 - 300\,\mathrm{MHz}$. It is located at Boolardy station in Western Australia, the future site of SKA1-Low, and covers much of the same frequency range. Those two factors make the MWA a valuable project for developing SKA1-Low.

The 2048 antennas are divided into 128 'tiles' of 4×4 antennas, where an MWA tile can be thought of as being equivalent to an SKA1-Low station. The tiles are beamformed with an analogue time-delay beamformer, meaning the accuracy of the beamforming does not have a frequency dependence. The tiles are mostly concentrated in a 1.5-km core but there are baselines of up to 3 km, allowing arcminute resolution [24].

Covering approximately the same frequency range as SKA1-Low, the MWA also

Figure 1.8: PAPER dipole antenna. The copper dipole is surrounded by a ground screen with side panels made of steel and wire mesh, which narrows the antenna's beam size to approximately the size of cold patches in the synchrotron sky [75]. (Image from `eor.berkeley.edu`).

has similar scientific aims. Specifically stated targets include imaging ionised bubbles of hydrogen, and measuring variations in the brightness of the 21-cm line, during the EoR. The MWA has also set out to conduct a Southern Hemisphere sky survey and hopes to detect a large number of faint supernova remnants, as current observations do not match the predicted number. Other science goals include searching for transient and variable events and monitoring space weather, such as solar bursts [74].

1.8.2 The Precision Array for Probing the Epoch of Reionisation (PAPER)

PAPER is an American project with two current arrays; one at Greenbank, West Virginia and the other in the Karoo desert, near the SKA South Africa site. The former has 32 dual-polarization antennas, the latter 64 (and currently expanding to 128). An image of a PAPER antenna is shown in Fig. 1.8.

Unlike SKA1-Low and the MWA, PAPER has no beamformer. Instead all the antennas are correlated with one another. This significantly increases the correlator operation costs, which increases with the square of the number of antennas, but precludes the possibility of errors in the beamformer and the grating lobes in the station beam introducing spectral structure that hinders attempts at foreground removal for

Figure 1.9: A station of the Long Wavelength Array (LWA) array, located in New Mexico. 256 antennas are visible in the core plus 1 outrigger antenna. The antennas have a dual-polarization design with low-noise amplifiers mounted on the top. Signals are routed underground to the correlator in the white cabin next to the array. The site is in close proximity to the VLA, which is visible in the upper right. (Image from `phys.unm.edu`).

EoR observations, as could occur in the case of the MWA's and LOFAR's (see below) gridded tiles.

As its name suggests, PAPER is focussed on measurements of the EoR. Results from PAPER have recently put upper limits on the EoR power spectrum, constraining it to below $(22.4\,\mathrm{mK})^2$ at $z = 8.4$ [76] and $(48\,\mathrm{mK})^2$ at $z = 7.55$ [77].

1.8.3 The Long Wavelength Array (LWA)

The LWA is another American project split across two sites: the VLA site in New Mexico and the Owens Valley Radio Observatory in California. The New Mexico site has 256 dual-polarization antennas in a compact core with four outrigger antennas on up to $\sim$ 300-m baselines from the core and beamforms four beams. An image of the core is shown in Fig. 1.9.

The main science drivers of the LWA are synchrotron radiation observations of supernova remnants, galaxies and clusters and ultra-high energy cosmic rays. It aims to measure cosmic evolution at high redshifts, including the EoR and study the

Figure 1.10: The international Low-Frequency Array (LOFAR) station in the UK. The Low Band Antennas, a 'droopy-dipole' design, can be seen in the foreground. The High Band Antennas are under the black weatherproof covering in the background. (Image from `astron.nl`).

transient phenomena. Observing down to 10 MHz also allows ionospheric studies [78].

The Owen's Valley LWA has 256 dual-polarization antennas. The Large Aperture Experiment to Detect the Dark ages (LEDA) digital signal processor correlates these antennas over an instantaneous bandwidth of $30 - 88$ MHz. The LWA with LEDA is studying the 21-cm line by attempting to measure its sky-averaged absorption against the CMB at $z \sim 20$ [79].

There are plans to expand the LWA New Mexico out to 53 stations around New Mexico, with baselines up to 400 km [80].

The groups behind the MWA, PAPER and LEDA plan to construct a next generation telescope called the Hydrogen Epoch of Reionisation Array (HERA). The collecting area of HERA will in concept be similar to SKA1-Low, to allow it to conduct detailed observations of the EoR [81].

1.8.4 The Low-Frequency Array (LOFAR)

Primarily a Dutch project, with support from international partners, LOFAR is the most ambitious of the current generation aperture arrays in terms of number of antennas and baseline lengths. It currently comprises 18 stations in a 2×3-km core in the Netherlands, 18 further Dutch stations on an approximate logarithmically spaced spiral distribution and 8 international stations in Germany, the UK, France and Sweden,

with a maximum baseline of $\sim 1500\,\mathrm{km}$.

Each station is split into low-band antennas (LBA, $10 - 90\,\mathrm{MHz}$) and high-band antennas (HBA, $120 - 240\,\mathrm{MHz}$). The frequency gap covers the FM radio band; indeed the LBAs are, by default, filtered to $30 - 80\,\mathrm{MHz}$ due to strong RFI at both ends of the LBA band. There are 96 LBA tiles per station and 48 or 96 HBA tiles (Dutch and international stations, respectively). An LBA tile is a dipole antenna whereas an HBA tile is a 4×4 beamformed grid of antennas, leading to a grating lobe in the station beam pattern. Each station typically has 8 beams, but the core stations can all be beamformed together to create up to 219 beams.

As with the other low-frequency telescopes, LOFAR is studying the EoR and transient events. It also aims to conduct surveys for both high redshift radio sources and pulsars and search for exoplanets [25].

1.8.5 Giant Metrewave Radio Telescope (GMRT)

Another telescope hoping to make EoR measurements is the GMRT in India. Unlike the other telescopes, the GMRT is an array of thirty 45-m dishes, not an aperture array. It has frequency bands centred on 50, 153, 233, 325, 610 and 1420 MHz. The 153 MHz band has been used to constrain the EoR power spectra to an upper limit of $(70\,\mathrm{mK})^2$ at $z = 8.6$ [82].

1.9 Thesis layout

The SKA will provided unprecedented sensitivity and survey speed capability to the radio astronomy community, by virtue of its large collecting area. However this high sensitivity is of little use if the observations are limited not by thermal noise but by various kinds of systematic error.

This thesis focuses on SKA1-Low, the SKA's low-frequency aperture array. The aim of this thesis is to analyse the performance of SKA1-Low, in order to determine limiting factors on its performance, and to understand how systematic errors of

various types will impact the results.

Chapter 2 assesses the impact of phase errors and broken antennas on the primary beams of low-frequency aperture arrays. Chapter 3 extends these simulations to the complete interferometer, to cover whole telescopes and the impact of gain and phase errors on the telescope's imaging capability. In Chapter 4 we create and analyse a range of telescope models conforming to the SKA1-Low specifications on the number of antennas and their radial distribution. These are assessed to see how sensitive they are to detecting the Epoch of Reionisation power spectrum in Chapter 5. Chapter 6 analyses the hardware requirements of implementing the beamformer. Chapter 7 summarises the conclusions of the thesis and discusses future work.

Chapter 2

Simulations of aperture array primary beam patterns

2.1 Introduction

SKA1-Low will be an aperture array comprising a large number of approximately identical stations. The stations will create their own primary beams (also known as station beams) which define the field of view of the telescope.

The shape of a primary beam is principally governed by the number and location of antennas within a station and the beam pattern of the individual antennas; the primary beam shape is the Fourier transform of the antenna layout and the individual antenna patterns. Stations with larger numbers of antennas typically have lower sidelobe levels in their primary beam, as shown in Fig. 2.1. A gridded layout will result in regular strong sidelobes known as grating lobes unless the array is fully sampled (i.e. the antenna spacing is $\leq \lambda/2$), in which case the grating lobes are pushed over the horizon. A randomised layout of antennas will more evenly distribute the power in the sidelobes, resulting in irregular sidelobes with the peaks of near-in sidelobes generally significantly weaker. It is proposed that SKA1-Low stations will have a random layout [60] using different layouts for each station (or the same layout but rotated), as this results in even lower sidelobes in the cross-product beam between two stations.

In this chapter the beam patterns of aperture array stations of various sizes will be

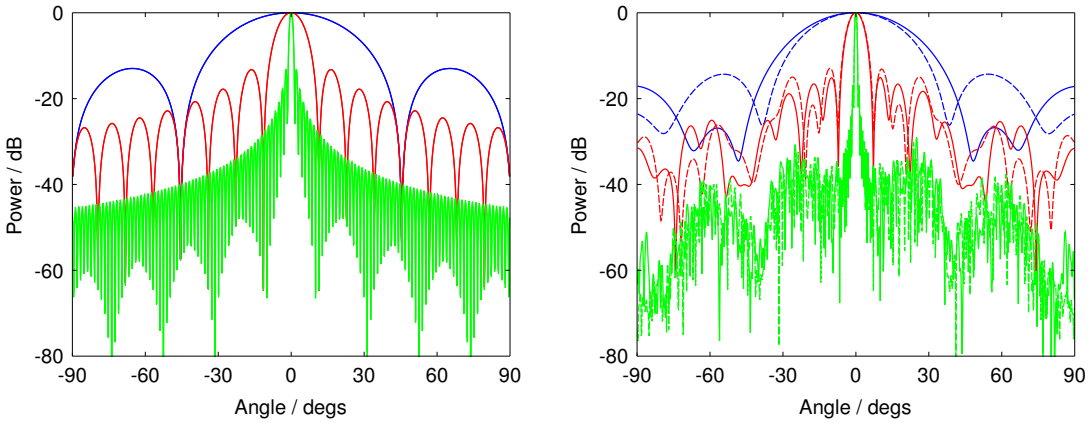

(a) Rectangular gridded layout of antennas. The blue line plots the beam pattern of a 4×4 antennas station, red 16×16 and green 128×128. Orthogonal cuts through the beam pattern that are aligned with the axes of the grid are identical.

(b) Randomised layout of antenna. The blue line plots the beam pattern of a 16-antenna station, red 256 and green 16 384. Solid and dashed lines show orthogonal cuts through the beam pattern.

Figure 2.1: One-dimensional cuts through the centre of primary beams of gridded and randomised array stations with differing numbers of antennas. Antennas are spaced at the critical frequency (on average, in the randomised stations). The gridded layouts give regular high-power grating lobes, whereas the randomised layout distributes the sidelobe power over a wider area. Orthogonal cuts through the gridded layout give identical beam patterns, but differing patterns in the randomised case. The difference between the orthogonal cuts of a randomised station tends to zero with increasing numbers of antennas.

simulated. The effect of station sizes on beam patterns will be discussed. In addition, the impact of phase errors and broken antennas on a range of beam pattern metrics will be investigated.

2.1.1 Primary beams

Understanding the shape and performance of the primary beams will be important to the success of SKA1-Low, as bright sources in sidelobes can interfere with measurements. By uniquely randomising the layout of each station, the detrimental effect of sidelobes will not be as strong as in the case of identical layouts. This is because a bright source located at the peak of a sidelobe of one station may be at a minima in a second station, so the correlation of the two will result in a weaker contribution to

the final image than if identical station layouts were used with the source at the peak of a sidelobe in both.

The size of a station's primary beam, and hence its field of view, $\mathrm{FoV}_{\mathrm{beam}}$, is also dependent on the location of antennas in the station, specifically the diameter of the station, d_{stat}. This is approximately described by Eq. (2.1) [60],

$$\mathrm{FoV}_{\mathrm{beam}} = \frac{\pi}{4}\left(\frac{1.3\lambda}{d_{\mathrm{stat}}}\right)^2\left(\frac{180}{\pi}\right)^2 , \tag{2.1}$$

where λ is the observing wavelength and $\mathrm{FoV}_{\mathrm{beam}}$ is in square degrees and d_{stat} is the effective station diameter (this can be different from the physical diameter depending on the illumination function of the station used, where the illumination function is the Fourier transform of the antenna's voltage reception pattern). Assuming a circular station and an average separation between antennas of d_{ant}, the diameter and number of antennas, N_{a}, in the station are related by

$$d_{\mathrm{stat}} = 2\sqrt{\frac{d_{\mathrm{ant}}^2\, N_{\mathrm{a}}}{\pi}} . \tag{2.2}$$

In this chapter the minimum separation between antennas in assumed to be 1.35 m, half the wavelength at the critical frequency, which is taken to be 111 MHz.[1] The mean value between antennas is set at 1.63 m; this is equivalent to the mean separation of an antenna from its 8 nearest neighbours in a gridded layout with adjacent separations of 1.35 m between antennas.

Studying the Epoch of Reionisation is a key science project of the SKA [84] and investigations have recommended that a field of view of 25 deg^2 at 100 MHz will be required to achieve this [85]. This is approximately the field of view achieved with a 35-m station at 100 MHz.

If larger station sizes than 35 m were created for SKA1-Low, theoretically the 25 deg^2 field of view could be achieved by mosaicing multiple beams. It has been argued that it may not be possible to detect the EoR with a multi-beaming approach as it

[1]The work in this chapter pre-dates more recent updates to the antenna design specifying the average antenna spacing as 1.93 m [83]. The conclusions are not significantly changed.

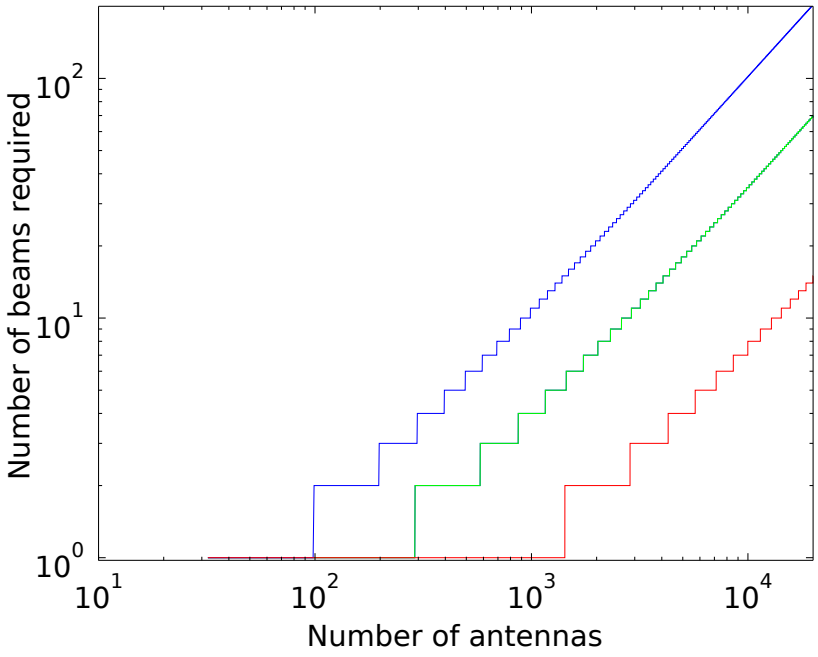

Figure 2.2: Number of beams required to maintain at least a 25 deg^2 field of view as a function of number of antennas at 50 MHz (red line), 111 MHz (green line) and 190 MHz (blue line). Plot assumes a mean antenna separation of 1.63 m (less than the 1.93 m currently specified for SKA1-Low).

is difficult to determine fluctuations and structures on angular scales comparable to that of the station beams [85].

Higher frequency observations have smaller fields of view, hence there is a requirement for more beams. Note that at 50 MHz the single beam area is already much larger than the 25 deg^2 requirement, whilst at 190 MHz multi-beaming is already required to maintain the 25 deg^2 beam. Due to the flat shape of the angular size-redshift relation, the angular size of the interesting features in the EoR signal does not vary much with frequency (see Chapter 5). Each extra beam would require an increased number of digital signal processing operations in the beamformer and correlator, adding expense to the telescope cost. Larger stations would result in fewer stations, however, reducing some digital signal processing costs. The trade-off between these factors is explored in Price et al. [86], who conclude that somewhat larger stations ($\sim$ 512 antennas) are optimal from a signal processing point of view.

2.2 Beam pattern metrics

There is a range of metrics that can be used to analyse a beam pattern, the pertinence of each depending upon the aims of the aperture array (or radio dish). For instance, an aperture array used for radar may aim to achieve a constant power level across a wide field of view. The directivity of the beam may help analyse this, as it measures the ratio of the peak power of a beam to its average power (hence an isotropic beam has a directivity of one). A ground-based dish for satellite communications may desire a narrow beam, to restrict communications to individual satellites and avoid nearby ones operating in the same frequency band, which could be measured by the half-power beam-width (HPBW). Astronomical telescopes, on the other hand, generally target low sidelobes and high sensitivities. The sidelobe levels of the station beams are considered by some authors to be of particular importance as to whether key science goals can be achieved with SKA1-Low [87, 88].

The typical sidelobe level in a station is roughly inversely proportional to the number of antennas in the station [51], however the actual locations and strengths of sidelobes are controlled by the station antenna layout. Sidelobes can be categorised as either 'near-in' or 'far-out'. Near-in sidelobes are caused by the overall illumination function of a station; a (uniformly weighted) station is only sensitive over a finite area, beyond which the sensitivity drops instantly to zero. As a consequence of this, near-in sidelobes appear around the main lobe of the beam. This can be mitigated by down-weighting antennas towards the edge of the station in a process known as apodization; see Chapter 4 for further discussion.

Far-out sidelobes are caused by the direction and spacing of the gaps between antennas within a station [89]. This is why gridded antenna layouts create strong grating lobes: the spacing between antennas is constant. Positioning antennas in a random manner creates a distribution of inter-antenna spacing and directions with the effect of redistributing the power in the grating lobes over a wider area but with a lower mean power.

Errors in the beamformer will distort the beam patterns, changing these metrics and impairing an aperture array's ability to perform as expected. With $\sim 250\,000$ antennas in the SKA1-Low alone, spread out along baselines of up to $\sim 100\,\mathrm{km}$ [60] some degree of error will have to be tolerated. Variations in the precision of mass-produced components, quantisation errors from finite bit depths and broken components, amongst other factors, will all contribute to decreasing the overall performance of the telescope.

Various theoretical analyses of the effect of phase errors on generalised arrays exist, for instance [51, 90]. Schediwy et al. [91] provide a quantitative analysis of the effect of phase errors on linear arrays using the (original) OSKAR simulator [92] and relate it to the 2-PAD (2-Polarizations All Digital) aperture array, a digital beamformer created as a demonstrator for SKA1-Low [93].

Here we will simulate the effect of phase errors, along with station size and fraction of broken antennas, on station beams. There are no definitive beam pattern measures to assess the quality of beam patterns in radio astronomy, however here three metrics are chosen to assess key parameters that strongly influence the outcome of radio astronomical observations. These are sidelobe level, sensitivity and beam-pointing accuracy.

2.2.1 The OSKAR-2 interferometer simulator

The OSKAR-2 interferometry simulator is designed to be a powerful tool in modelling the performance of radio telescopes. OSKAR-2 provides the opportunity to simulate and determine the effect of flaws in the beamformer on the overall performance of the telescope. This can provide constraints on the tolerable errors in the SKA beamformer, such that it can meet its key science goals.

OSKAR-2 calculates visibility data using the Jones matrix-based formalism of the Radio Interferometer Measurement Equation [44, 94]. A station's beam pattern is calculated by $\mathbf{E} \cdot \mathbf{G}$, where $\mathbf{E}$ and $\mathbf{G}$ are the Jones matrix representations of the

array factor and the element factor, respectively.

Errors in the amplitude gain and phase are incorporated into the beamforming weights. The beamforming weight, W, defines how the telescope's beam is pointed, and is given by

$$W(\theta_b, \phi_b, x, y) = W_{\text{geo}}(\theta_b, \phi_b, x, y) \cdot G\mathrm{e}^{i\phi} \ , \tag{2.3}$$

where θ_b and ϕ_b are the azimuth and altitude beam-pointing directions, x and y are the antenna position, G is the gain and ϕ is the phase. W_{geo}, the geometric beamforming weight, is a complex value dependent on the station layout and beam-pointing direction.

2.3 Station size simulations

Models for each of 13 station sizes from $32 - 11\,200$ antennas, or station diameters of $\sim 9 - 180$ m (the upper bound equalling the number of antennas of SKA stations in the 2010 *Preliminary system description* [95]) were created by specifying the antennas' geometric locations. One hundred unique models were generated for each station size with the antennas in randomised layouts, yet kept within a diameter specified by Eq. (2.2) and with at least 1.35 m (half the wavelength at the critical wavelength) separation between antennas.

Beam patterns from each model were simulated using OSKAR-2. Four one-dimensional (1-D) slices through the centre of each beam pattern were taken at 45° angles to one-another. The median sidelobe level is determined from measuring the value of each 1-D slice from beyond the first minima either side of the main lobe. Beam patterns are simulated at 111 MHz in the zenith direction. Pointing away from the zenith would increase the sidelobe levels.

2.3.1 Results of station size simulations

As shown in Fig. 2.3, the sidelobe levels of the station beams are inversely proportional to the number of antennas in a station, following a trend of k/N_a where $k = (6.7 \pm$

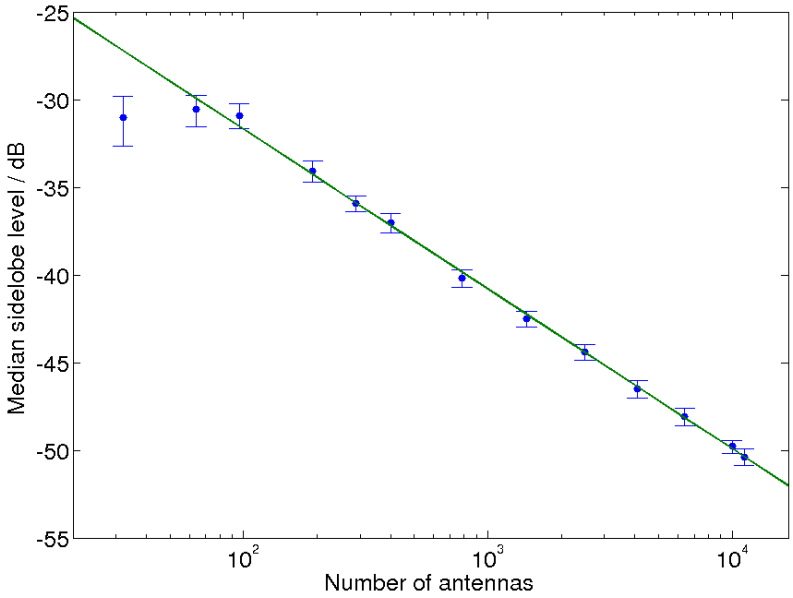

Figure 2.3: Median sidelobe level of beampattern as a function of station size. The sidelobe level is inversely proportional to the station size. The range of values at each station size shows the standard error of the sampling distribution from hundred randomised antenna layouts.

0.3) $\times 10^{-2}$ is a constant from a χ^2 fit. A station of 100 antennas yields a median sidelobe level of ~ -32 dB, 1000 antennas lowers it to ~ -42 dB and 10 000 results in ~ -52 dB. This is consistent with the theoretical prediction that the typical voltage response of the station will go as $1/\sqrt{\mathrm{Na}}$ when the antenna contributions are adding up with quasi-random phases and hence that the power pattern will go as $1/\mathrm{Na}$.

The value for $N_a = 32$ is not included in the fit. This is because at low frequencies and with small station sizes, the field of view of a station covers a large fraction of the sky. The median sidelobe value of the 32-antenna station was artificially lowered by having the peaks of at least one strong grating lobe beyond the horizon and thus not measured. This is a limitation of the definition of sidelobes used here (the beam pattern beyond the first minima)

The median sidelobe level required for SKA1-Low is not specified in the SKA's *Requirements specification* but the beamformer is listed as requiring a dynamic range of better than 40 dB [96]. For the station beam to output data with a signal-to-noise

ratio of better than 40 dB the median sidelobe level should be lower than 40 dB, as otherwise the background noise level will be too great. This suggests station sizes of $\gtrsim 1000$ should be constructed. However this does not account for the potential benefits of correlating station beams later in the signal processing chain, nor the effect of each station having a different random configuration, which introduces an additional drop in power on each baseline.

2.4 Phase error simulations

Investigations into the effect of phase errors in the signal processing chain from each antenna on the beam pattern were also carried out. The phase error applied to each antenna is generated from a Gaussian distribution; the standard deviation of the Gaussian distribution is taken as the overall phase error of the telescope. The errors are uncorrelated between antennas. The detrimental effects from the uncorrelated errors will, to an extent, cancel out across the station, however there will be a residual error which will affect station performance.

Phase error standard deviations from $0 - 180^\circ$ are used in 56 logarithmically spaced steps. The phase error of each antenna was regenerated for each of the 1000 simulations conducted for each station size. Telescope models of three different station sizes were used: 16 (4×4), 256 (16×16) and 16 384 (128×128) antennas. Gridded station layouts were used. This is different to the randomised layouts specified for SKA1-Low, however these simulations compliment the work of Schediwy et al. [91] by extending their analysis to 2-D arrays.

Schediwy et al. use three linear arrays of 4, 16 and 128 antennas. The station sizes used in this study extend these linear arrays into 2-D grids, with 4^2, 16^2 and 128^2-antenna stations. These station sizes are also relevant to current and future radio telescopes. A 16-antenna station is the size of a tile in the MWA, 256 antennas is the current specified station size for SKA1-Low [96] and 16 384 antennas is of order

the number of antennas per station originally planned for SKA1-Low [95] (and of order the recommended size by some studies [87]).

The loss in sensitivity is measured by comparing the brightest pixel in a simulated beam pattern with a given phase error with that of a beam pattern from an ideal telescope with the same number and layout of antennas but no phase error.

The beam-pointing offset is determined by comparing the position of the brightest pixel in a simulated telescope model with phase errors to that of a telescope with no phase error. The ideal telescope's HPBW is determined from a 1-D cut through the beam centre; the HPBWs are used as a reference for the size of the pointing offset.

2.4.1 Results of phase error simulations

There are power law relations for the loss in sensitivity from 0° to $\sim 20^\circ$ phase error, beyond which greater phase errors have diminishing additional detrimental effects on sensitivity, as shown in Fig. 2.4a.

The loss in sensitivity for each phase error increases with the number of antennas in a station. The two larger stations, 128 and 16 384 antennas, give results within $\sim 2\%$ of one another and $15 - 20\%$ of the 16-antenna station. This, along with the power law results, is consistent with the results found by Schediwy et al. However, the 2-D gridded stations show a greater loss in sensitivity than the linear arrays in Schediwy et al., for a given phase error. The 16-antenna 2-D grid loses approximately twice as much sensitivity as the 16-antenna linear case in Schediwy et al. In both cases antennas were separated from adjacent antennas by the half-wavelength distance, although the difference may (at least partially) arise from different antenna beam patterns; here a square cosine taper was applied to an ideal dipole beam pattern; it is unclear how Schediwy et al. modelled their antennas.

According to the *Requirements specification* the difference between the parameterised station beam model and the actual station beam is required to be no greater than $\sim 1\%$ relative to the main beam peak power. To keep the sensitivity loss below

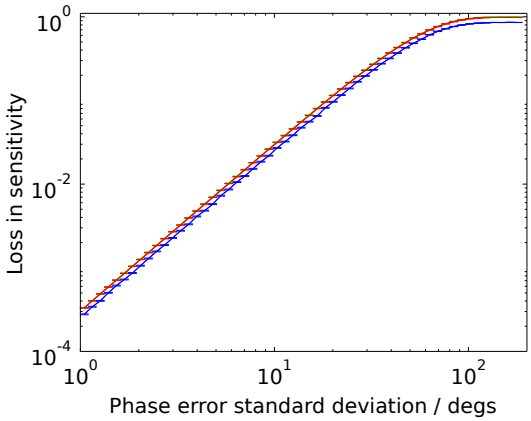

(a) Loss in beam sensitivity, relative to beam pattern with no phase errors.

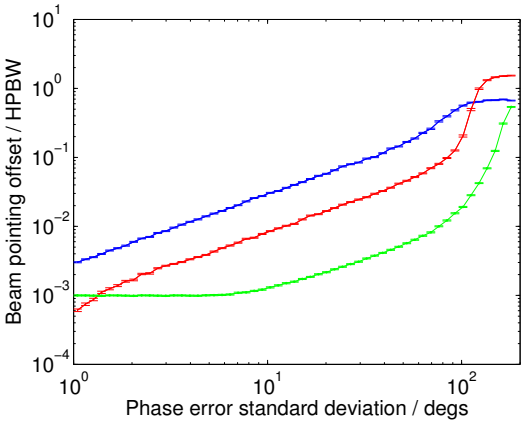

(b) Beam offset from pointing direction as a fraction of half power beam width (HPBW) of the ideal beam.

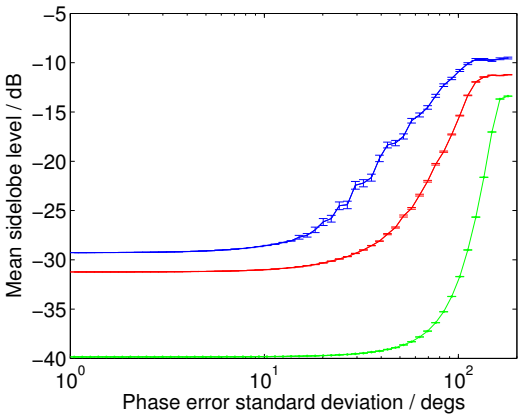

(c) Mean sidelobe level, relative to peak power of a beam pattern with no phase errors.

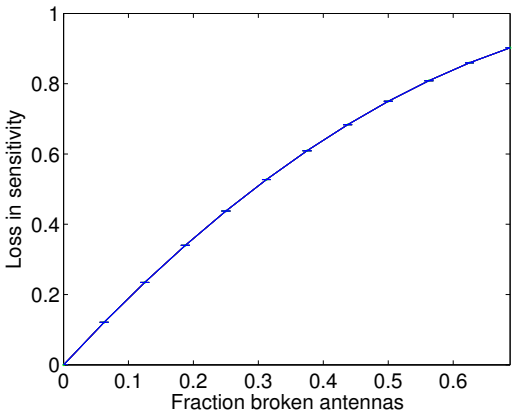

(d) Loss in beam sensitivity, relative to beam pattern with no broken antennas. Results for the three station sizes overlay one another almost exactly.

Figure 2.4: Beam pattern metrics as a function of standard deviation of antenna phase errors and broken antennas. Gridded telescope layouts were used. The blue line is for a 4×4 antenna layout, red 16×16, and green 128×128.

1%, phase errors needs to stay below $\sim 5.8°$ for the 256 and 16 384-antenna stations and $\sim 6.4°$ for 16-antenna stations.

The beam-pointing offset simulations, plotted in Fig. 2.4b, also show a general power law trend over a range of low phase errors. However there is a clear deviation from the power law trend at low phase errors for the 16 384 and 256-antenna stations. This can be attributed to the difficultly of accurately measuring and interpolating

the change in position of the beam's brightest pixel. Again, the results are consistent with Schediwy et al., including the deviation from the power law at low phase errors.

At high phase errors, the beam offsets increase at a significantly larger rate than the power law. This is also present in Schediwy et al. and is a result of the large phase errors significantly distorting the beam patterns. Furthermore, the large loss in sensitivity, combined with this distortion could result in some sidelobes becoming more powerful than the main lobe. This would result in the perceived position of the main lobe, as measured by the brightest pixel in the station beam pattern, rapidly shifting in position and hence offset. However, such large phase errors are unrealistic in practice.

To keep the beam-pointing offset below 1% of the HPBW, the phase error is required to be kept below $\sim 3.5°$, $\sim 15°$ or $\sim 80°$ for a 16, 128 or 16 384-antennas station, respectively. These levels of phase error are quite achievable in practice.

The mean sidelobe level, shown in Fig. 2.4c, is resilient to change for phase errors up to $\sim 10°$. The phase errors may be moving the power in the sidelobes around, but the overall power and hence the mean power may have stayed approximately consistent. Beyond the approximately constant sidelobe level regime, the sidelobe level rises very sharply with phase error, potentially caused by the power in the main lobe being spread to the sidelobes.

2.5 Broken antenna simulations

Similar simulations are conducted with selected antennas 'turned off' by setting their gains to zero, to allow a measurement of the effect of broken antennas in the station. The same sized gridded station layouts were used as the previous phase error simulations. Beam patterns were simulated for 0 to 75% of antennas broken, with 1000 simulations for each station size and number of broken antennas. The antennas to be broken are selected randomly for each iteration of the simulation, without repetition,

so that each simulation has a unique set of broken antennas. In some cases the number of possible unique samples of antennas, such as a single broken antenna out of a station of 16 antennas, is fewer than 1000, meaning the maximum possible number of unique samples are taken instead.

2.5.1 Results of broken antenna simulations

The effect of broken antennas on sensitivity is shown in Fig. 2.4d. The loss in sensitivity is non-linear, contrary to what might be expected. The largest proportional loss in sensitivity, however, is at low numbers of broken antennas: 1% of antennas broken causes a $\sim 2\%$ loss in sensitivity and 5% broken leads to a $\sim 10\%$ loss, but increasing the fraction of broken antennas to 30% and 40% increases the loss in sensitivity to 51% and 64%, respectively, a proportionally lower loss.

Broken antennas distort the beam pattern, turning the fully dense array into a partially sparse one, for instance, which may explain the non-linearity. The result of a given fraction of broken antennas is the same, regardless of station size.

SKA1-Low stations are unlikely to have stations with a large fraction of broken antennas, as broken antennas should be fixed or replaced. Alternatively data from a whole station may simply be flagged and unused if a significant fraction of its antennas were broken and not able to be corrected. The large sensitivity losses at low numbers of broken antennas suggests that SKA1-Low will need to ensure either its antennas are durable and robust or that there is an efficient system in place to detect and replace broken antennas, as they may otherwise prevent SKA1-Low detecting weak sources.

These simulations do not take into account the possibility of broken antennas creating spurious signals or otherwise inaccurate data. It would be prudent to periodically check signals from individual antennas to ensure this does not occur.

2.6 Summary

It is shown that increasing the number of antennas in a station will reduce the sidelobe level of the station beam and mitigate the impact of phase errors on the performance of the primary beam. For the beam sensitivity to be reduced by less than 1% and pointing offset to be within 1% of its HPBW, the phase error should be less than $\sim 3.5^\circ$ for a 16-antenna gridded station and $\sim 6^\circ$ for one with at least 128 antennas in a gridded layout. The mean sidelobe level of a station does not increase significantly for phase errors below $\sim 10^\circ$ for gridded stations with 16 or more antennas, but it is probable that the sidelobe positions are shifted. In addition to restricting phase errors, broken antennas in stations should be corrected for swiftly (within the economical means of the telescope) as they can strongly affect the sensitivity in a station, at a rate faster than the simply the fraction of broken antennas.

SKA1-Low stations will have a randomised layout as opposed to the gridded layouts used in the phase error and broken antenna simulations presented here. We would expect qualitatively similar performance from randomised layouts as the gridded layouts, although they may be quantitatively different.

The level of tolerable gain and phase errors seem quite achievable when considering the effect on simple metrics of individual station patterns. However this does not take into account the full effects of interferometric imaging, which will be considered in the next chapter.

Chapter 3

Effect of array errors on imaging

3.1 Motivation

SKA1-Low is proposed to cover the frequency range $50 - 350$ MHz. At the lower end of the frequency band images are created with relatively wide station beams, and hence also wide near-in sidelobes. This could present difficulties in creating the large imaging dynamic range desired by SKA1-Low.

It has been suggested that the upper end of the frequency band could be extended to 650 MHz by sampling in the second Nyquist zone of detected signals, and that this could be a relatively inexpensive upgrade to the current instrument design [97]. However, there are concerns over the fidelity of the images that would be produced at such high frequencies due to the very sparse sampling of the aperture array [98].

The first revision of the SKA's *Requirements specification* [99] called for a critical sampling frequency of 110 MHz for SKA1-Low, which is followed in this chapter unless otherwise stated. This specification has since been excised, however the current antenna design requires an average spacing of 1.93 m [83], which corresponds to an approximate dense-sparse transition frequency of 78 MHz, implying any problems caused by very sparse sampling of the aperture array may be worse than simulated here.

The aim of this chapter is to investigate how imaging is affected across a frequency range of $50 - 600$ MHz by simulating the effect of phase and gain errors on the

images produced by an interferometer matching the SKA1-Low specifications. This investigation compliments the work of Razavi-Ghods et al. [100] which assesses the far-sidelobe confusion noise of sources in an ideal SKA1-Low telescope. In addition to modelling phase and gain errors, baselines beyond the SKA1-Low core are included and the frequency range extended down to 50 MHz.

3.2 Telescope models

The Nyquist sampling theorem states that a band-limited function can be completely determined by sampling it at a frequency of twice the bandwidth. Hence to sample SKA1-Low signals between $350 - 650$ MHz it is possible to filter signals to only this band-pass and sample at 600 MHz; this is known as sampling in the second Nyquist zone.

The Nyquist sampling theorem can also be applied in spatial terms, where it corresponds to sampling at every half-wavelength of the shortest wavelength present. This means that the distance between antennas in an aperture array is significant: if antennas are spaced further than half the observed wavelength apart, the station beam is spatially aliased, resulting in strong sidelobes (see, for example, [101]). The frequency at which this transition occurs is the critical sampling frequency.

A source located in a sidelobe of a radio telescope is detected at a level dependent on the sidelobe strength. It is theoretically possible to remove the signal from sources in far-out sidelobes in synthesised images with post-processing. However, this requires a perfect understanding of the telescope's behaviour, and of the sources in the sky. When uncertainties are present in the signal processing chain, the resultant data is perturbed from the error-free case. By investigating the differences between images from an ideal (i.e. error-free) telescope and those from a telescope with errors, the effect of errors on the images can be quantified.

The OSKAR-2 simulator is used to generate simulated visibilities for all cases, while CASA is used to make images of simulated visibilities.

To simulate SKA1-Low, a telescope model is created with station locations matching the description given in the *Requirement specifications, 1st revision*. The layout and instantaneous *uv*-coverage are shown in Fig. 3.1.

The *Requirements specification* describes a telescope containing 1024 thirty-five metre diameter stations, with 256 antennas per station. Here 768 ($= 3\times256$) antennas per station are used, giving stations a factor $\sqrt{3}$ times the specified diameter. The station beams are apodized using a Hanning window such that their HPBW is equal to a 256-antenna unapodized station. This is discussed in [89] but is briefly summarised here.

The baseline design of 35 m diameter separate stations assumes all antennas within the station have equal weight, in order to achieve the required total effective area of the array. However, a station with uniform weighting generates high sidelobes in the station beam, significantly increasing the sensitivity to off-axis sources. Apodizing stations to improve the sidelobe level results in reduced overall collecting area and larger station beams. Alternatively a scheme of apodized overlapping stations in a contiguous distribution of antennas can give efficient use of collecting area along with good beam performance. The scheme used in this chapter gives a good approximation to this concept in so far as utilising apodized stations to reduce sidelobe levels, but a contiguous distribution of antennas is not used; this concept is discussed further in Chapter 4.

A randomised antenna layout which is common to each station is used, but each of the 1024 stations are rotated to a unique angle, $(360n/1024)^\circ$, where n is an integer between 0 and 1023, so that the position of each antenna within a station varies relative to the centre; the orientation of each individual antenna remains the same, however. This rotation mitigates the effect of non-radially symmetric sidelobes, hence reducing the sensitivity to off-axis sources, when correlating the station beams. This approach is used by LOFAR [25].

In line with the *Requirements specification*, log-periodic antennas with a directivity of almost 8 dBi are used. Antenna models matching those of the current SKA1-Low

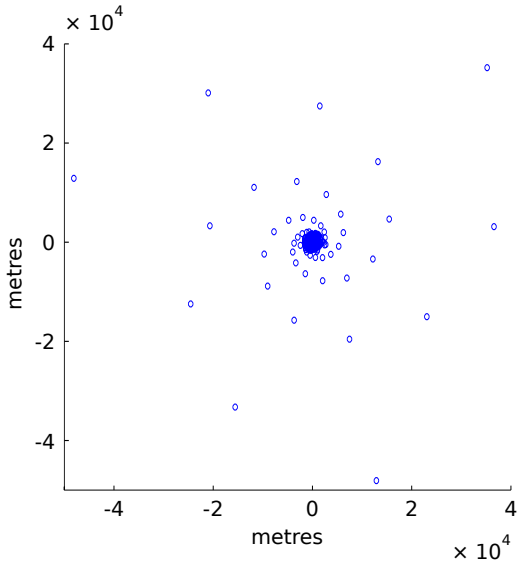

(a) Layout of the telescope model used. There are 1024 stations, with 50% within 600 m of the centre and 75% within 1 km. There are three spiral arms of 15 stations each.

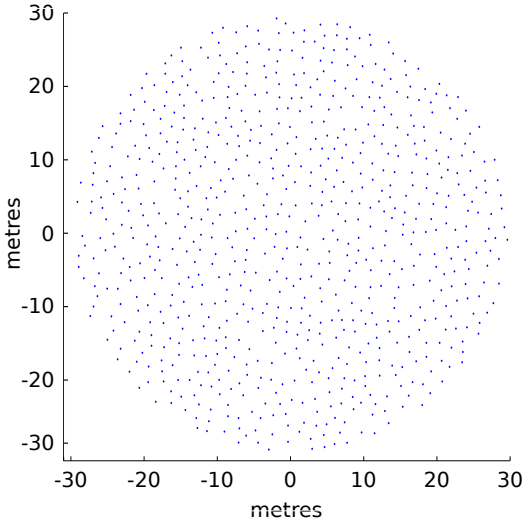

(b) Layout of one station. 768 (= 3 × 256) antennas within a $\sqrt{3} \times 35$ m diameter circle, randomly distributed, but separated by a minimum of 1.35 m (critically sampled at 110 MHz). Every other station has the same layout, but the layout is rotated around the station centre by a unique angle.

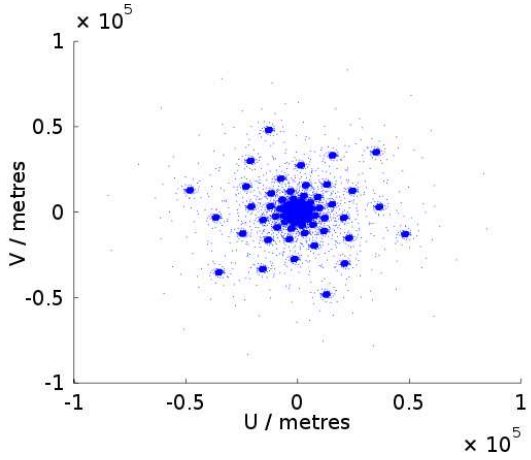

(c) Snapshot *uv*-coverage (at the zenith) of the telescope layout.

Figure 3.1: The telescope model used for the simulations in this chapter. Station locations conform to the stipulations in the *Requirements specification* [102].

antennas ('SKALA antennas') are used at each frequency simulated, except at 50 MHz where an antenna model is not available; the antenna response at 100 MHz is used instead; log-periodic antennas are designed to have approximately the same antenna beam pattern across their frequency range. Antennas have a minimum separation of 1.35 m, corresponding to a critical sampling frequency of $\sim$110 MHz specified in the

original baseline design.

The SKALA antenna patterns [103] are expected to vary by up to $\sim 3\,\mathrm{dB}$ across the antennas in a station, however the average antenna beam pattern of a station is well constrained from one station to the next [104]. These variations are not simulated; only the mean antenna patterns are available. Furthermore, OSKAR does not model mutual coupling between antennas.

3.3 Sky models

Two metrics for measuring the impact of uncorrelated gain and phase errors on imaging are now defined: residual in-field noise (RIN) and residual sidelobe noise (RSN). These are root mean square (RMS) deviations between a dirty image made from an ideal, error-free observation and a 'dirty image' (the synthesised sky intensity distribution) from an identical observation with gain and/or phase errors in the signal processing chain. RIN occurs from sources located within the HPBW of the station beams and RSN occurs from sources in the sidelobes of the station beams.

In the simulations, the difference between the gain/phase error-affected image and the ideal image is entirely due to the gain and/or phase errors present. These metrics hence quantify the negative effect of gain and phase errors on synthesised images.

Four sky models are used in these simulations, which are divided into two types. The first type of sky model has sources placed only in the main lobe of the station beam; no sources are placed more than the HPBW away from the observation's phase centre, allowing the RIN to be measured. This boundary distance is decreased proportionally with the reduction in field of view at higher frequencies. In the second type of sky model all sources within twice the HPBW of the station beam are removed; sources populate the rest of the sky, allowing RSN to be investigated. Further details of the sky models are given as follows:

- RIN: sources only within the HPBW

  - Model 1: A grid of nine 1 Jy sources. The on-sky separation between the sources is scaled with the observation frequency, such that the nine sources are located on the same image pixels at every observation frequency.
  - Model 2: A point-source sky model using data from the VLA Low-Frequency Sky Survey (VLSS). All sources outside the station beam are removed. The VLSS was measured at 74 MHz; a spectral index of $\alpha = -0.7$ (with $S(\nu) \propto S^{\alpha}$, where S is the flux) is hence applied to all sources to determine their fluxes at different frequencies. This spectral index is within the typical range for synchrotron radiation in this frequency range [52]). Figure 3.2 shows how the sources included in the sky model change with frequency.

- RSN: No sources within twice the HPBW.
  - Model 3: A grid of 1 Jy sources placed at 10° intervals in RA, Dec across the sky, except in the direction of the synthesised beam. Source separation does not scale with frequency.
  - Model 4: The VLSS sky model with all sources within the HPBW removed.

The VLSS catalogue sky model is a low-frequency radio survey of the northern sky conducted by the VLA, covering 95% of the sky above latitude −30° (most of the gaps in the survey are between −20 to −30° latitude). It provides a list of ∼ 70 000 point sources detected above 5σ from the noise level (∼ 0.1 Jy/beam).

The six brightest sources (Cassiopeia A, Cygnus A, Taurus A, Hercules A, Virgo A, Hydra A) are removed from the VLSS sky model; these sources account for approximately 40.5 kJy of total flux density at 74 MHz. These six strong radio sources have the potential to unduly distort the results if they remain in the sky model. For example, if Cassiopeia A is located in a strong sidelobe, it could significantly affect the measured results compared to a weaker source, which are far more common and thus more typical of future SKA1-Low observations. All bright sources, however, will

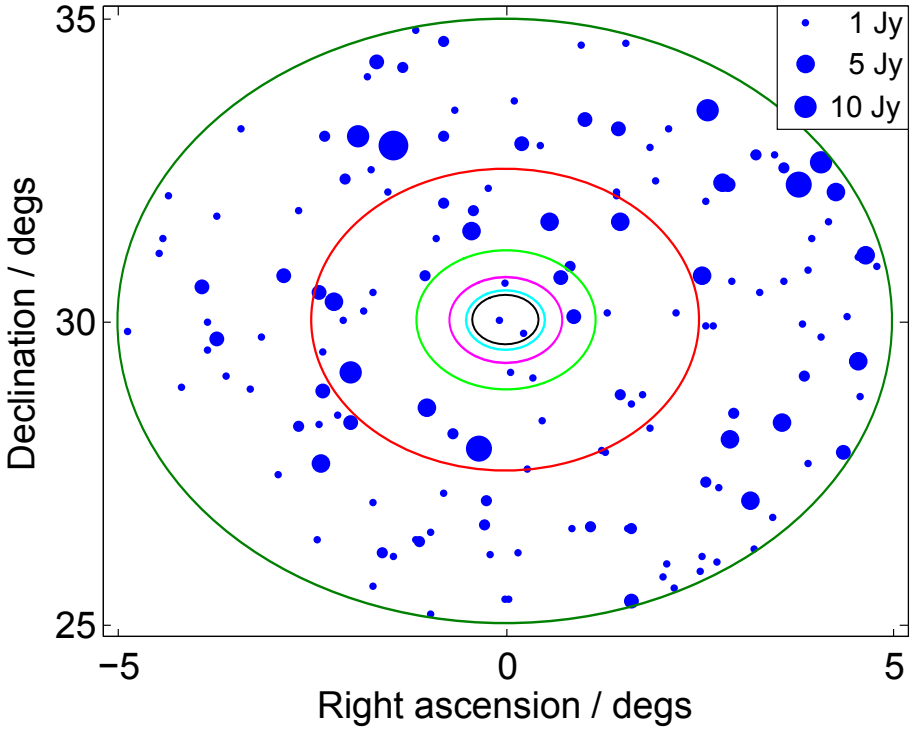

Figure 3.2: Half power beam width observed by the station beam of the telescope. The dots represent sources in the VLSS sky model, with their size corresponding to the source strength, as shown in the legend. Source fluxes are shown at 74 MHz; there is a power law spectral index of -0.7 applied to determine fluxes at other frequencies. The outer (dark green) circle shows the field of view at 50 MHz, with subsequent internal circles increasing in frequency (100, 220, 350, 500, 600 MHz).

need to be well understood and accounted for when conducting actual observations with the SKA.

3.4 Simulations

The simulations use a phase centre at the zenith (except where stated otherwise) and the array centre positioned at $(0, 30)$ deg in (longitude, latitude), the approximate latitude of the SKA1-Low site (albeit in the opposite hemisphere; this is to allow the use of the northern sky VLSS catalogue as a sky model). Simulations are run at 50, 100, 220, 350, 500 and 600 MHz, each with 13.7 s of integration (except where otherwise stated) and 100 kHz bandwidth. 13.7 s corresponds to approximately 1 milliradian of earth rotation. Images are made with Stokes I.

Uncorrelated uncertainties in the gain and phase are added to the signals 'detected'

at each antenna. These represent the cumulative errors acquired along the signal chain. The uncertainties are assigned from random values selected from Gaussian distributions and are different for each antenna. The error values are kept constant throughout the integration time, whether it is a 13.7-s snapshot or $\sim$1-hour long observation. However, the error values are regenerated for repeated iterations to determine errors on the images using the Monte Carlo method.

The standard deviations of the uncertainties are the error values; for instance a phase error of 1° means that the phase value applied to an antenna is $\phi \pm \phi_{\rm err}$, where ϕ is the ideal phase value and $\phi_{\rm err}$ is the error value, which is randomly selected from a Gaussian distribution with a standard deviation of 1°. Likewise a gain error of 10% represents an additive error selected from a Gaussian distribution of 0.1 times the true value. For large gain errors this means that the gain is potentially negative. This is only significant for gain errors of order unity, which are unrealistic, but are included in these simulations to show the upper bound on the effect of such errors.

3.5 Thermal noise

It is useful to compare the RMS deviations to the thermal noise level of the observations. The thermal noise level is the weak-source limit of the telescope, per polarization, and is given by

$$S_{\rm min} = \frac{1}{\eta_s} \frac{SEFD}{\sqrt{2\,\Delta\nu\,t}}\,, \tag{3.1}$$

where η_s is the system efficiency factor, $\Delta\nu$ is the bandwidth and t is the integration time. *SEFD* is the system equivalent flux density and is calculated by

$$SEFD = T_{\rm sys} \frac{2k_B}{\eta_A A_e}\,, \tag{3.2}$$

where k_B is the Boltzmann constant, η_A is the antenna efficiency (assumed to be 0.9), A_e is the effective area of the antenna and $T_{\rm sys}$ is the system temperature [105].

The SKA1 *Baseline design* [60] gives a model for $A_{\rm e}$ per antenna of 3.2 m^2 below 110 MHz and decreasing in proportion to λ^2 above. This model has been followed in this chapter.

The *Baseline design* also gives a model for $T_{\rm sys}$:

$$T_{\rm sys} = T_{\rm sky} + T_{\rm rcvr}\,, \tag{3.3}$$

where $T_{\rm sky}$ is the sky temperature and $T_{\rm rcvr}$ is the receiver temperature, which is modelled as $0.1T_{\rm sky} + 40\,{\rm K}$.

The Milky Way is the brightest radio background at these frequencies and thus its temperature can be considered as a good approximation for $T_{\rm sky}$. The Haslam 408-MHz all-sky map [106, 107] provides temperatures for the galactic foreground at this frequency. By measuring the average value of pixels within a station beam-sized area, $T_{\rm sky}$ can be determined for any telescope pointing direction in the sky.

The temperature within a station beam-sized area of a cold patch of sky (i.e. away from the galactic plane), along with the temperature in a typical temperature area of sky can hence be determined by the Haslam map.

To convert the 408-MHz temperatures to other frequencies a power law of

$$T(\nu_2) = T(\nu_1)\left(\frac{\nu_2}{\nu_1}\right)^{\beta}\,, \tag{3.4}$$

for β = -2.4, -2.55 and -2.7 was assumed, which covers the approximate variation in the spectral index over this frequency range.

By combining the range of temperatures given by the typical and cold patches of sky, along with the different values of β, upper and lower bounds on $T_{\rm sky}$ are determined.

The antenna $A_{\rm e}/T_{\rm sys}$ values found from the effective areas and system temperature outlined above are consistent with the range of values given in de Lera Acedo et al. [83], which gives estimates for the SKALA antenna $A_{\rm e}/T_{\rm sys}$ values.

The thermal noise level is the theoretical best weak-source sensitivity of the telescope. Hence if the noise from gain and phase errors exceeds this level, the performance of the telescope will be impaired. A comparison of the two provides an upper bound on the allowable gain and phase errors.

Due to computational resource limits, simulations are mostly run with snapshot observations of 13.7-s duration. The RMS thermal noise level in a radio astronomy

observation will decrease with the square root of the length of observation. One may also expect the RSN to also decrease with $t^{\frac{1}{2}}$, at least initially, as the rotation of the Earth moves sources through the sidelobe errors of the telescope in a quasi-random way. This effect should continue for around 12 hours until the uv-tracks of the baselines begin to repeat themselves and the new noise components are not independent of those previously observed. To check that this occurs, further observations of the RSN using the VLSS sky model are simulated with integration times of up to an hour (longer integrations were not simulated due to limits of available computing facilities.

These simulations model only the effect of gain and phase errors and not all the causes of noise. When accounting for different effects, unpublished preliminary simulations suggest the noise may decrease at a slightly slower rate than $t^{\frac{1}{2}}$ [108].

3.6 Analysis of simulations

3.6.1 Residual in-field noise

The addition of uncorrelated gain and phase errors to the antennas results in distortions of the station beam patterns, which adds errors to the images, ultimately increasing the RMS deviation between synthesised images from a telescope with errors and images from one without. An example of an ideal station beam pattern, a beam pattern with gain errors of 100%, and the difference between the two, are shown in Fig. 3.3.

The results for the dense grid sky model (sky model 1) shown in Fig. 3.4 are a useful test case. The RMS deviation between the ideal and error images remains very nearly constant across the entire frequency range. As this sky model contains only sources in the station beam, where the source separation scales perfectly with the field of view, this result is as expected; sparse stations cause differences in station beam sidelobes, rather than the main lobe of the beam.

An example of the difference between the ideal and error images for the dense sky model is shown in Fig. 3.5b, alongside the ideal image in Fig. 3.5a. The largest differences between the ideal and error images are found at the locations of the sources.

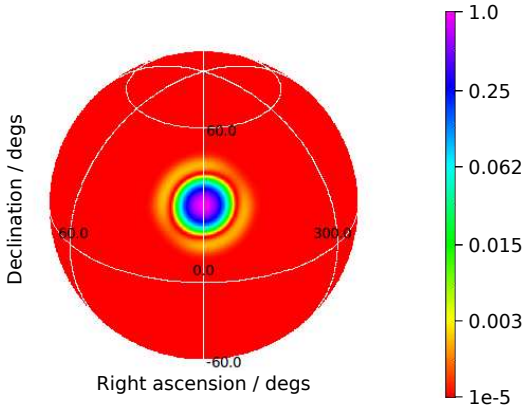

(a) Station beam pattern of ideal telescope

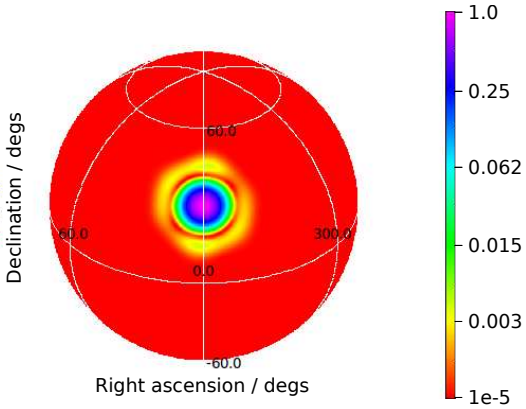

(b) Station beam pattern of telescope with gain errors of 100%

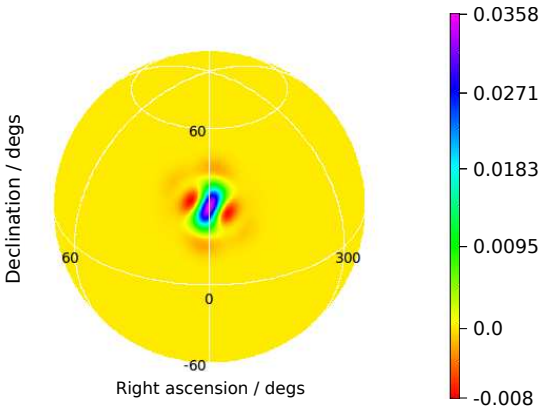

(c) Difference between Fig. 3.3a and Fig. 3.3b

Figure 3.3: Station beam patterns for ideal telescope, one with errors, and the difference between the two. The colour scale shows the gain of the beam. Note the change in scales between Fig. 3.3a and Fig. 3.3b (logarithmic), and 3.3c (linear).

This is not surprising, as these pixels measure the largest flux values; a 1% error, for example, will cause a much greater absolute discrepancy here than in pixels recording near-zero flux levels. As the noise is not randomly distributed across the image, it may call into question the extent of the applicability of using RMS deviation as a metric, which should strictly be used for Gaussian noise. However, the RMS deviation is used here purely as a measure of the difference between the two images' respective pixel values; other metrics such as the total difference or maximum difference would show the same trend in results.

Figure 3.6 shows the results for sources in the station beam from the VLSS sky model (sky model 2). There is a trend of decreasing RMS deviation with increasing

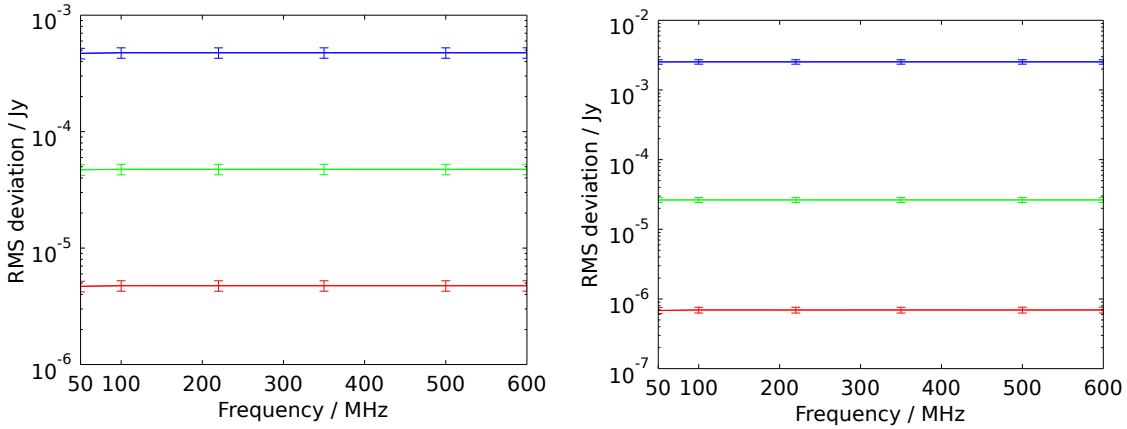

(a) Gain errors of 100% are shown in blue, 10% in green and 1% in red

(b) Phase errors of 10° are in blue (error bars ×100) 1° in green (error bars ×20) and 0.1° in red (error bars ×2)

Figure 3.4: RMS deviation between images simulated by an ideal telescope and one with errors for a sky model of a 3×3 grid of 1-Jy point sources (sky model 1) in the station beam plotted against observation frequency. The separation between the sources is decreased with frequency, such that the source centres occupy the same pixels at every frequency. An example image from the ideal telescope is shown in Fig. 3.5a.

frequency. As the field of view decreases at higher frequencies, so do the number of sources in the field of view. As shown above, it is the pixels corresponding to the source locations that contribute the most to the RMS deviation, so fewer sources in the smaller beams leads to smaller RMS deviations. Furthermore, the spectral index of the VLSS sky model is set to -0.7, so the flux of the sources in the beam is diminished at the top of the frequency band relative to the flux at the bottom. As a semi-realistic sky model is used here, the RMS deviations give an approximation for the lower bound of source fluxes measurable from on-sky observations of 13.7-s duration with 100-kHz bandwidth.

3.6.2 In-field flux density and the RMS deviation

To investigate the relationship between the total flux density in the field and the RMS deviation, further simulations were run, the results of which are shown in Fig. 3.7. These simulations restrict all sources to be within 80% of the HPBW, to minimise the differences in sensitivity depending on the source positions in the beam.

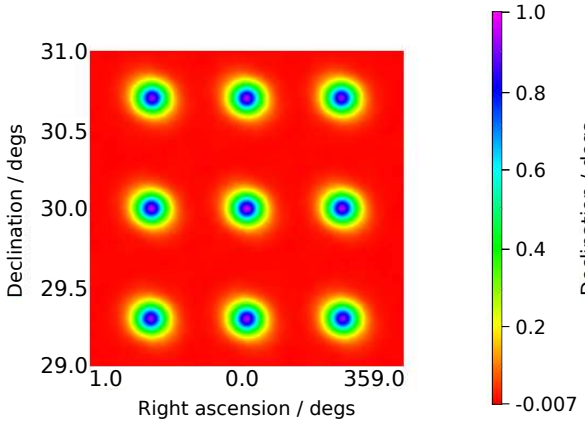

(a) Simulated image of a dense grid of sources positioned within the station beam of the telescope (sky model 1). The sources are separated by 0.7° and the observation simulated at 50 MHz. There are no phase or gain error effects included in this image.

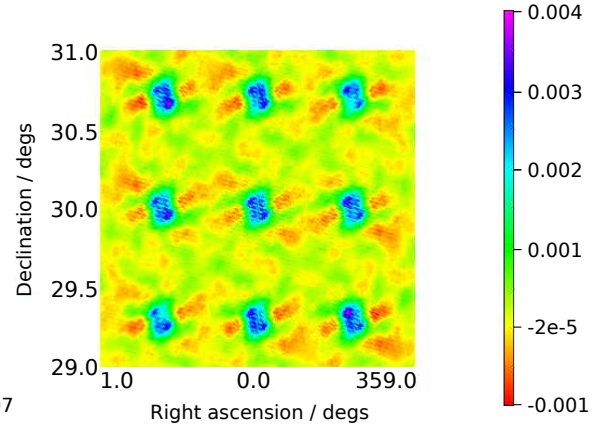

(b) The difference between Fig. 3.5a and an image taken with a telescope with antenna gain errors of 100%.

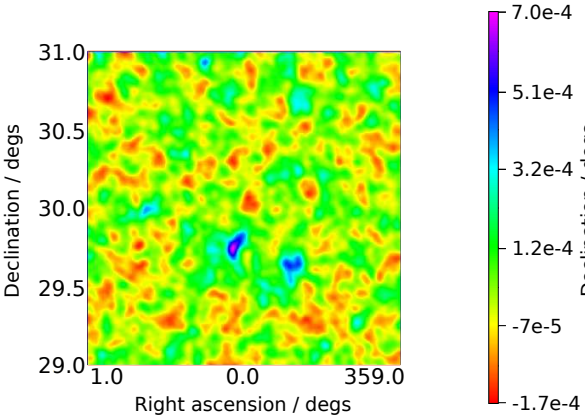

(c) Simulated image of the VLSS catalogue sky located in the sidelobes of the station beam pattern, with the phase centre pointed at an empty patch of sky (sky model 4). No sources are located within twice the HPBW of the observation's phase centre. Image simulated at 100 MHz. There are no phase or gain error effects included in this image.

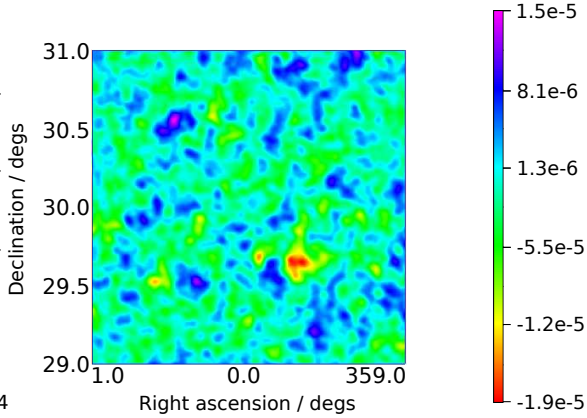

(d) The difference between Fig. 3.5c and an image taken with a telescope with antenna gain errors of 100%.

Figure 3.5: Simulated images for two different sky models. Ideal images and the difference between ideal and error images are shown.

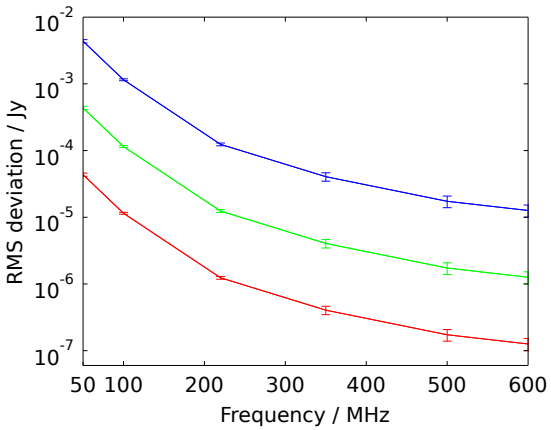

(a) Gain errors of 100% are shown in blue, 10% in green and 1% in red

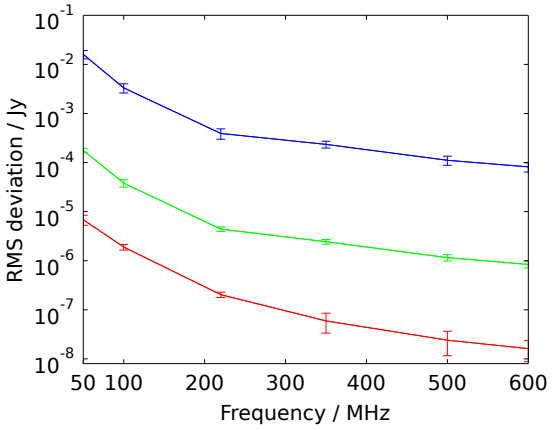

(b) Phase errors of 10° are in blue (error bars ×100), 1° in green (error bars ×20) and 0.1° in red (error bars ×10)

Figure 3.6: RMS deviation between images simulated by an ideal telescope and one with errors for a sky model from the VLSS catalogue. Sources beyond the HPBW were removed from the sky model, as shown in Fig. 3.2. The simulated telescopes were pointed at the zenith.

When separate 1-Jy sources are added to the field of view (in a gridded pattern, like the sources shown in Fig. 3.5a), the sources' RMS deviation contributions added in quadrature, as the deviations are incoherent, resulting in a trend of RMS deviation $\propto \sqrt{B}$, where B is the total field brightness (dark blue line, Fig. 3.7). On the other hand, when the field brightness is increased by means of a single source with increasing flux, the RMS deviations were coherent and a linear relationship between RMS deviations and field brightness results (red line Fig. 3.7).

Further simulations use the same source positions as the 1-Jy gridded case (sky model 1), but with sources of differing flux levels. With weak sources (magneta line Fig. 3.7), the $\sqrt{B}$ trend is followed, but with stronger fluxes (green data Fig. 3.7), the relationship initially starts as somewhere between $\propto \sqrt{B}$ and B, before settling into the $\propto \sqrt{B}$ trend when there are a greater number of sources. Simulations using the VLSS catalogue are also shown in Fig. 3.7, the results of which follow the $\sqrt{B}$ trend once a sufficient number of sources have been included in the field; the low field flux data points are more susceptible to be skewed by individual strong sources.

From these simulations it is possible to predict that the approximate RMS devi-

ation produced by a large number of weak sources is given by

$$\Delta S = \left[\sum_{i}^{\text{sources}} (kS_i)^2 \right]^{1/2}, \tag{3.5}$$

where S_i are the flux values for sources in the field and k is a constant of proportionality. For these simulations, with a gain error of 10%, $k = (5.4 \pm 0.1) \times 10^{-7}$, as determined by χ^2 minimisation to the data sets (discounting the first two and three data points, respectively, for the strong sources in the grid (green line) and VLSS catalogue (black line) data sets, where strong sources skew the trend away from $\sqrt{B}$). The accuracy of Eq. (3.5) decreases when there are few sources, and when sources are far from the centre of the beam.

3.6.3 Residual sidelobe noise

The other class of sky models have no sources within twice the HPBW of the station beam, but populate the rest of the sky with sources. Figure 3.5d shows an example image of the difference between the ideal and error-case images, alongside the corresponding ideal image Fig. 3.5c. Figure 3.8 shows the results for the gridded model (sky model 3), where sources were spaced at 10° intervals. The RMS deviation is seen to stay approximately constant across the frequency range, with the values slightly increasing and decreasing across the frequency range. As the pattern of the sidelobes varies with frequency, different sources move in and out of the sidelobes, potentially explaining this variation.

The other RSN simulations use the VLSS point-source model (sky model 4). As the exact number of sources and their position with respect to the sidelobes affects the measured RMS deviations, observations are made in five different directions: $(0, 0)$, $(10, 10)$, $(20, 15)$, (-10, 30) and (-5, 20); values are degrees in RA and Dec away from the zenith.

As shown in Fig. 3.9, a spread of results can be seen depending on the beam direction, but each follows the same trend of decreasing RMS deviation with frequency. In Fig. 3.9b the RMS deviation is significantly greater in two observing directions at low frequencies than the other three, with the difference decreasing with frequency.

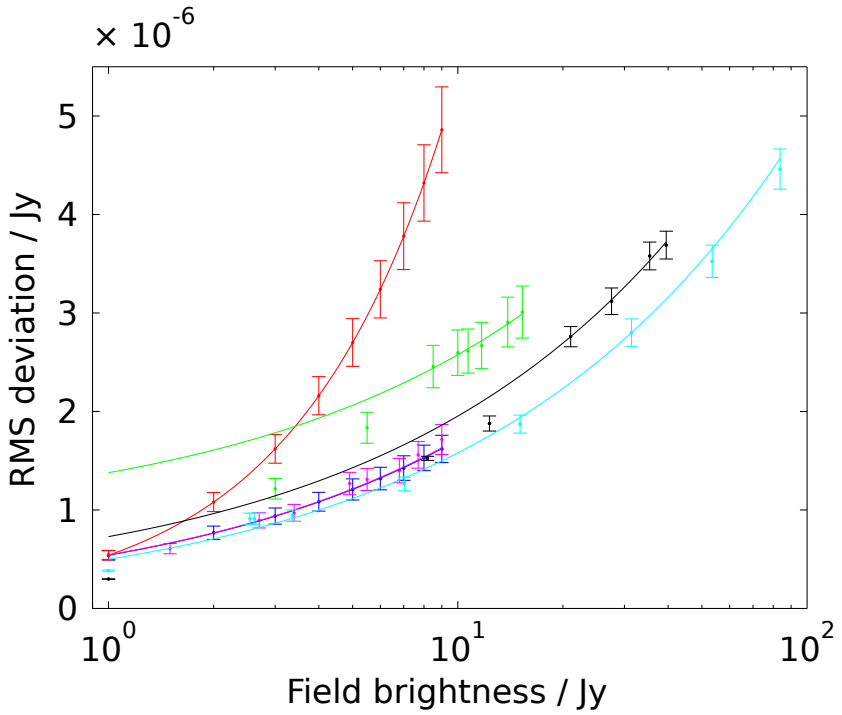

Figure 3.7: RMS deviation from residual in-field noise between ideal and error-telescope images, as a function of field brightness. The dark blue data points show the case of separate 1-Jy sources positioned in a grid (sky model 1) in the HPBW of the beam (trend line $\propto \sqrt{B}$, where B is the brightness in the field), whereas the red points are for the case when a single source of increasing brightness was used (trend line $\propto B$). The magenta and green data points also use the grid source locations, but with varying source fluxes: the green data shows stronger fluxes and magenta weaker. The magenta data points follow a $\sqrt{B}$ scaling and, after an initial trend lying between linear and $\propto \sqrt{B}$, the green data points scale with $\sqrt{B}$. The light blue data shows a $\sqrt{B}$ relationship for a randomly generated source fluxes and positions within 80% of the HPBW. The black data points show the values using a small number of sources from the VLSS catalogue, which also scales with $\sqrt{B}$ once a sufficient number of sources were included. Gain errors of 10% were used at 100 MHz.

This suggests a bright source was located in a strong sidelobe in these two observing directions, which caused a much larger RMS deviation to be measured. At higher frequencies (≥ 350 MHz) the beam size will have decreased such that the bright source was no longer in a strong sidelobe. This result supports the decision to remove the six brightest sources from the sky model, as their presence in a sidelobe could cause an even greater difference in the RMS deviations.

In general, the RMS deviations decrease approximately by a factor of 10 across the frequency range, with around 50% of the drop occurring between 50 and 100 MHz.

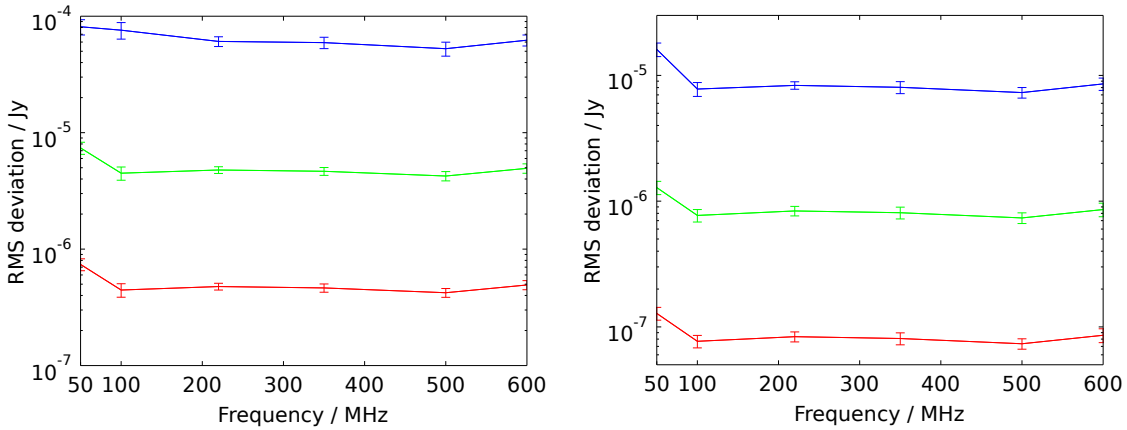

(a) Gain errors of 100% are shown in blue, 10% in green and 1% in red

(b) Phase errors of 10° are in blue, 1° in green and 0.1° in red

Figure 3.8: RMS deviation between simulated images by an ideal telescope and one with errors for a sky model of a sparse grid of point sources in the sidelobes. The simulated telescopes observed an empty patch of sky with a sparse grid of 1-Jy point sources, spread across the sky at 10° separations, in the sidelobes. An example image from the ideal telescope is shown in Fig. 3.5a.

This is partially attributable to the decreased sky brightness at higher frequencies, as shown by the black dot-dashed lines in Fig. 3.9. These lines display the RMS values according to the power law spectral index of -0.7, relative to the RMS deviation at 100 MHz.

The reduced average sidelobe level at higher frequencies could explain the rest of the reduction. The near-in sidelobes are the strongest sidelobes; at low frequencies they cover wide angles on the sky.

The magenta lines in Fig. 3.9 shows the approximate range of the thermal noise floor of the telescope model for a 13.7-s integration with 100-kHz bandwidth, as calculated by Eqs. (3.1) to (3.3). Thermal noise levels are plotted for both typical and relatively cold patches of sky. The thermal noise limit shows that gain errors of 10% and phase errors of 1° are comfortably within the thermal noise level, but gain errors of 100% and phase errors of 10° can produce noise levels exceeding the thermal noise. From inspection of these results, an approximate upper bound of $\sim 20\%$ in gain error and $\sim 5°$ in phase error would keep the RMS deviations below the thermal noise levels at all frequencies for snapshot images. It is important to stress that this is

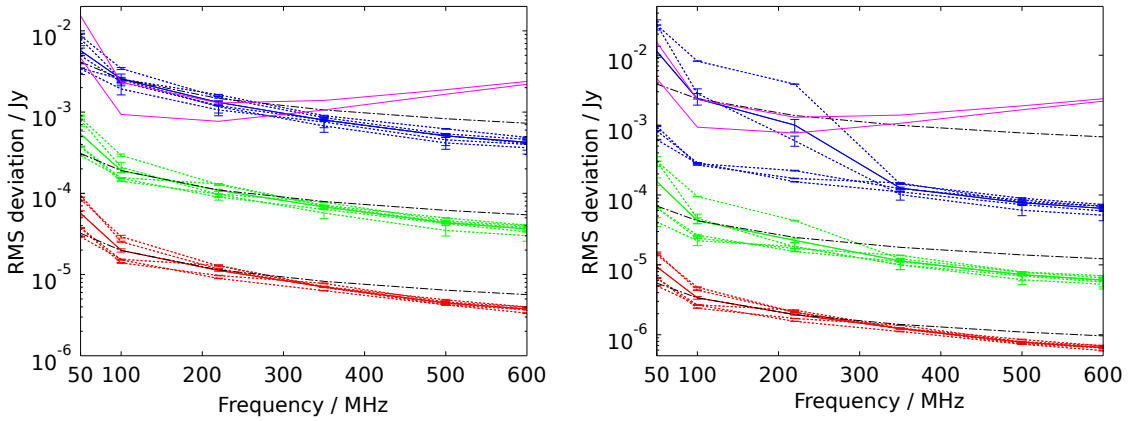

(a) Gain errors of 100% are shown in blue, 10% in green and 1% in red

(b) Phase errors of 10° are in blue, 1° in green and 0.1° in red

Figure 3.9: RMS deviation between images simulated by an ideal telescope and one with errors. The simulated telescopes were pointed at empty patches of sky with the VLSS point source sky model in the sidelobes. The size of the empty patch of sky was twice the HPBW of the station beam. Observations were taken in five different directions: $(0, 0)$, $(10, 10)$, $(20, -15)$, $(-10, 0)$, $(-5, -10)$, where the values are the (RA, Dec) away from the zenith in degrees. The dashed lines show different observing directions, and the mean is shown by the solid lines. The black dot-dashed lines show the RMS values if they were to decrease according to the spectral index (-0.7) of the sky model, referenced to the mean RMS value for all the observing directions at 100 MHz. The magenta lines show the probable range of thermal noise level at each frequency.

a best-case scenario, assuming the position and flux of all sources are known perfectly at each observation frequency.

The RMS deviations in Figs. 3.6 and 3.9 indicate that it is the low frequencies at which SKA1-Low should be most concerned to ensure the telescope can meet the sensitivities required by the science goals. The results suggest that at low frequencies, the large beam sizes (and hence large number of sources in the beams) and prominent sidelobes result in greater distortions of the visibilities due to errors in the telescope model. Based on these results, if images can be satisfactorily produced at low frequencies, then the higher frequencies should not present greater difficulties.

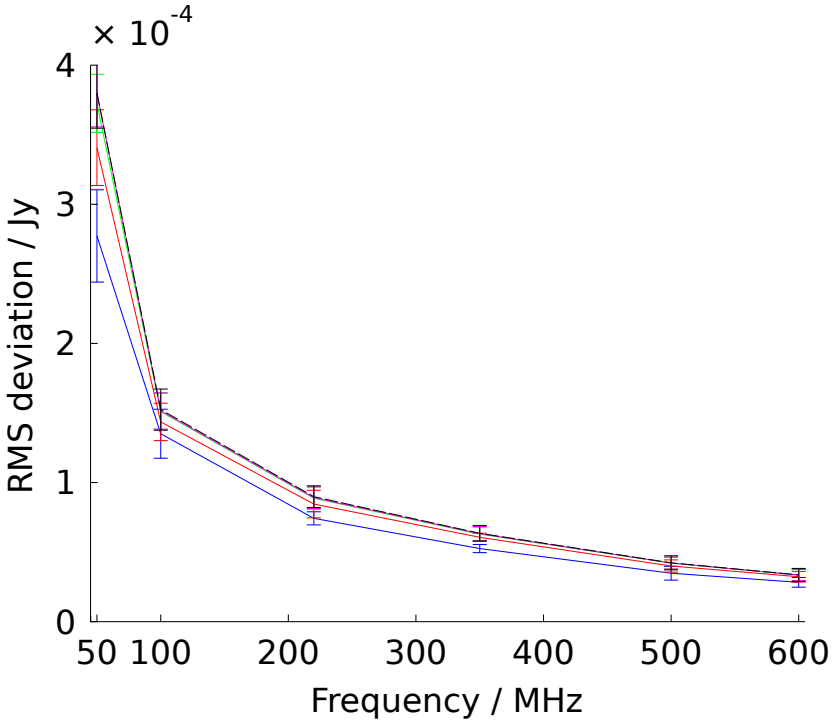

Figure 3.10: RMS deviation for sky models with different minimum fluxes in the sky model. The VLSS sky model was used, but only the brightest 1% (i.e. >14 Jy at 70 MHz), 10% (>3.7 Jy), 50% (>1.2 Jy) and 100% of sources were used; these are shown by the blue, red, green and black data points respectively. The magenta data points shows the RMS when all sources within double the weakest source flux in the VLSS sky model (i.e. sources <0.79 Jy) were removed. There is only a small difference between the 50% (green), 100% (black) and double the weakest source (magenta) data, hence their respective lines in the plot overlap. A gain error of 10% is used.

3.6.4 Far-field flux density and RMS deviation

To assess the impact of the strongest sources on the RMS deviation, further simulations can be run to remove varying fractions of the weakest sources in the sky model. Figure 3.10 shows that the RMS deviations are dominated by the brightest 10% of sources in the sky. This is encouraging, as in reality no sky model will be able to perfectly characterise all the sources in the sky; it will be most accurate with the brightest sources.

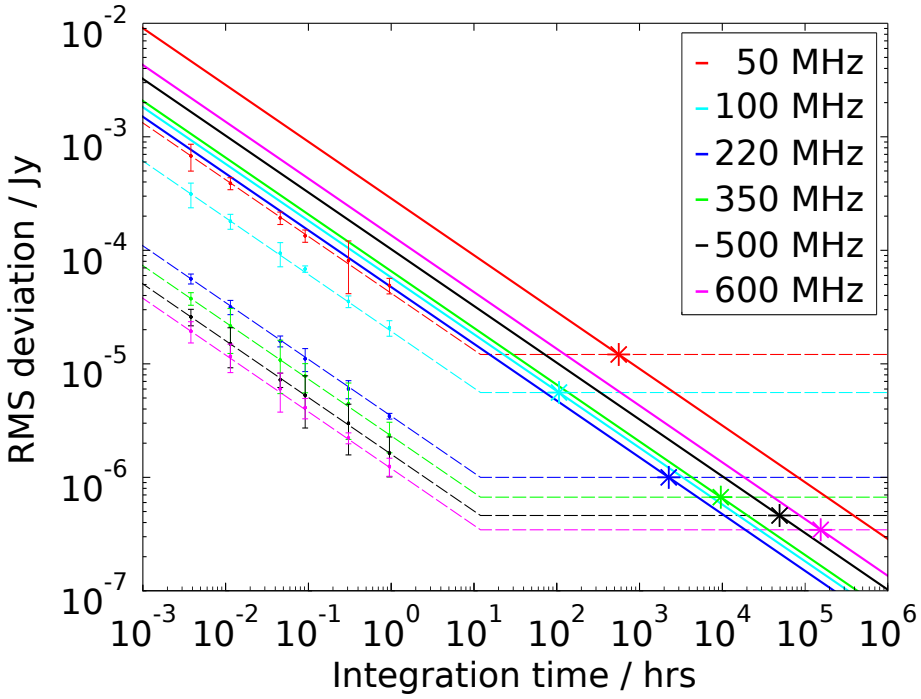

(a)

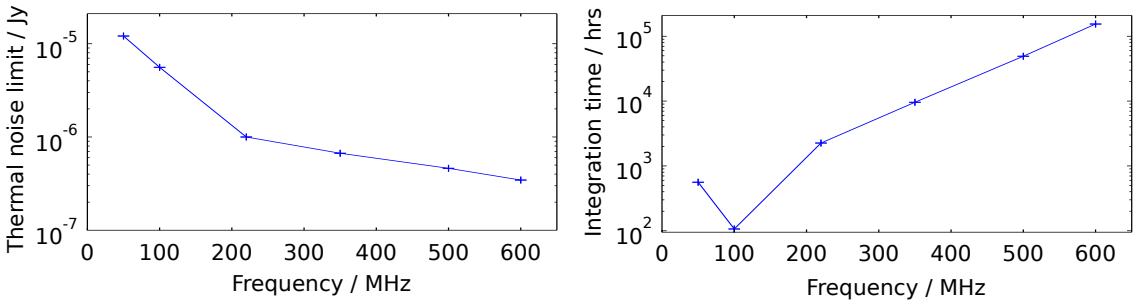

(b) (c)

Figure 3.11: (a) RMS deviation between the ideal and error image as a function of integration time shown by the data points (error bars $\times 10$) for a range of frequencies. The dashed lines through the data points show a $t^{-\frac{1}{2}}$ trend line. Thermal noise levels are shown in solid lines. The stars highlight the integration time at which the thermal noise drops beneath the RMS deviation; these values are also plotted as a function of frequency in (b) for clarity. Different coloured lines show the RMS deviation and thermal noise at different frequencies from $50 - 600\,\mathrm{MHz}$. Simulations were run with the VLSS sky model with an empty area of sky in the main beam. A gain error of 10% was used; equivalent results would be obtained for a phase errors of $\sim 6°$. (c) shows the integration time required to reach the thermal noise limit as a function of frequency, using data extracted from (a).

3.6.5 Effect of integration time

A further simulation can be conducted to test the effect of integration time on the RMS deviations. Figure 3.11 shows how the RMS deviation level changes for extended integration times, of up to 57 min, across the frequency band. The results follow the expected $t^{-\frac{1}{2}}$ progression, where t is the observation duration, across the frequency range. The 50-MHz results suggest a possible deviation from the $t^{-\frac{1}{2}}$ trend, reducing the improvement in noise reduction with time, but longer integration time simulations would be required to establish this.

The dashed lines in Fig. 3.11 extrapolate the RMS deviation to 12 hours of integration according to the $t^{-\frac{1}{2}}$ reduction. Beyond 12 hours, the uv-coverage no longer increases as the Earth rotates, rather it repeats itself, which diminishes the benefit from extended observations. The thermal noise level, shown in solid lines, also decreases by $t^{-\frac{1}{2}}$ as it is simply Gaussian noise, but without the 12-hour restriction. Once the thermal noise decreases below the RMS deviation it limits the usefulness of longer integrations.

By interpolating the impact of gain and phase errors on the RMS deviation in Fig. 3.9, the RMS deviation created by a gain error of 10% is equivalent to a phase error of $\sim 6°$. Hence, from Fig. 3.11 it is found that for gain errors of 10%, or phase errors of $\sim 6°$, the achievable thermal noise limit is between $\sim 0.3 - 10\,\mu$Jy and reaches after $\sim 100 - 100\,000$ hrs, depending on frequency. The shortest achievable integration time is for 100 MHz, with a thermal noise limit of 6 μJy achieved after a 100-hour integration. At 50 MHz the best noise level is higher, at 10 μJy, but reached after a much longer integration time (~ 600 hr) due to the much higher system temperature. Above 100 MHz the best achievable noise level drops rapidly. Reducing the gain and phase errors would further decrease the thermal noise limit.

An important caveat is that 57 mins is a limited time period compared to the multi-hour observations often performed with interferometers. Longer duration simulations were not practical owing to available compute resources, but simulating much greater integration times to verify the accuracy of the extrapolations would be beneficial.

As with all simulations, these have their limitations. The performance of electronic components is not explicitly modelled, rather the noise introduced by components such as the LNAs is assumed to be included in the gain and phase errors introduced on an antenna basis. This neglects the fact that electronics noise is likely to vary with frequency, especially across a wide bandwidth of $50 - 600\,\mathrm{MHz}$.

Other errors in the station beams are also not explicitly modelled, with these simulations assuming that no extra uncertainties are added when the antenna signals are combined to create station-level beams. Quantisation and noise from electronic components will prevent this idealised scenario.

It is also worth noting that for time-dependent errors, such as imprecise antenna locations, the resultant phase errors will increase with frequency. However, it is difficult to know precisely the contribution of time-dependent and time-independent phase errors in the final signal, so modelling this effect is not straightforward. These simulations assume the phase errors specified are the same across the frequency range.

The VLSS sky model used for these simulations is an incomplete model of the sky. The VLSS model is a point-source sky model, hence extended sources, including the galactic plane, are not accounted for within these simulations. It is not clear how these would have affected the results. Furthermore, applying a spectral index of -0.7 to model the change in flux with frequency is a simplification; in reality source fluxes vary in a range of ways with frequency. Finally, the differences in RMS of the beam directions in Fig. 3.9b suggests that very bright sources can produce RMS differences that are an order of magnitude bigger than what we see here. As stated previously, the six brightest sources were removed perfectly from the VLSS catalogue used as a sky model; the RSN could potentially rise by amounts greater than an order of magnitude if one of these sources were located in a strong sidelobe.

3.7 Summary

These simulations suggest that if the signal processing is of a high enough accuracy so that errors are dominated by antenna gain and phase errors, the effect of uncorrelated

errors in the beam pattern of SKA1-Low will make imaging more difficult at the low frequency end of the band compared to the high frequency. This can be attributed to larger field of view and increased flux from sources at low frequencies.

The RMS deviations between ideal and error images decrease with increasing frequency and decreasing magnitude of the errors. These simulations indicate that the thermal noise limit on images will range between $\sim 10 - 0.3\,\mu$Jy and reached after $\sim 100 - 100\,000$ hours, for observations from $50 - 600\,$MHz.

The results also suggest that, in order to mitigate the effect of large beam areas on imaging at the lowest SKA1-Low frequencies, consideration be given to schemes which increase the effective station size at low frequency.

Chapter 4

Station models

4.1 Introduction

SKA1-Low is described in the SKA's *Requirements specification* [96] as comprising 1024 stations of 35-m diameter, each containing 256 antennas. The antennas will be randomly positioned within each station, within the constraints of the antennas' physical dimensions. Half of the stations will be located in a 600-m radius of SKA1-Low's centre and a further quarter between $600 - 1000$ m. The remaining stations will be distributed beyond 1000-m radius, forty-five of which will be in three spiral arms which continue to radial distances of $\sim$ 50 km. The total collecting area of SKA1-Low is necessary to achieve the required sensitivity for key science projects such as mapping the EoR [55].

Simulations have shown that the strong, near-in sidelobes of the lowest frequency primary beams, which cover wide angles on the sky, cause residual noise which can significantly affect the telescope's imaging (see Chapter 3). The strength of the sidelobes can be ameliorated by apodization. This is the process in which the signal detected by each antenna in a station is given a weighting according to its radial position in the station.

This chapter considers a range of telescope layouts for SKA1-Low and assesses their total residual noise levels and their ability to detect the power spectrum of the EoR. Conventional discrete station designs are considered with different apodization

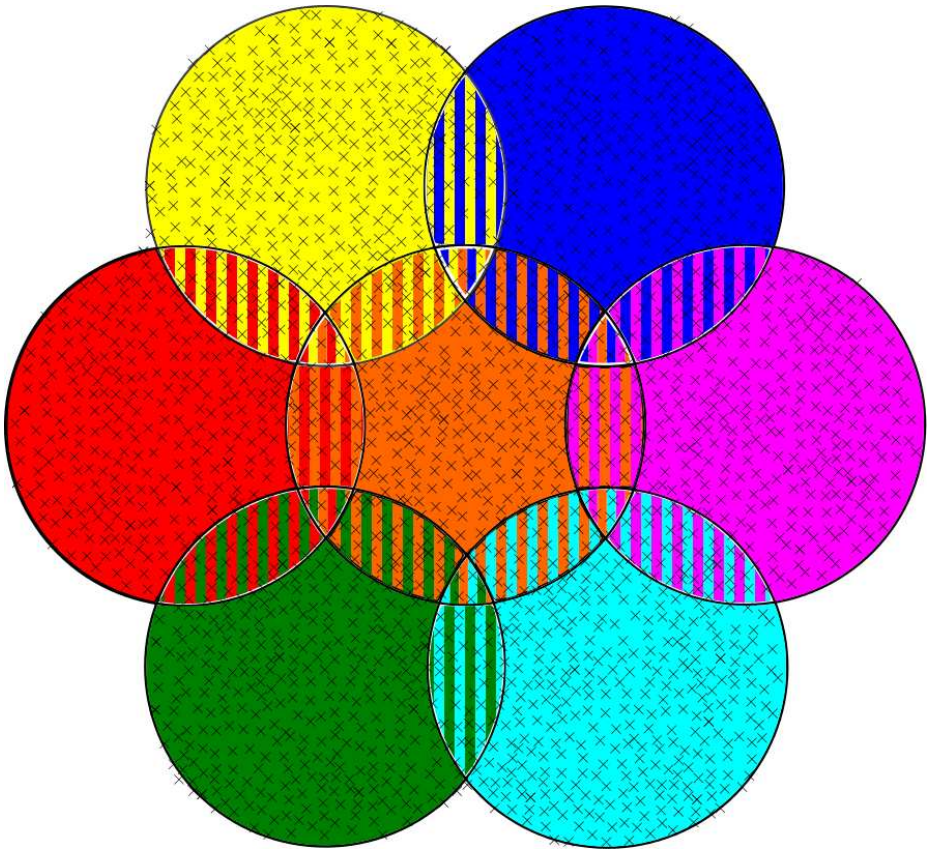

Figure 4.1: An illustration of the flexible station concept suggested by Grainge [89]. Individual stations are shown by coloured circles and antennas represented by crosses. Some antennas are used in multiple stations, shown by the antennas in striped regions where stations overlap. By sharing antennas in this manner, a greater number of antennas can be used per station without an increase in the overall number of antennas. This permits the use of apodization by down-weighting antennas without a loss in sensitivity.

schemes, along with station layouts with shared antennas.

The idea of shared antennas between stations for SKA1-Low was suggested by Grainge [89] and an example is illustrated in Fig. 4.1. Sharing antennas between stations also opens up the possibility of maintaining an approximately constant beam size over SKA1-Low's frequency range, as well as creating elliptical stations to observe sources away from the zenith, such that the stations appear circular to an observer at the source's position; this is illustrated in Fig. 4.2.

The station beam from an elliptical station will vary with time as individual antennas enter and leave the flexible station, as the station's geometrical footprint

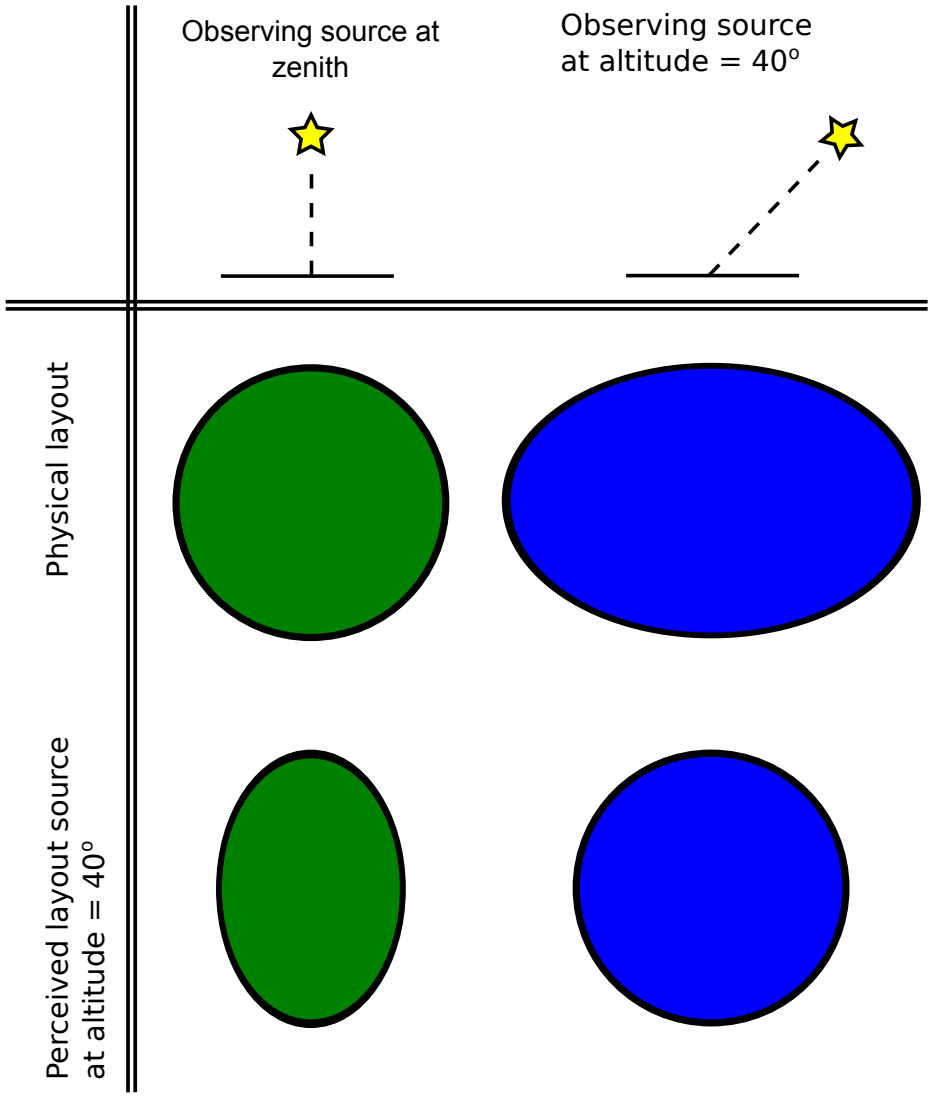

Figure 4.2: An illustration of the geometric layout of stations and the perceived layout of stations from observers in the far-field in different directions. By manipulating the axis sizes of the station, the perceived appearance of a station can remain circular.

changes. However, with apodization, an antenna entering from the edge of the station would have a very low gain weighting. Its weighting would increase as the station boundaries varied to place the antenna in a more central position. This change in weighting will diminish sudden step changes in the station beam patterns as antennas joined or left. If apodization is not applied to a flexible station telescope model this effect may prove to be more problematic. Station illumination functions that are not apodized are likely to have poor imaging properties.

4.2 Theory of station models

An aperture array is composed of a large number of individual antennas. The signals of these antennas are beamformed into stations. The stations are equivalent to the dishes in a more conventional interferometer. The *Requirements specification* [96] envisions 1024 discrete stations, each creating their own beam by combining the signals from their 256 constituent antennas. Each station's beam is then cross-correlated with the beams of the other stations.

It is possible to change the number of stations and their sizes (i.e. number of antennas). Larger stations will typically have a smaller average sidelobe level (see Chapter 2), hence confusion noise levels should decrease. The *uv*-coverage of a telescope containing larger stations will be less uniform, however, as the number of baselines are reduced. Larger stations will also have larger minimum baseline lengths, decreasing sensitivity to the largest on-sky angles. Furthermore, the primary beam's field of view is inversely proportional to station size.

Studies have advocated both large station sizes ($\sim$ 180-m diameter), to help the SKA reach the thermal noise limit [87] and small stations ($\sim$ 35-m diameter), so that the primary beam field of view is large enough to trace the relevant size scales for the EoR key science project [85]. The field of view for any station size could be increased by mosaicing multiple beams together, however there are concerns as to whether the resultant data would be accurate enough for the EoR key science project [85].

Station designs which comprise spatially separated groups of antennas, where each group is one station, will be described as 'discrete' stations. An alternative approach is to place antennas together in large contiguous areas. These areas of antennas provide a versatile space in which to create stations. In this approach, antennas could be present in multiple stations and station shapes and sizes would be flexible, hence they are termed here as 'flexible' stations. The versatility of flexible stations would be limited by the digital signal processing connections between antennas.

4.2.1 Apodization

One advantage of sharing antennas between stations is that stations could be apodized without a decrease in the overall sensitivity. Steep gradients in spatial sensitivity lead to ringing, or sidelobes, in the beam. In an aperture array, this steep gradient occurs at the edge of a station, where an area of constant sensitivity is adjacent to one of zero sensitivity (i.e. the radius at which there are no further antennas). The ringing manifests itself in the sidelobes of a primary beam. Simulations of SKA1-Low observations have shown that strong, wide-angle sidelobes associated with lower frequencies could pose problems for imaging (see Chapter 3). Apodization is the technique of down-weighting the illumination in an aperture as a function of radial distance from the centre, resulting in a smooth transition at the aperture edge and reduced prominence of sidelobes. In an aperture array this is achieved by changing the beam coefficients used in the beam tuning as a function of radial distance. This should help limit the impact of confusion noise.

Down-weighting antennas decreases the sensitivity of a station. This can be counteracted by including more antennas in the station, so that there is an inner circle of antennas of equal radius to an unapodized station which are not down-weighted, and an outer ring of antennas that are down-weighted according to a suitable window function (although this does increase the station size). By overlapping the outer rings between stations, and using the signals from each antenna in the overlap area in both stations, it hence becomes unnecessary to increase the total number of antennas in the aperture array to preserve sensitivity. Autocorrelation of signals will be a problem if antennas are shared between stations, since the autocorrelated noise adds total power to the visibility and would dominate over the cross-correlation noise. It is hence imperative that baselines with shared antennas are not formed.

In this investigation, antennas are down-weighted in one of two methods. One is

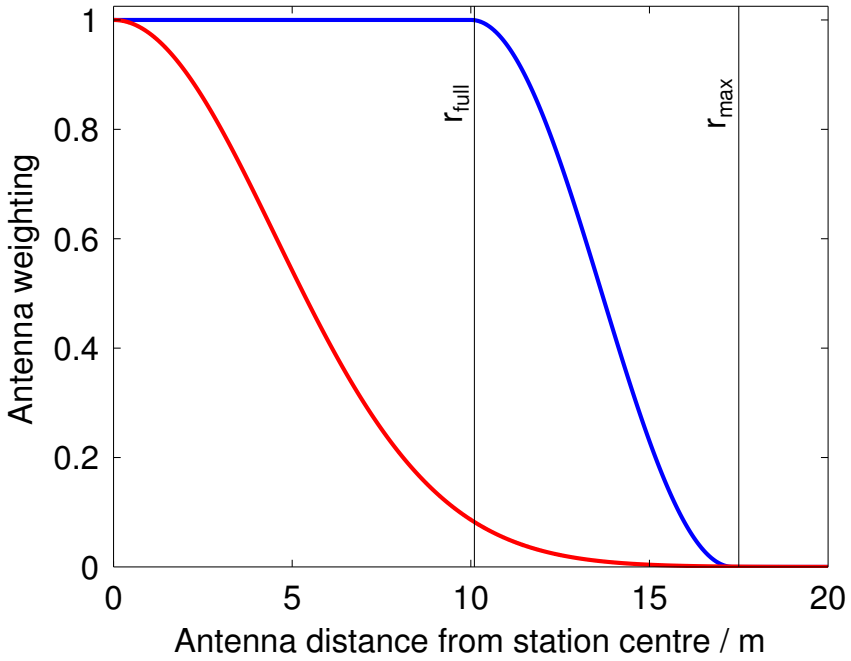

Figure 4.3: Weighting of antennas as a function of radial distance from station centre for the apodization of a 17.5-m radius circular station. The blue line shows the Hanning window described in Eq. (4.1) and the red line shows a Gaussian function. r_{max} is the maximum station radius and $r_{\mathrm{full}} = r_{\mathrm{max}}/\sqrt{3}$ is the radius of the unapodized antennas in the Hanning window.

by a Hanning window function, as suggested by Grainge [89]:

$$w(r) = \begin{cases} 1 & \text{if } r \leq r_{\mathrm{full}}\,, \\ \frac{1}{2}\left(1 + \cos\left(\frac{\pi(r - r_{\mathrm{full}})}{(\sqrt{3}-1)r_{\mathrm{full}}}\right)\right) & \text{if } r_{\mathrm{full}} \leq r \leq r_{\mathrm{max}}\,, \\ 0 & \text{if } r > r_{\mathrm{max}}\,, \end{cases} \tag{4.1}$$

where $w(r)$ is the weighting given to each antenna, r is the radial distance of each antenna from its station centre and r_{max} is the station's maximum radius. r_{full} is the radius of the unapodized antennas and $r_{\mathrm{full}} = r_{\mathrm{max}}/\sqrt{3}$. In the other method, antennas were down-weighted more aggressively by means of a Gaussian function, such that $w(r_{\mathrm{full}}) = 0.09$ and $w(r_{\mathrm{max}}) = 0.0007$. Both these functions are shown in Fig. 4.3.

4.2.2 Elliptical stations

A further advantage of flexible stations is the potential ability to compensate for the reduction in each station's collecting area when observing sources away from the zenith. When observed from directly above, a circular station appears circular, but as an imaginary observer moves from the zenith to the horizon, the circle is projected to be an ellipse of greater and greater eccentricity, until eventually the observer perceives only a line. If, however, the station were in fact geometrically elliptical, this same effect could be used to make station appear circular at nearly any given angle. Flexible stations could be used to change the geometric shape of each station, depending on the direction of an observation.

When tracking a source, individual antennas would be added and removed from the station so that it would most closely approximate the ideal elliptical station to observe a source in a given direction. However, the random placement of antennas would allow the potential for sudden changes in a station's beam pattern every time an individual antenna enters or leaves the station, as the station's geometrical footprint changes.

To determine the geometrical layout of a station required to give the appearance of a circle to an arbitrarily positioned observer on the celestial sphere, the general equation of an ellipse is used

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1\,, \tag{4.2}$$

where a and b are the semi-major and semi-minor axes of the ellipse, respectively, and x and y define the position of the station's boundary. a is a taken to be a constant equal to $r_{\max}$ and b is defined by

$$b = \frac{a}{\sin(\theta)}\,, \tag{4.3}$$

where θ is the altitude of the observation. The station is then rotated such that the two ends of its semi-minor axis are equidistant from the source. The rotated position

of each antenna is defined by

$$x' = x\cos(\phi) - y\sin(\phi)\,, \tag{4.4}$$

and

$$y' = x\sin(\phi) + y\sin(\phi)\,. \tag{4.5}$$

where ϕ is the azimuth of the observation.

The weighting to apply to antennas is not simply a function of radius in elliptical stations. The position of each antenna must be normalised by scaling down the semi-minor axis direction, to where it would be positioned in an equivalent circular station to determine its weighting with Eq. (4.1). For a station with $\phi = 0$, this is given by

$$y_{\text{norm}} = y\left(\frac{a}{b}\right)\,. \tag{4.6}$$

Flexible stations also allow the option of maintaining a constant field of view across wide frequency ranges. The primary beam field of view is dependent upon effective station diameter, d_{stat}, and observation wavelength, λ, and is approximately calculated by Eq. (2.1), which is restated here:

$$\text{FoV}_{\text{beam}} = \frac{\pi}{4}\left(\frac{1.3\lambda}{d_{\text{stat}}}\right)^2\left(\frac{180}{\pi}\right)^2\,. \tag{4.7}$$

Hence the field of view remains constant if d_{stat} is scaled in proportion to λ. The effective station diameter is dependent upon the apodization; the Gaussian weighting function produces a smaller effective diameter station than the Hanning function in Eq. (4.1). To create larger effective stations, steeper apodization functions could be applied to stations of a fixed diameter. Alternatively, stations with a greater number of antennas could be used.

A significant advantage of flexible stations is that they could allow the confusion noise and residual noise levels of the telescope to attain lower minimum values than discrete stations. The confusion noise level of a telescope tends to diminish as approximately $t^{-\frac{1}{2}}$, where t is time (the simulations in Chapter 3 predict that the sidelobe

residual noise will decrease as $t^{-\frac{1}{2}}$; other (unpublished) simulations on confusion noise suggest that confusion noise will decrease at less than $t^{-\frac{1}{2}}$ for the SKA [108]). The rotation of the Earth continuously changes the directions of the telescope's baselines and hence 'fills in' larger areas of the *uv*-plane with extended observations, compared to a single snapshot observation. After a maximum of 12 hours (and hence the Earth rotating 180°) the baselines no longer occupy unique positions on the *uv*-plane, instead they repeat themselves. The synthesised beam pattern (and its confusion noise causing sidelobes) is no longer a unique shape at future instants in time. Consequently confusion noise remains constant beyond 12 hours.

A flexible station, on the other hand, could potentially be reconfigured after 12 hours, moving the station centres. This could create new baselines, allowing the confusion noise to decrease further. This would be limited by the flexibility of the telescope's signal processing to create stations in different positions. The density of stations in the SKA1-Low core suggests that it will be unlikely that any new set of baselines created would be unique; there would almost inevitably be a repetition of some baselines. This would diminish the benefits of extending observations.

4.3 Deconvolution

In addition to producing a telescope configuration that is optimised for detecting the EoR power spectrum signal, it is strongly desirable that the synthesised beam of the telescope has low sidelobes and it is useful for the beam to be circularly symmetric; a Gaussian beam is generally considered optimum. To accurately determine the sky's intensity distribution the dirty image must be deconvolved using the synthesised beam to give a good approximation to the sky image. The dirty image is the image created by Fourier transforming; it is the true sky convolved by the synthesised beam of the telescope. A beam with low sidelobes will result in a dirty image that is less distorted from the true sky, allowing a more accurate sky intensity distribution to be recovered

[109, 110]. This is useful for all observations with SKA1-Low, including those of the EoR.

The response function of an interferometer is not sampled everywhere in the uv-plane, but only in the places corresponding to the baselines over the length of the observation. An interferometer yields a dirty image, I_D, according to

$$I_D(x,y) = \iint I(u,v)\, S(u,v) \mathrm{e}^{2\pi i(ux+vy)}\, \mathrm{d}u\, \mathrm{d}v\,, \tag{4.8}$$

where $I(u,v)$ is the full uv-response function and $S(u,v)$ is the sampling function, given by

$$S(u,v) = \begin{cases} 1\,, & \text{if } (u,v) \text{ sampled by the interferometer.} \\ 0, & \text{elsewhere.} \end{cases}$$

An ideal telescope would sample (u,v) at all points, leaving Eq. (4.8) as

$$I_D(x,y) = \iint I(u,v)\, \mathrm{e}^{2\pi i(ux+vy)}\, \mathrm{d}u\, \mathrm{d}v\,, \tag{4.9}$$

Applying the convolution theorem to Eqs. (4.8) and (4.9) leads to

$$I_D(x,y) = I(x,y) * B(x,y)\,, \tag{4.10}$$

where $*$ is the convolution symbol and $B(x,y)$ is the synthesised beam and Fourier transform of $S(u,v)$, given by [111],

$$B(x,y) = \iint S(u,v)\, \mathrm{e}^{2\pi i(ux+vy)}\, \mathrm{d}u\, \mathrm{d}v\,. \tag{4.11}$$

Equation (4.10) has an infinite number of possible solutions, as the unsampled visibilities in the uv-plane could theoretically take any value [112], however a range of techniques have been developed to deconvolve Eq. (4.10) to provide plausible solutions.

4.3.1 CLEAN

Perhaps the most popular deconvolution algorithm is CLEAN, first published by Högbom in 1974 [113]. The CLEAN algorithm takes the dirty image and subtracts

a fractional component of the synthesised beam away from a position centred on the dirty image's brightest pixel. The brightest pixel's position is recorded and the process repeated until all structure is removed from the dirty image and only noise remains (this point can be determined by measuring the RMS level of the image and noting when it fails to decrease from one iteration to the next). At every position the dirty beam was subtracted, a restoring beam (or 'CLEAN' beam) is placed instead. The restoring beam would typically be a Gaussian of HPBW equal to that of the dirty beam [38].

CLEAN's enduring popularity is partially attributable to its relative simplicity, but this simplicity also results in drawbacks. CLEAN algorithms generally work on the assumption that every radio source is a point source. Consequently CLEAN may not accurately recover $I(x, y)$ in images containing extended sources, especially sources with low brightnesses relative to the background. Variations on Högbom's original algorithm attempt to improve CLEAN, such as the Clark algorithm [114] and the Cotton-Schwab algorithm [115].

4.3.2 Maximum Entropy

An alternative approach is that of the Maximum Entropy Method, a statistical method used in a variety of fields which was introduced to radio astronomy at about the same time as CLEAN [116, 117, 118]. The Maximum Entropy Method attempts to take all possible solutions to Eq. (4.10) and narrow it down to the distribution with the minimum information content, as measured by the entropy, which is consistent with the visibility data. The Maximum Entropy Method does not assume point sources like CLEAN and is superior at imaging extended sources. Furthermore Maximum Entropy Method images are smoother than images made from the same visibilities with CLEAN [119]. Conversely, the Maximum Entropy Method performs less favourably than CLEAN when imaging point sources. The resolution varies across Maximum Entropy Method images depending on the signal-to-noise level [38].

4.3.3 Ideal beam shape

Regardless of the specific deconvolution algorithm used, a circularly symmetric beam with low sidelobes is strongly preferred in radio astronomy as it will help facilitate accurate deconvolution. The response of sidelobes will vary with, for instance, gain and phase errors in the signal processing path. Minimising sidelobes will minimise the extent to which they can impair imaging.

4.4 Generation of telescope models

4.4.1 Overview of telescope models

A range of telescope models were created to evaluate the suitability of different types of telescope models for SKA1-Low. Three classes of telescope models were created, which will be termed discrete, 'starfish' and 'cluster'. The starfish and cluster models are flexible telescope models.

Discrete telescope models are telescope layouts in the conventional aperture array sense, comprising antennas grouped into geographically localised stations in which each antenna is placed in a single station. This is the type of telescope model currently planned for SKA1-Low, as detailed in the *Requirements specification* [96].

In contrast, flexible telescope models allow antennas to be shared between multiple stations. The geometries of the stations can also be recast between observations, within the confines allowed by the signal processing architecture, to allow the geometries to be tailored to the observation at hand. Flexible stations allow stations to be apodized whilst retaining the same number of antennas and approximate sensitivity in the telescope model.

The starfish models are based upon the geometry proposed by Grainge [89]; a three-legged starfish-shaped telescope model, similar to that shown in Fig. 4.4. When creating a flexible telescope model there are two competing demands on the telescope model. The first is the desire for a large contiguous area of antennas in which to have a versatile space to create stations, which is optimised by concentrating the antennas

in a compact core. The second desire to have stations scattered over a wide range of distances to produce a *uv*-plane coverage which is sensitive to a wide range of angular scales; this is why the *Requirements Specification* states 50% of antennas should be within a 600-m radius of the telescope centre and 75% within 1000 m. A Y-shaped array compromises between these demands.

The Y-shape also allows circles to be filled in the *uv*-plane in only 4 hours, one of the reasons it has been used for existing interferometers such as the VLA and GMRT. However, the Y-shape of the VLA is known to cause problems with its synthesised beam, creating strong sidelobes and reducing the accuracy of deconvolution [120]. Curving the arms of the Y allows degeneracies in the baselines to be broken, enhancing the *uv*-coverage and hence improving the synthesised beam.

The starfish models considered in this work consist of a contiguous block of antennas with the legs tapered so as to conform to the specification of 50% of antennas within a 600-m radius of the station centre and 75% within 1000 m. The remaining 25% of antennas were placed in thirty clusters of 2185 antennas along three spiral arms diverging from the starfish. Clusters were concentrated towards the array centre by radially placing them with logarithmic spacing [96].

For circular stations, as used in these investigations, the target value was 768 antennas per station; this is the number of antennas in a *Requirements specification* station, 256, multiplied by the factor of increased area of the flexible stations, 3, which allows for apodization without down-weighting the sensitivity of a station below that of a 256-antenna station. In the outlying clusters a minimum distance between station centres of 12 m was imposed to increase utilisation of different antennas.

The second variation of flexible telescope models, the cluster models, uses the same basic premise as the starfish models. However, rather than using a starfish-shaped contiguous block of antennas as the framework on which to build stations, a number of large circular clusters of antennas are used instead. As with the starfish designs, the density and radial distribution of antennas conform to the *Requirements specification.*

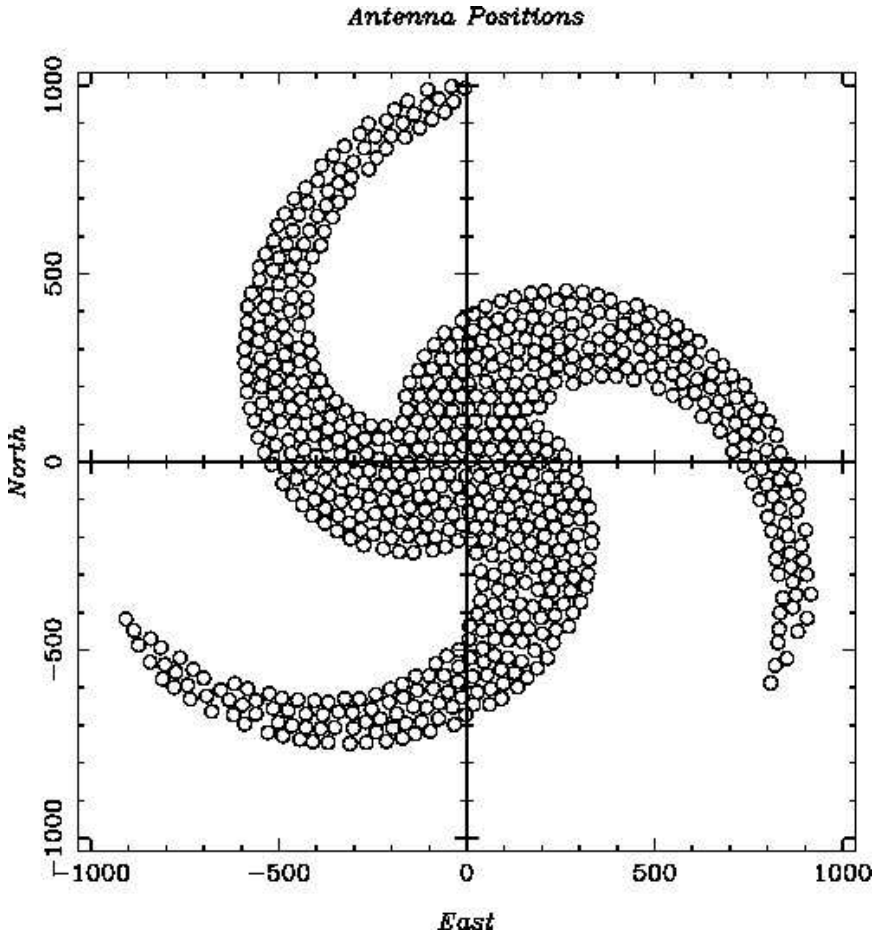

Figure 4.4: Example of starfish-shaped telescope model proposed by Grainge. Image from [89].

Each cluster contained 8, 16, 32 or 64 station's worth of antennas, depending on the telescope model. A variation on this approach is to place a large cluster at the centre of the telescope layout, surrounded by smaller clusters. Positioning the stations is done in the same manner as with the outlying clusters in the starfish telescope models, except that the minimum distance between station centres varies between $12 - 17\,\mathrm{m}$, depending on cluster sizes; larger clusters allow for greater separation.

4.4.2 Sub-stations

To prevent the autocorrelation of signals, the visibility data of baselines between stations with shared antennas were discarded. This meant, however, that there would be no baselines of less than $2r_{\mathrm{stat}}$, where r_{stat} is the station radius (which is a factor of $\sqrt{3}$ bigger than the $17.5\,\mathrm{m}$ specified in the *Requirements specification*). If the station layout is approximated to be hexagonally tiled, the consequence was that six baselines were lost per station, one baseline for each of a station's nearest neighbours, as illustrated in Fig. 4.5.

However, by splitting each station into four smaller sub-stations, which are correlated with only the other sub-stations within their station, six additional short baselines can be created per station. This is illustrated in Fig. 4.6. The radius of the sub-stations was set to utilise the maximum number of antennas whilst maintaining the same shape as the stations themselves. For circular stations, the substation radius, r_{substat}, is given by

$$r_{\mathrm{substat}} = \frac{r_{\mathrm{stat}}}{1 + \sqrt{2}}. \tag{4.12}$$

In each station, the four sub-stations were rotated about the station centre by a unique angle, to give the largest range of instantaneous non-redundant baselines across the telescope. With the maximum radius of sub-stations as given in Eq. (4.12), $\sim 69\%$ of antennas per station are used in both a sub-station and the station. The baseline

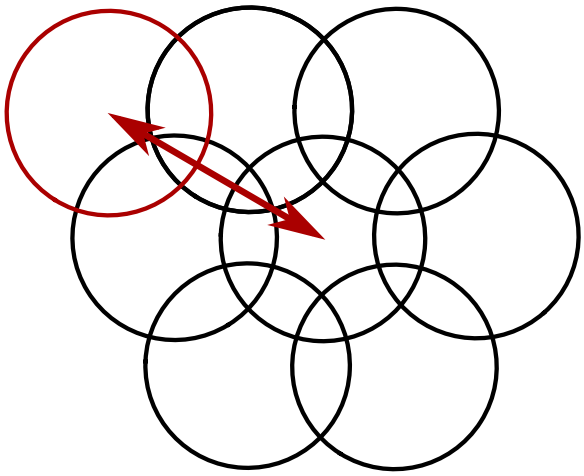

Figure 4.5: Illustration of flexible stations, showing stations which share antennas are not correlated. Each circle represents a station. If the layout of stations is considered to approximate hexagonal packing, each station is not correlated with its six nearest neighbours, but does correlate with every other station. The arrow shows the only permissible baseline of the central black station in this telescope model; only the red station does not share any antennas with it.

lengths between two sub-stations within a station is either $0.83\,r_{\mathrm{stat}}$ or $1.17\,r_{\mathrm{stat}}$, depending on whether it is a baseline with an adjacent or diagonal substation. This provides more data regarding angular scales from baselines below $2\,r_{\mathrm{stat}}$, but does not provide baselines approximately continuously down to $\sim 0.57\,r_{\mathrm{stat}}$ as would be the case without the increased radius of stations for apodization.

4.4.3 Antenna and station positions

In this work, telescope models of all types were created with stations positioned at randomised polar coordinates to mitigate the powers of individual sidelobes. In the discrete models the radial position of each station was selected from a Gaussian distribution. To allow the radial distribution of stations to conform to the *Requirements specification* the peak of the Gaussian distribution was at a hand-picked radius . The angle coordinate for the stations was randomly generated from a uniform distribution.

Each station was provisionally placed at the generated trial coordinates. If any other station centres were within twice the station radius (i.e. the stations overlap) it was discarded and new coordinates generated. The process was repeated until the

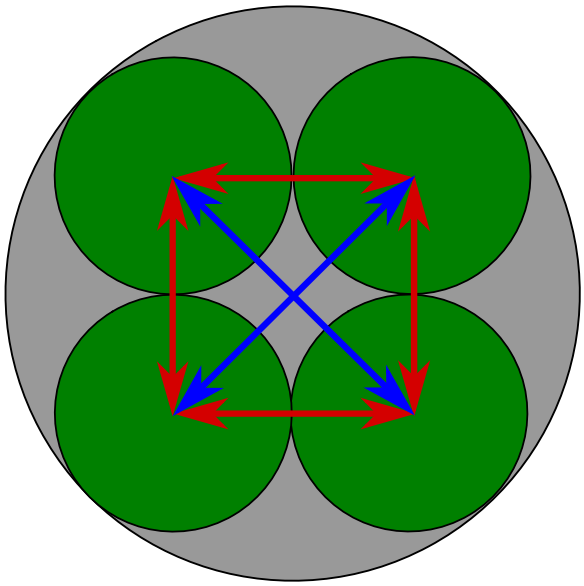

Figure 4.6: Illustration of sub-stations within a station. Sub-stations are shown by the smaller, green circles, with they larger, grey circle showing the layout of the station. The red arrows show baselines between adjacent substations and blue arrows show baselines between diagonal substations. The radius of each sub-station is $r_{\rm stat}/(1+\sqrt{2})$.

requisite number of stations was positioned.

Some station positions were changed by hand to ensure the station distribution conform to the 50% stations within 600 m, 75% within 1000 m requirement.

The outer stations of the SKA, according to the *Requirements specification*, are to be distributed along three spiral arms, which start at a radius $\gtrsim 1000$ m and finish at ~ 50 km. To generate station positions along the outer spiral arms, a logarithmically separated list of angles (going beyond 2π radians) was produced with the list length equal to the number of stations in each spiral arm. The radial position of each station in a single spiral arm was determined by the equation of a logarithmic spiral,

$$r = a\mathrm{e}^{b\theta}\,, \tag{4.13}$$

where r and θ are the polar coordinates and a and b are constants. The positions of each station were rotated about the telescope centre to produce the required number of spiral arms.

The same positioning process was used to place antennas within each station as generating the central station positions, except the antenna radial distributions were uniformly radially weighted rather than having a Gaussian distribution. This resulted in a greater density of trial coordinates nearer the centre of each station, as small differences in angle at large radii result in larger physical separations than at small radii. By discarding overlapping antennas a quasi-uniform distribution of antennas across the station resulted, due to the high density of antennas in a station. A typical station layout is illustrated in Fig. 3.1b.

A different approach was taken to generate the flexible station models to the discrete stations. The initial step was to generate a quasi-uniformly distributed square area of antennas. Antenna separations were allowed to randomly vary over a predefined range of $1.19 - 3\,\mathrm{m}$, but the mean spacing corresponded to $1.93\,\mathrm{m}$. Approximately 100 000 antenna locations were generated in this way, in an area of $\sim 250 \times 250\,\mathrm{m}^2$. This template square of antennas was then replicated to cover an area of $\sim 6000 \times 6000\,\mathrm{m}^2$.

The next step was to crop the antenna locations list so that only antennas within predefined areas remained. For the cluster models, this meant all antennas within the circular cluster areas. These areas were randomly generated according to the desired distribution of clusters. To get antenna locations for clusters outside of the core, a circular area of equal size to each cluster was selected from a random location in the antenna layout and the distribution of those antennas was copied to the cluster centre.

For the starfish models, antennas located within the core's starfish border remained.

The starfish models were structured around three arcs, spiralling out from the telescope centre and separated from one another by $2\pi/3$ rads. The starfish border was defined by

$$R = A - Br\,, \tag{4.14}$$

where r is the radial distance from the telescope centre to a point along an arc, R is the distance from this point along the arc to the starfish border (as measured perpendicularly to the arc) and A and B are constants. In addition, a sharp cut-off was applied to the end of the arcs, to accelerate the tapering of the tips of the arms, as antennas located in a long, narrow taper would be rarely utilised in the flexibly placed stations.

Through these techniques the geometric locations of 262 144 antennas were placed for each of a variety of cluster and starfish layouts. The size, number and radial distribution of clusters were varied between models, as were the thickness and lengths of the starfish arms.

The final step was to assign antennas to stations. Whilst the geometric locations of the antennas are fixed, their grouping into stations is not if a sufficiently flexible signal processing backend is implemented. Station sizes and shapes were dependent upon the number of antennas per station and whether ellipsoidal stations were used (only circular stations were created here). Trial station centres were guided by the approximate antenna locations; if a minimum number of antennas were located in the station and no other stations were within a defined minimum separation distance, the station and its antennas were added to the telescope model. This was repeated until the desired number of stations was created. The weightings of the antennas within each station were calculated according to the apodization scheme utilised.

4.4.4 Telescope models created

4.4.4.1 Discrete models

Three discrete telescope models were created: one with 50 stations of 5000 antennas each ('50×5000'), one with 200 stations of 1250 antennas each ('200×1250') and one with 1024 stations of 256 antennas each ('1024×256'). Station radial distributions

conform to the *Requirements specification* of 50% within a radius of 600 m of the telescope centre, 75% within 1 km and three spiral arms out to a radius of 50 km. The 1024×256 model conforms exactly to the *Requirements specification*, 50×5000 presents large stations akin to those advocated by Braun [87] and 200×1250 is a half-way point between the two and similar to the efficient signal processing approach suggested in Price et al. [86]. These telescope models are illustrated in Figs. 4.7a, 4.7b, 4.8a, 4.8b, 4.9a and 4.9b.

4.4.4.2 Starfish models

A large number of flexible telescope models were also created, to allow the exploration of a wide parameter space. Models created in the starfish design had a range of arm thicknesses, arm taper rates, core radii and arm curvature, as illustrated in Figs. 4.10a, 4.10b, 4.11a, 4.11b, 4.12a and 4.12b. Three different specifications are used for the arm curvature: 'no curvature', 'moderate curvature' and 'tight curvature', respectively.

Ideally each antenna would be used equally, however the geometry of these solutions prevents this. The starfish telescope models do not utilise all available antennas equally in any one configuration. Antennas at the edges of the starfish arms are typically used in only one station where they were heavily down-weighted, whereas those in the centres of clusters and arms may be used in several stations. Although this is not ideal, it is necessary, to an extent, to allow flexible stations to be formed. This gives a weighting to the signals detected at certain antennas, along with their inherent errors. Apodizing antennas in discrete models under-utilises antennas to a greater extent than with the flexible models, however.

Attempts were made to increase the utilisation rate of antennas by restricting the placement of stations. This was implemented by placing the stations on a grid within the boundaries of the antenna-covered starfish and then perturbing the positions of the stations. Unfortunately the underlying gridded structure of the station placement

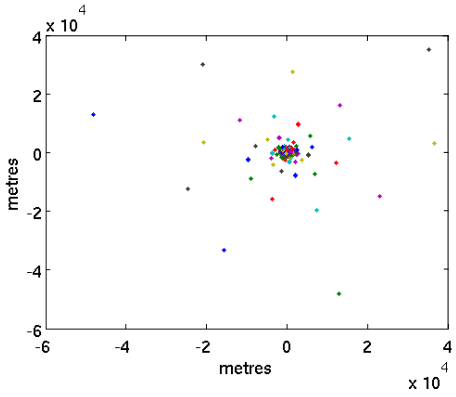

(a) Positions of all stations

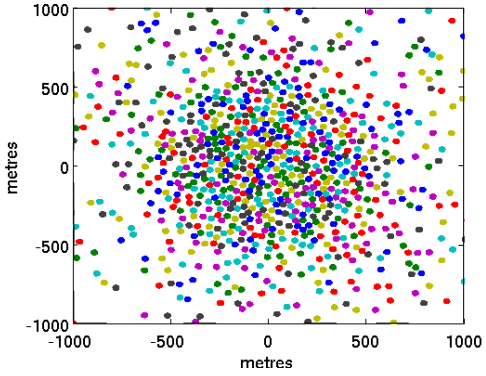

(b) Positions of central stations only

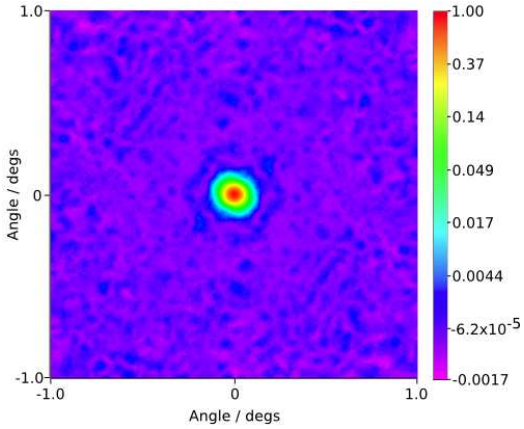

(c) Synthesised beam at 110 MHz

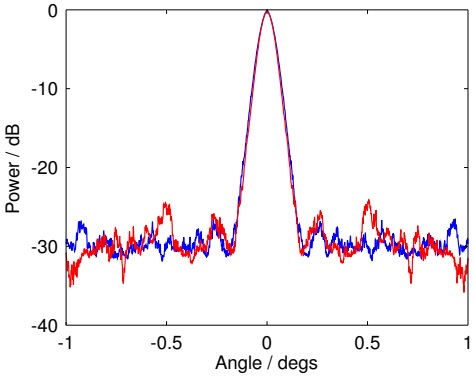

(d) Orthogonal one-dimensional cross-sections of synthesised beam at 110 MHz

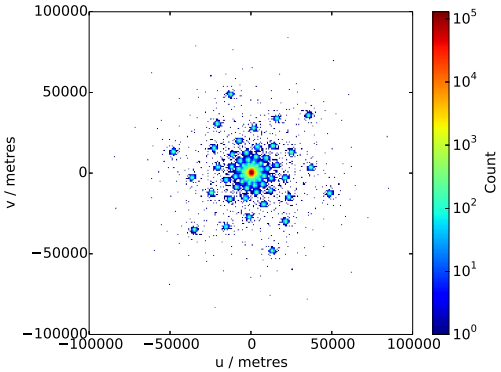

(e) Density of *uv*-coverage. *uv*-coverage is binned into spatial areas; the count of the total number of baselines in each bin is displayed.

Figure 4.7: Discrete telescope model with 1024 stations, each with 256 antennas.

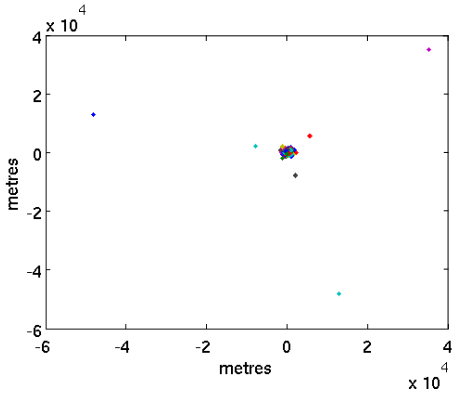

(a) Positions of all stations

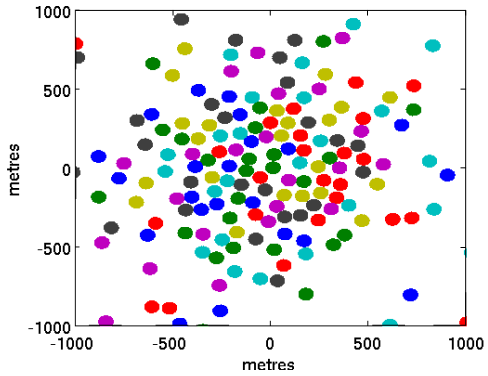

(b) Positions of central stations only

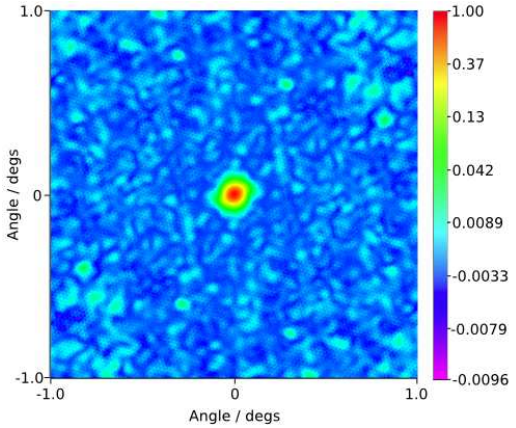

(c) Synthesised beam at 110 MHz

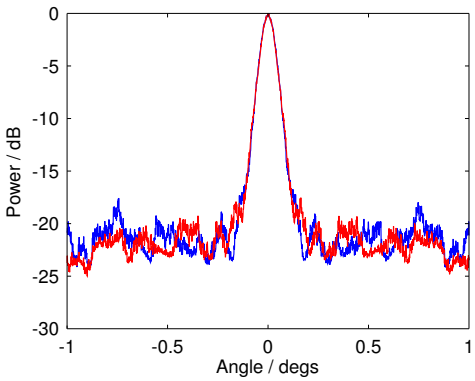

(d) Orthogonal one-dimensional cross-sections of synthesised beam at 110 MHz

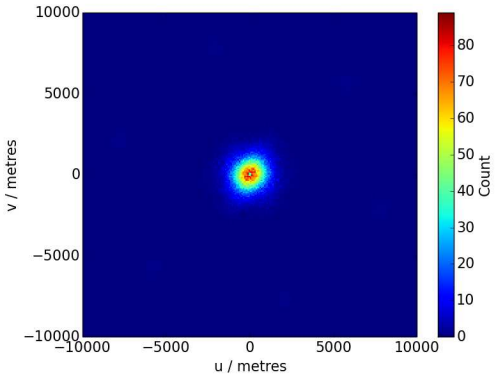

(e) Density of *uv*-coverage. *uv*-coverage is binned into spatial areas; the count of the total number of baselines in each bin is displayed. *uv*-coverage from the longest baselines are not shown.

Figure 4.8: Discrete telescope model with 200 stations, each with 1250 antennas.

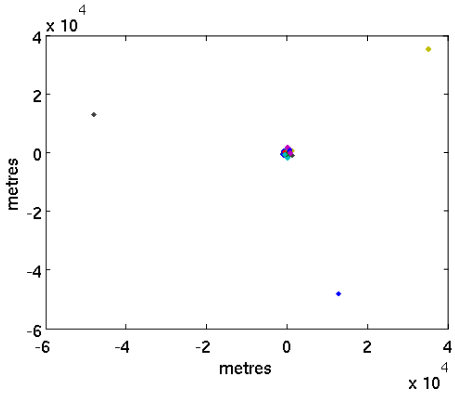

(a) Positions of all stations

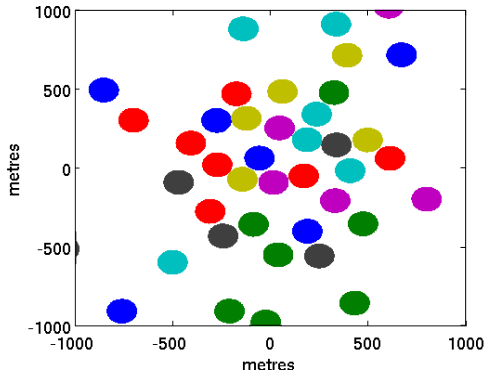

(b) Positions of central stations only

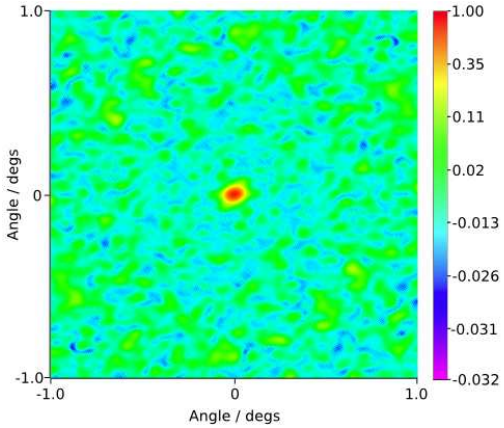

(c) Synthesised beam at 110 MHz

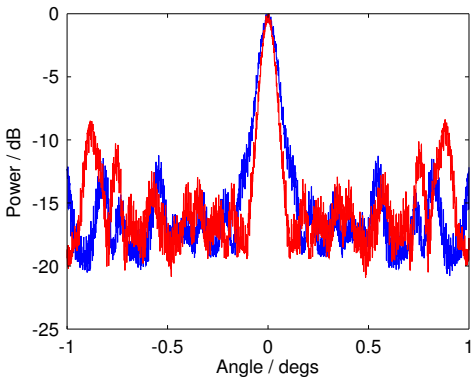

(d) Orthogonal one-dimensional cross-sections of synthesised beam at 110 MHz

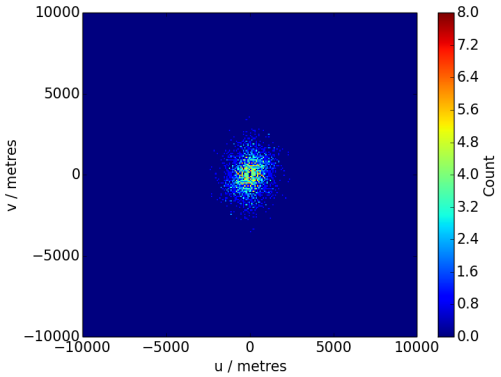

(e) Density of *uv* coverage. *uv* coverage is binned into spatial areas; the count of the total number of baselines in each bin is displayed. *uv* coverage from the longest baselines are not shown.

Figure 4.9: Discrete telescope model with 50 stations, each with 5000 antennas.

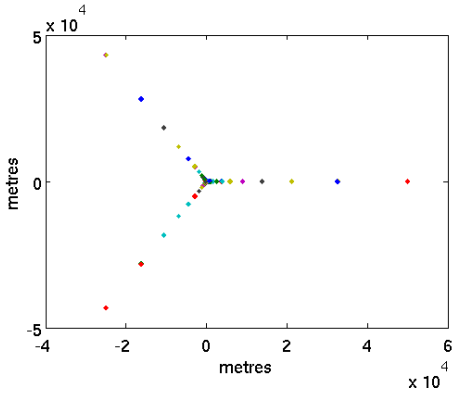

(a) Positions of all stations. Stations are set-up for observing at the zenith (i.e. they are circular).

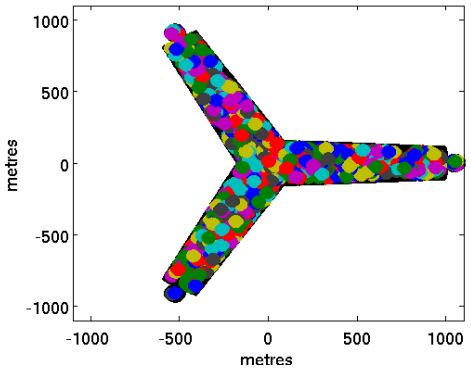

(b) Positions of central stations only. The non-circular black areas at the edge of the central starfish area show positions of antennas that are not utilised (n.b. the grey circular areas are stations)

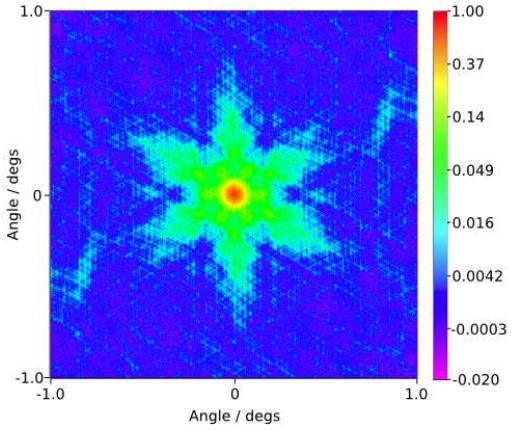

(c) Synthesised beam at 110 MHz

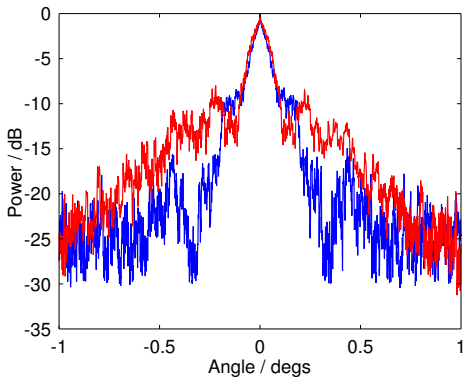

(d) Orthogonal one-dimensional cross-sections of synthesised beam at 110 MHz

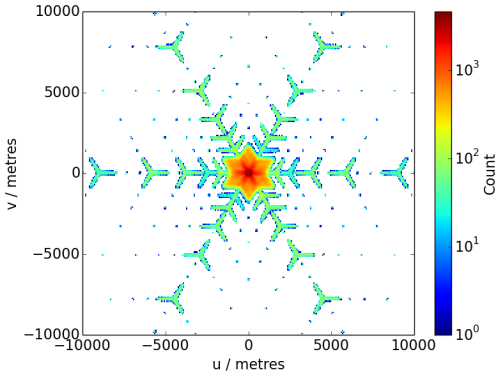

(e) Density of *uv*-coverage. *uv*-coverage is binned into spatial areas; the count of the total number of baselines in each bin is displayed. *uv*-coverage from the longest baselines are not shown.

Figure 4.10: Flexible starfish telescope model with no arm curvature.

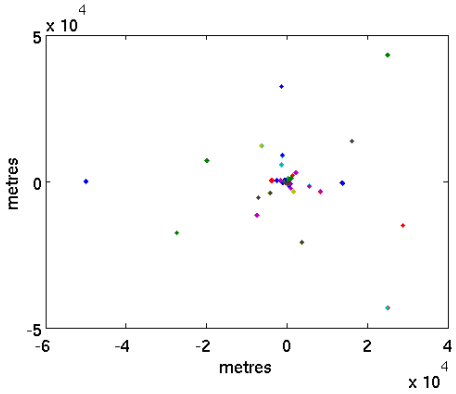

(a) Positions of all stations. Stations are set-up for observing at the zenith (i.e. they are circular).

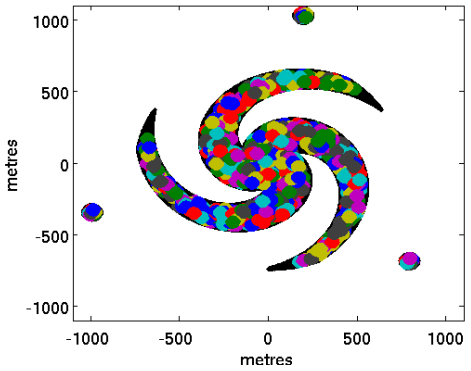

(b) Positions of central stations only. The non-circular black areas at the edge of the central starfish area show positions of antennas that are not utilised (n.b. the grey circular areas are stations).

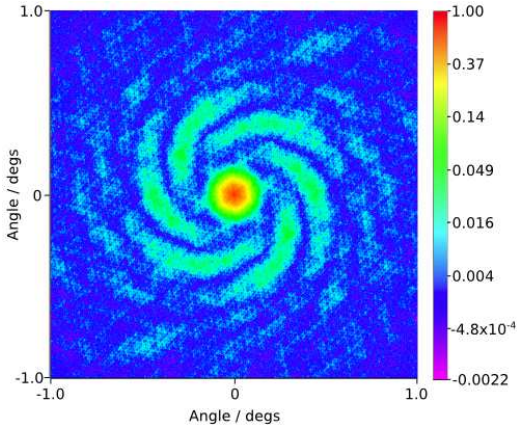

(c) Synthesised beam at 110 MHz

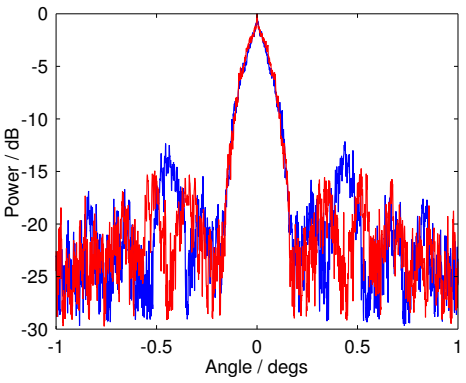

(d) Orthogonal one-dimensional cross-sections of synthesised beam at 110 MHz

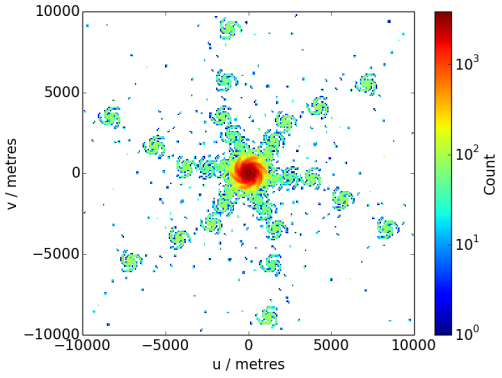

(e) Density of *uv*-coverage. *uv*-coverage is binned into spatial areas; the count of the total number of baselines in each bin is displayed. *uv*-coverage from the longest baselines are not shown.

Figure 4.11: Flexible starfish telescope model with moderate arm curvature.

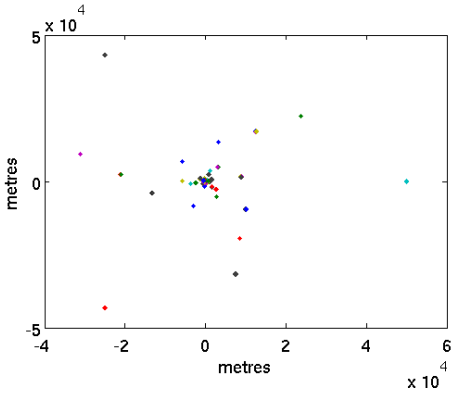

(a) Positions of all stations. Stations are set-up for observing at the zenith (i.e. they are circular).

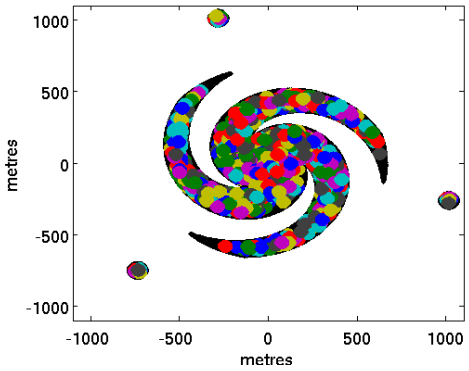

(b) Positions of central stations only. The non-circular black areas at the edge of the central starfish area show positions of antennas that are not utilised (n.b. the grey circular areas are stations).

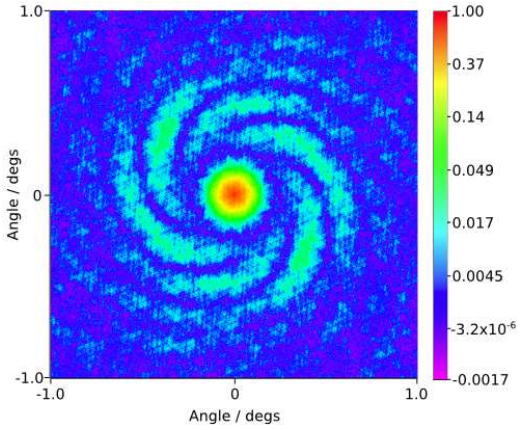

(c) Synthesised beam at 110 MHz

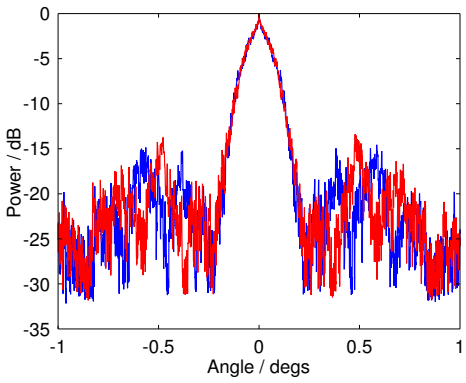

(d) Orthogonal one-dimensional cross-sections of synthesised beam at 110 MHz

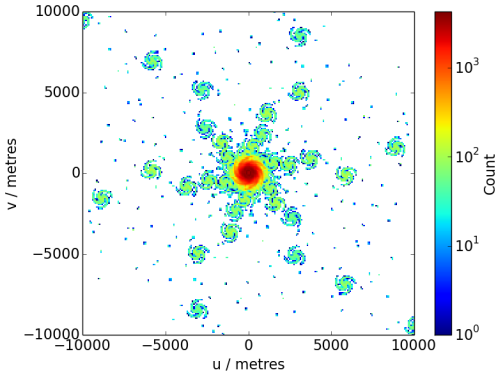

(e) Density of *uv*-coverage. *uv*-coverage is binned into spatial areas; the count of the total number of baselines in each bin is displayed. *uv*-coverage from the longest baselines are not shown.

Figure 4.12: Flexible starfish telescope model with severe arm curvature.

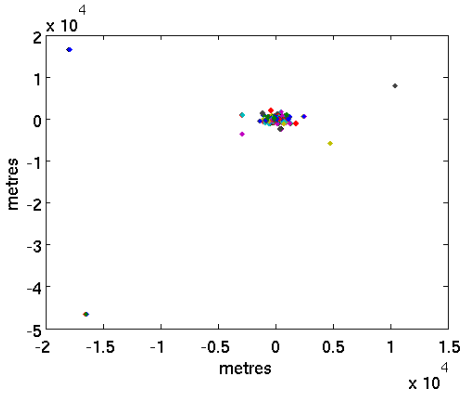

(a) Positions of all stations. There were 448 in central supercluster and 567 surrounding the supercluster, with 8 stations per cluster. Of the surrounding clusters, 9 were in one spiral arm.

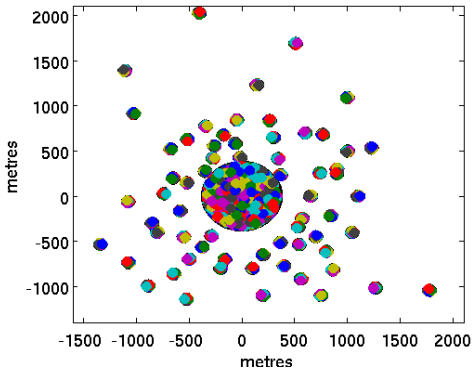

(b) Positions of central stations only. The central supercluster is visible as the large cluster in the centre. Stations are set-up for observing at the zenith (i.e. they are circular).

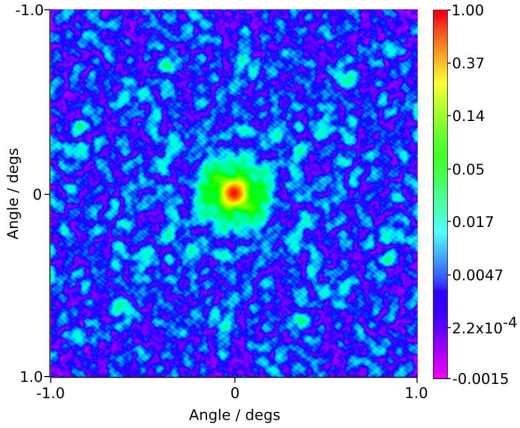

(c) Synthesised beam at 110 MHz

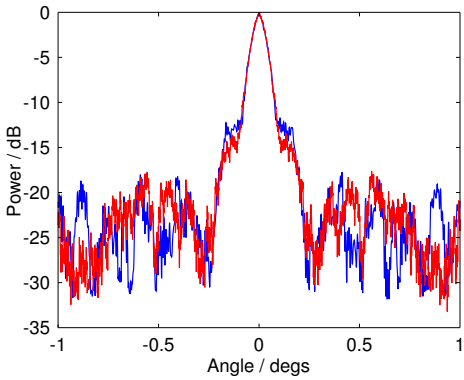

(d) Orthogonal one-dimensional cross-sections of synthesised beam at 110 MHz

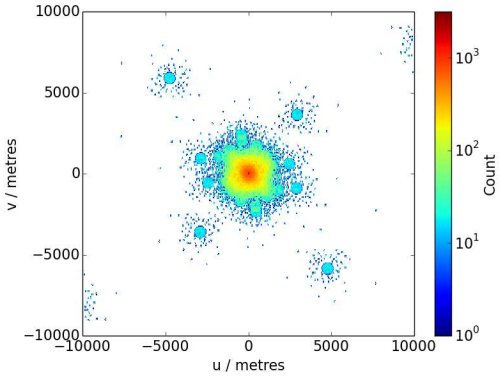

(e) Density of *uv*-coverage. *uv*-coverage is binned into spatial areas; the count of the total number of baselines in each bin is displayed. *uv*-coverage from the longest baselines are not shown.

Figure 4.13: Flexible cluster telescope model A.

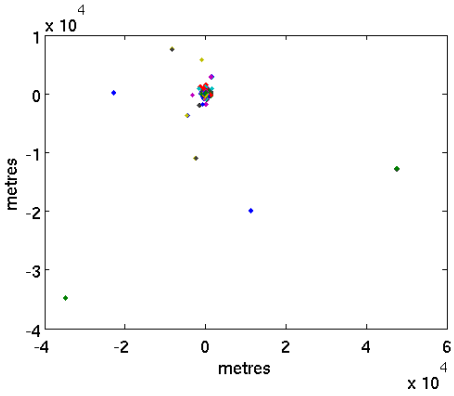

(a) Positions of all stations. There were 448 in supercluster and 562 surrounding the supercluster, with 8 stations per cluster. Of the surrounding clusters, 14 were in two spiral arms.

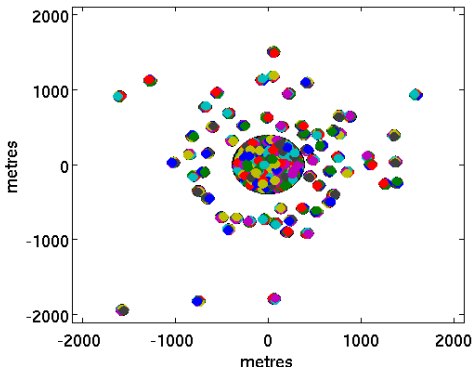

(b) Positions of central stations only. The central supercluster is visible as the large cluster in the centre. Stations are set-up for observing at the zenith (i.e. they are circular).

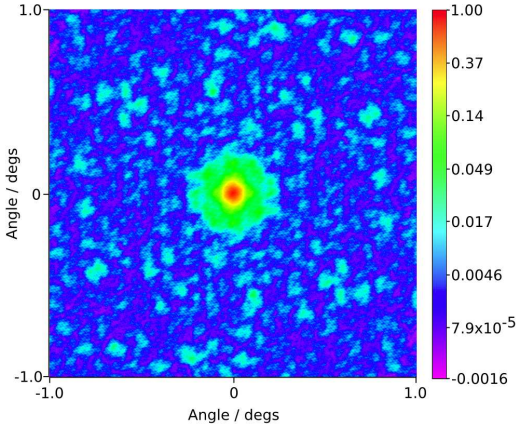

(c) Synthesised beam at 110 MHz

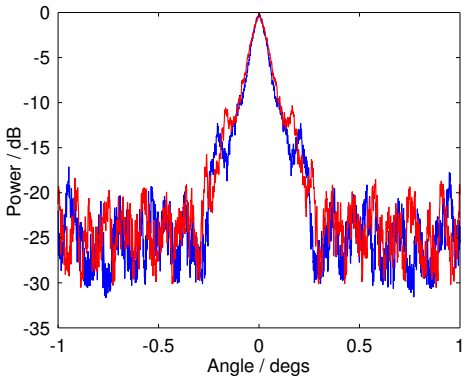

(d) Orthogonal one-dimensional cross-sections of synthesised beam at 110 MHz

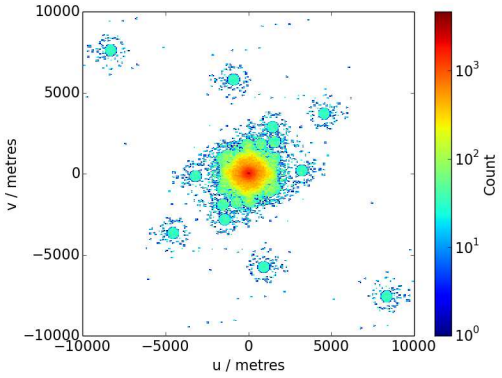

(e) Density of *uv*-coverage. *uv*-coverage is binned into spatial areas; the count of the total number of baselines in each bin is displayed. *uv*-coverage from the longest baselines are not shown.

Figure 4.14: Flexible cluster telescope model B.

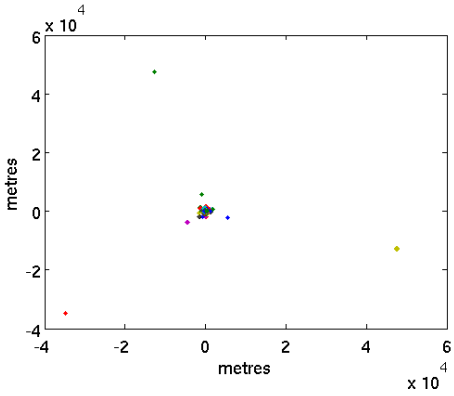

(a) Positions of all stations. There were 448 in supercluster and 567 surrounding the supercluster, with 8 stations per cluster. Of the surrounding clusters, 9 were in three spiral arms.

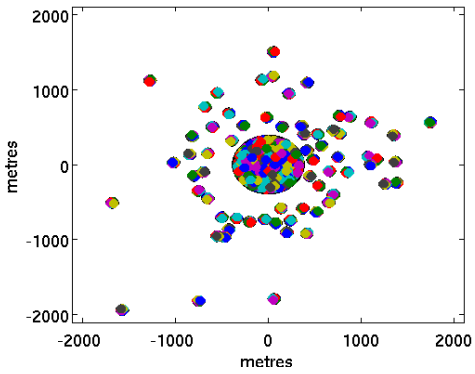

(b) Positions of central stations only. The central supercluster is visible as the large cluster in the centre. Stations are set-up for observing at the zenith (i.e. they are circular).

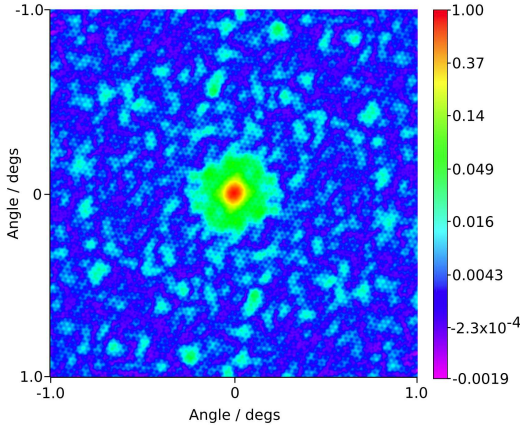

(c) Synthesised beam at 110 MHz

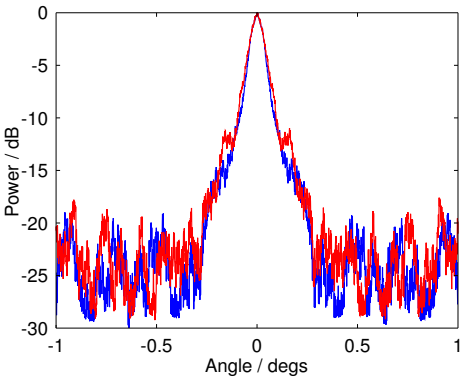

(d) Orthogonal one-dimensional cross-sections of synthesised beam at 110 MHz

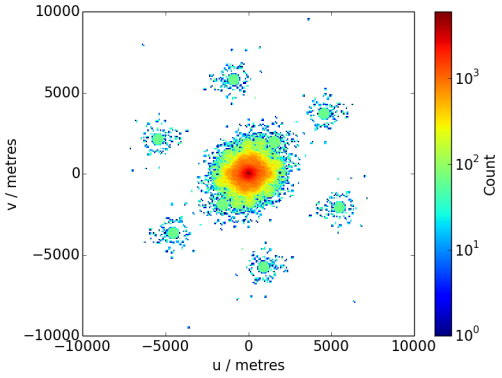

(e) Density of *uv*-coverage. *uv*-coverage is binned into spatial areas; the count of the total number of baselines in each bin is displayed. *uv*-coverage from the longest baselines are not shown.

Figure 4.15: Flexible cluster telescope model C.

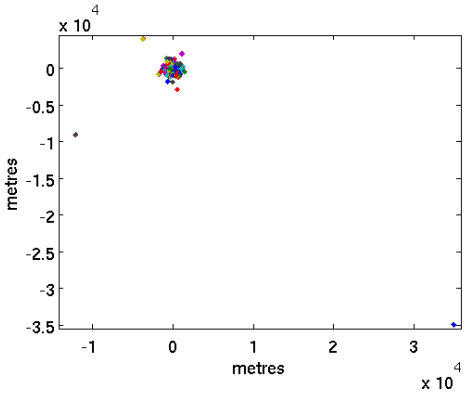

(a) Positions of all stations. There were 512 in supercluster and 507 surrounding the supercluster, with 8 stations per cluster. Of the surrounding clusters, 5 were in one spiral arms.

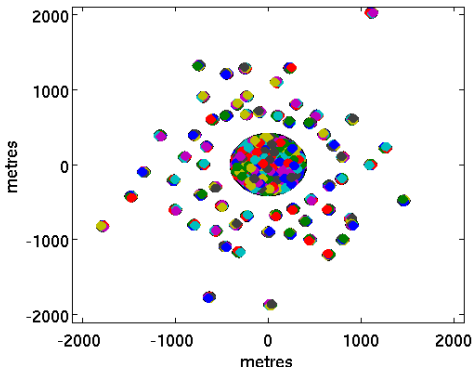

(b) Positions of central stations only. The central supercluster is visible as the large cluster in the centre. Stations are set-up for observing at the zenith (i.e. they are circular).

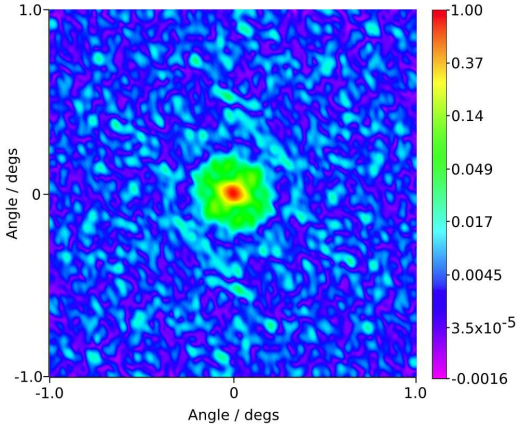

(c) Synthesised beam at 110 MHz

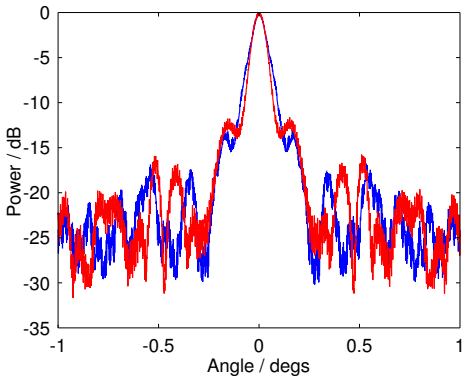

(d) Orthogonal one-dimensional cross-sections of synthesised beam at 110 MHz

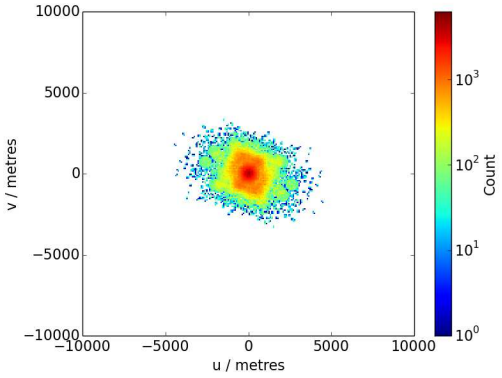

(e) Density of *uv*-coverage. *uv*-coverage is binned into spatial areas; the count of the total number of baselines in each bin is displayed. *uv*-coverage from the longest baselines are not shown.

Figure 4.16: Flexible cluster telescope model D.

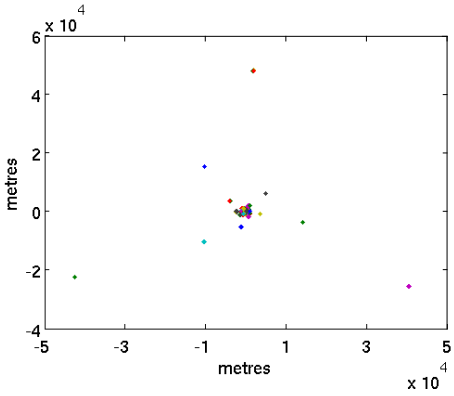

(a) Positions of all stations. There were 448 in supercluster and 576 surrounding the supercluster, with 16 stations per cluster. Of the surrounding clusters, 15 were in three spiral arms.

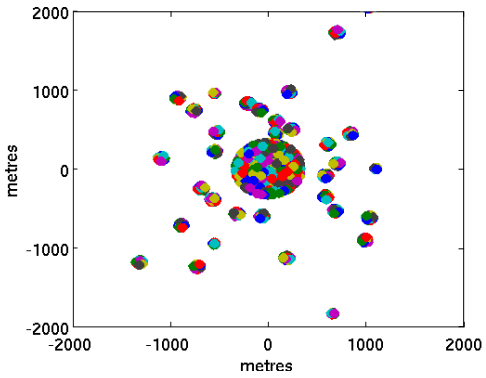

(b) Positions of central stations only. The central supercluster is visible as the large cluster in the centre. Stations are set-up for observing at the zenith (i.e. they are circular).

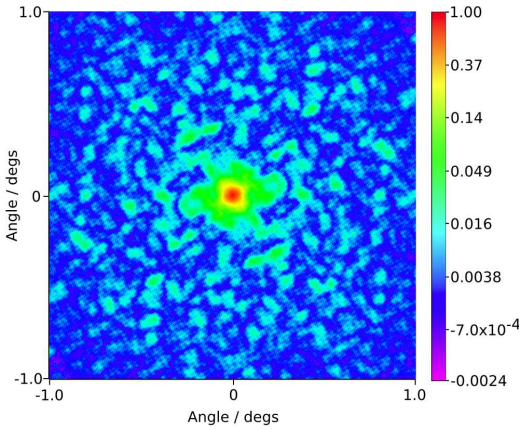

(c) Synthesised beam at 110 MHz

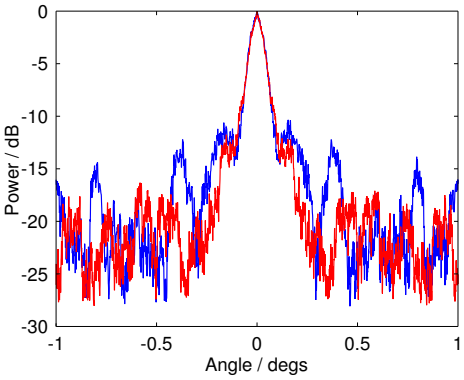

(d) Orthogonal one-dimensional cross-sections of synthesised beam at 110 MHz

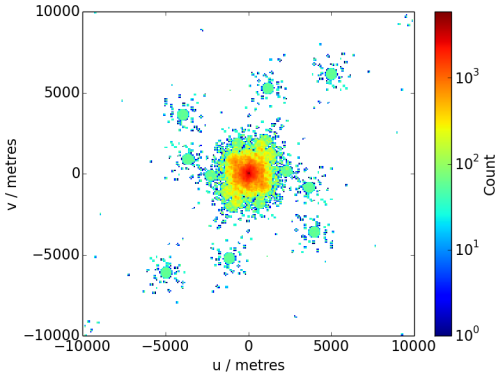

(e) Density of *uv*-coverage. *uv*-coverage is binned into spatial areas; the count of the total number of baselines in each bin is displayed. *uv*-coverage from the longest baselines are not shown.

Figure 4.17: Flexible cluster telescope model E.

consistently revealed itself in the form of strong grating lobes in the synthesised beams, hence this method was not used to create the final telescope models.

Another approach, not implemented, to increase the utilisation rate of antennas would be to have flatter or rounder tips to the starfish arms, as stations cannot be placed near the tips in the illustrated models.

When generating the starfish telescope models, stations are placed in locations if a sufficient number of antennas fall within the station's perimeter; the specific number of antennas used per station is allowed to vary within a range of ± 5 antennas. This has the advantage of allowing stations to occupy a greater number of positions within the starfish antenna area, but has the disadvantage of each station having a slightly different overall sensitivity. This could be accounted for by imposing a stricter limit on the permitted number of antennas per station or by weighting each station in the correlator. However, this is not incorporated here, as the deviation in sensitivity between stations is considered small, as there are 768 ± 5 antennas, where some or all of the extra antennas could be significantly down-weighted for apodization. It is also worth noting that the random positions of the antennas within a station mean that, when apodized, the total sensitivity of each station will vary by a small amount even if each station contains an equal number of antennas.

The large, dense core of the starfish models could potentially have an issue of access (especially vehicular) to the central antennas, hindering maintenance, for instance. Gaps for roads could be left between the antennas, but this would impinge on the flexibility of forming station. It would be prohibitive in terms of flexibility and cost of the beamforming system to allow complete freedom to form stations out of any combination of antennas within this flexible station concept, so the inclusion of roads would perhaps not be a significant extra hindrance.

4.4.4.3 Cluster models

The cluster design telescope models were created with variations in cluster sizes (that is, how many stations per cluster), how many stations occupy a central 'supercluster'

and how many spiral arms are used; some examples are shown in Figs. 4.13a, 4.13b, 4.14a, 4.14b, 4.15a, 4.15b, 4.16a, 4.16b, 4.17a and 4.17b. For each cluster size, five telescope models are generated with different random cluster and station positions to try and find more optimal symmetric synthesised beams from station distributions.

The cluster telescope models suffer from some of the potential problems noted for the starfish designs: access to central antennas, variations in number of antennas per station and different utilisation rates of antennas.

4.5 Beam patterns and uv-coverage of models

The beam patterns and uv-coverages of the following telescope models all use circular stations and zenith observations.

4.5.1 Discrete telescope models

The three discrete telescope models all give approximately the same synthesised beam pattern in terms of them being approximately Gaussian and of similar HPBWs as shown in Figs. 4.7c, 4.8c and 4.9c. However, the sidelobe level away from the peak of the beams changes from approximately -15 dB in the 50×5000 model to -30 dB in the 1024×256 model. The sidelobe level within the 50×5000 model oscillates significantly over the central 4 deg^2, ranging from -10 dB to -20 dB; the sidelobe range of 200×1250 model and 1024×256 model is considerably less.

Whilst all three models give approximately Gaussian beams, it is also clear from the 1-D cross-sections in Figs. 4.7d, 4.8d and 4.9d that the greater the number of stations, the more Gaussian the beams. This is a product of the more symmetric and uniform uv-coverage allowed by a greater number of stations. The deviations from symmetry could be mitigated by optimising the station locations, rather than picking the angular coordinate of the station randomly (the radial position is controlled and stations placed in polar co-ordinates).

The three apodization schemes (Gaussian, Hanning and none) were applied to the stations in the telescope models, but the effect on the synthesised beam patterns are not significantly different between apodization schemes in the central 4 deg^2 of the synthesised beams. Apodization will have affected the beam patterns far more significantly away from the beam centres.

The density of the *uv*-coverage of the discrete telescope models are shown in Figs. 4.7e, 4.8e and 4.9e. The *uv*-coverage is concentrated at short baselines in the 50×5000 and 200×1250 telescope models. The relatively low baseline counts are a result of the relatively low numbers of stations.

For the 1024×256 model the three spiral arms of the telescope are reflected in the *uv*-coverage. These longer baselines could result in sparse sampling of long baselines for snapshot observations, but with longer integrations Earth-rotation aperture synthesis will fill in the *uv*-coverage more evenly.

4.5.2 Starfish telescope models

The synthesised beams of the starfish telescope models are significantly non-Gaussian, as shown in Figs. 4.10c, 4.10d, 4.11c, 4.11d, 4.12c and 4.12d. For the straight-armed starfish telescope model (Fig. 4.10) the synthesised beam takes on a snowflake-like appearance. The starfish telescope models with curved arms result in a Gaussian central beam surrounded by a spiral; this is a result of the starfish's three curved arms. The synthesised beams of each of the starfish telescope models could make deconvolution difficult. The VLA and GMRT have similar instantaneous synthesised beams, but do not have the same sensitivity to detect as many point sources as the models here.

The starfish telescope models give poor snapshot *uv*-coverage, as shown in Figs. 4.10e, 4.11e and 4.12e, but the *uv*-plane would be filled in over a 4-hour observation. Deep science observations, such as those which will be used for the EoR key science project, will require long observations, which will ensure the *uv*-plane is filled in. However,

calibrating the data will require snapshot observations taken periodically throughout the observation, hence a Gaussian-like synthesised beam will be preferred.

The sidelobe levels of these beams are not as low as the 1024×256 discrete model, despite having the same number of stations. This can be attributed to the less uniform coverage of the *uv*-plane in the starfish models, as baseline separations are concentrated towards the arm spacing.

The significant non-Gaussian nature of the beam could perhaps be overcome by more thoroughly exploring the parameter space of arm curvature and thickness. There are limits on these parameters, for example the arms need to be a minimum thickness to allow sufficient antennas to form stations flexibly (if the arms are too thin it will not be possible to populate a whole station with the antennas). A small number of further starfish telescope models were created which varied the arm thickness and curvature within these limits, but the spirals remained prominent in the synthesised beams.

Further changes to the thickness, tapering and curvature of the arms would affect the radial distribution of the stations, meaning it would be necessary to deviate from the specification of 50% stations with 600 m and 75% stations within 1 km prescription.

4.5.3 Cluster telescope models

The synthesised beam patterns of various cluster-based telescope models are shown in Figs. 4.13c, 4.13d, 4.14c, 4.14d, 4.15c, 4.15d, 4.16c, 4.16d, 4.17c and 4.17d. The beam patterns and *uv*-coverage of the better performing cluster telescope models are shown (a greater range of telescope models were created and tested). The cluster beam patterns showed a tendency to form a central peak surrounded by a broad secondary beam at approximately -10 to -15 dB, which in turn was surrounded by the typical sidelobe level of approximately -20 to -25 dB. This is a result of the large compact cores used in most of the cluster models. This very strong weighting towards short

baselines is evident in the *uv*-densities Figs. 4.13e, 4.14e, 4.15e, 4.16e and 4.17e. The *uv*-densities also highlight the asymmetries present at long baselines if one or two spiral arms are used for the outer stations.

The *uv*-densities show that, in general, a greater number of small clusters give a more symmetric *uv*-coverage compared to a smaller number of large clusters. However, smaller clusters means a larger fraction of shared antennas (i.e. those antennas used in more than one station) are required in each cluster, as a larger fraction of antennas in each cluster will be situated towards the cluster edge and therefore down-weighted for apodization; this is because the ratio of circumference to area of a circle increases for smaller radii.

4.5.4 Comparison of models

To assess the potential usefulness of each telescope model, the similarity to Gaussian of their respective synthesised beams were measured. Gaussian-like beams facilitate more accurate deconvolution of images and calibration of the telescope as they have low sidelobe levels and Gaussian beams are typically chosen as restoring beams [38]. The least squares fit of each synthesised beam to Gaussians of full width half maximum from 6×10^{-4} to 5 degs, in step sizes of 6×10^{-4} degs, was calculated. The best fit values are plotted in Fig. 4.18.

It is clear that the 1024×256 telescope models (Gaussian, Hanning and no apodization windows) most closely approximated a Gaussian. The 50×5000 models were the least Gaussian; this is a result of the relatively low number of baselines in the 50×5000 model cases limiting the uniformity of the *uv*-coverage and hence the synthesised beams' Gaussianity.

The non-Gaussian shape of the starfish models' synthesised beams are also shown by this metric. The cluster models all have a smaller deviation from Gaussian than the starfish models, with the most Gaussian-like cluster models, A and C, having $\sim 70\%$ the RMS deviation of the most Gaussian starfish model.

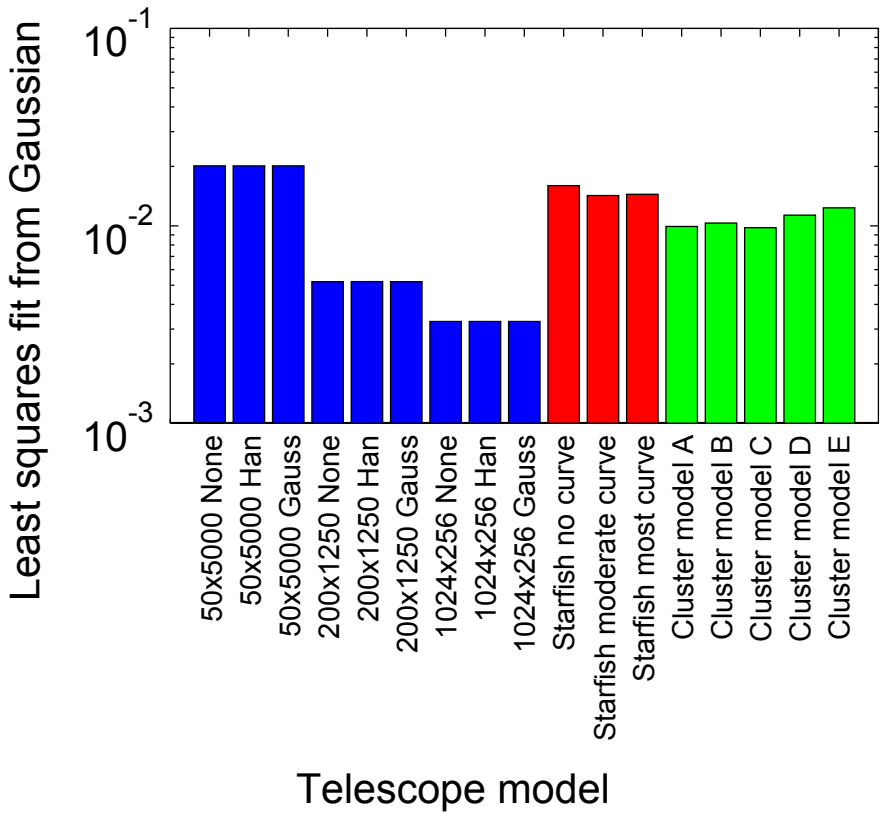

Figure 4.18: A least-squares fit was performed of Gaussians to each synthesised beam. The RMS deviation of the beams to the best-fit Gaussians are shown here. Only the central 4 deg^2 of each synthesised beam was considered. The bars of the discrete telescope models are coloured blue, starfish models in red and cluster models in green.

The synthesised beams of the starfish models tested, within the constraints of flexible stations as described here, did not perform favourably in comparison to the *Requirements specification*'s discrete 1024×256 telescope model. Not utilising a small fraction of antennas ($\sim 5 - 15\%$, depending on the parameters of the starfish model) and considerable overlapping of some stations were also concerns for this approach. Consequently, the starfish telescope models, as they are implemented here, are unlikely to be viable for SKA1-Low.

The beam patterns of the cluster models are not as promising as the 1024×256 model beam pattern, but the potential advantages of flexible stations, such as elliptical stations, apodization and calibrating your station set to optimise the telescope for specific science cases warrants further investigation.

4.6 Summary

The flexible stations concept was discussed and potential application to SKA1-Low explored. A range of telescope models, flexible and discrete, were created which conform to the SKA *Requirements specification* on antenna number and distribution, but with varying station sizes and allowing for antennas to be shared between stations. The effect of apodization was also discussed.

Fitting the synthesised beams of the telescope models to Gaussians indicated that the 1024 station with 256 antennas per station telescope model (the current SKA-low1 design) yields the most Gaussian beams, whilst the starfish shaped flexible telescope models and the 50 stations with 5000 antennas per station model were the least Gaussian. The non-Gaussianity of the starfish models' beams suggests images may prove difficult to deconvolve so the telescope model designs generated here should not be pursued further.

Chapter 5

Sensitivity of station models to the Epoch of Reionisation power spectrum

5.1 Introduction

The Cosmic Dawn and Epoch of Reionisation mark an important period in the early universe. The Cosmic Dawn marks the formation of the first stars and the Epoch of Reionisation saw radiation from early galaxies begin to ionise the majority of the intergalactic medium (IGM). SKA1-Low will attempt to observe these eras by means of redshifted 21-cm line emission; neutral hydrogen emits 21-cm line radiation whilst ionised hydrogen does not. In this chapter, the Cosmic Dawn and Epoch of Reionisation will collectively be referred to as the EoR, unless otherwise indicated.

Measurements of the EoR power spectrum will be one of several techniques used for these observations. A range of telescopes, such as PAPER, MWA and LWA are currently attempting to measure the EoR power spectrum. In theory, these telescopes have the sensitivity to detect the EoR power spectrum, but their success will be dependent on systematics and the actual strength of the EoR power spectrum. Whether they are successful or not, knowledge gained from these telescopes, for instance calibration techniques and treatment of sky models, will help guide SKA1-Low. The greater sensitivity of SKA1-Low will, in theory, allow the EoR to be imaged.

The EoR signal will vary in strength over a range of redshifts and angular scales in

the sky. In this chapter, various telescope models for SKA1-Low will be considered to ascertain their suitability for measuring the EoR signal. This is achieved by evaluating the ability of each telescope model to detect the EoR power spectrum, whilst also accounting for various systematic effects which introduce noise signals.

Both discrete and flexible telescope models, as discussed in Chapter 4, are investigated. The thermal sensitivity of the telescope as a function of angular scale will be compared to the expected EoR signal. Also considered are the ability of the telescope models to detect and resolve foreground point sources, the expected residual noise levels of unsubtracted point sources within the observing beam, and the level of residual noise due to imperfect calibration as studied in Chapter 3.

5.2 The 21-cm line

A hydrogen atom has two hyperfine energy levels in the 1s ground state, arising from the magnetic interaction between the proton and electron spins. When the spin of the electron flips from parallel to the proton's spin to antiparallel, a photon of 21-cm wavelength is emitted. Likewise, 21-cm radiation can be absorbed through the electron spin flipping in the opposite direction.

Whether 21-cm radiation is seen in emission or absorption depends on whether the spin-temperature of neutral hydrogen is greater than or less than that of the photons which illuminate it; in the early Universe, these are Cosmic Microwave Background (CMB) photons. The greater the difference in spin temperature, the more easily the 21-cm radiation can be detected.

Over the history of the Universe, 21-cm radiation has varied in strength and from absorption to emission. The brightness temperature of 21-cm radiation, relative to the CMB, is given by [121, 122]

$$\delta T_{\rm b} = 27 x_{\rm HI}(1+\delta)\sqrt{\frac{1+z}{10}}\left(\frac{T_{\rm s}-T_{\rm CMB}}{T_{\rm s}}\right)\left(\frac{\Omega_{\rm b}}{0.044}\frac{h}{0.7}\right)\sqrt{\frac{\Omega_m}{0.27}} \times \left(\frac{1-Y_{\rm p}}{1-0.248}\right)\left(1+\frac{1}{H(z)}\frac{dv_\parallel}{dr_\parallel}\right)^{-1} {\rm mK}\,, \quad (5.1)$$

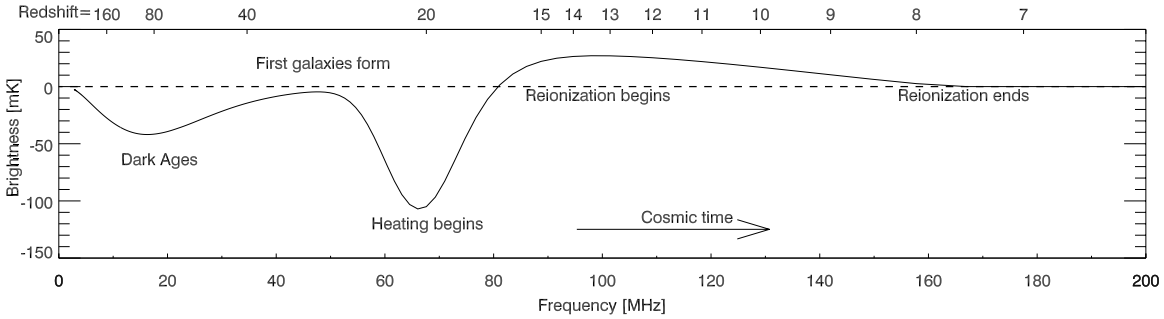

Figure 5.1: Sky-averaged 21-cm radiation brightness from the beginning of the Dark Ages to the end of the Epoch of Reionisation. Figure from Pritchard & Loeb [123].

where $x_{\rm HI}$ is the neutral fraction of hydrogen, δ is the fractional overdensity in baryons, T_s the spin-temperature of HI, $T_{\rm CMB}$ the CMB temperature, h the dimensionless Hubble parameter, $Y_{\rm p}$ the primordial Helium mass fraction, $H(z)$ the Hubble constant and $dv_\parallel/dr_\parallel$ the velocity gradient along the line of site. The expected brightness, as a function of redshift, is plotted in Fig. 5.1.

The first hydrogen atoms were formed from protons and electrons combining as the universe cooled following the Big Bang. This, the Epoch of Recombination, occurred at $z \sim 1000$ [124]. Between the Epoch of Recombination and $z \sim 200$, the difference in spin-temperature between HI and the CMB photons was unobservable, leading to this period being aptly known as the Age of Ignorance. The Dark Ages followed ($30 \lesssim z \lesssim 200$) in which the Universe's adiabatically cooled neutral gas absorbed 21-cm line radiation.

Towards the end of the Dark Ages, around $z \sim 30 - 40$, the expansion of the Universe caused the gas density to decrease sufficiently so that the HI spin temperature couples back to the CMB photons rather than the neutral gas; the ensuing decrease in the spin-temperature difference between HI and the CMB once again made 21-cm radiation difficult to detect.

At $z \sim 30$ the first stars and X-ray sources began to form. The 21-cm transition is a forbidden transition, but Lyman-α coupling, in which neutral hydrogen is raised from its $n = 1$ to its $n = 2$ state and then back to $n = 1$, with the spin of the electron flipping in the process, is far more likely to occur in the presence of stars.

This coupled the spin-temperature of HI back to the cool neutral gas, leading to the 21-cm line being seen in absorption [123]. This period is the Cosmic Dawn.

Gradually the stars and X-ray sources heated their surrounding neutral gas, increasing the spin-temperature of HI to greater than the temperature of CMB photons, allowing the 21-cm line to be seen in emission for the first time at $z \sim 15$. The radiation and heating also saw neutral gas become ionised. Ionised gas originally appeared in small pockets, or bubbles, surrounding the most powerful ionising sources. Overtime the bubbles expanded and merged with one another until they eventually encompassed the majority of the Universe around $z \sim 5-6$; this period is the Epoch of Reionisation [122].

5.2.1 21-cm science

Measuring reionisation with 21-cm radiation will provide valuable insight into the early evolution of galaxies. Reionisation is theorised to be a function of the combined extreme ultra-violet flux in the IGM, allowing estimates of cosmic star formation to be made. The types of galaxies principally responsible for reionisation can also be established by mapping reionisation, as well as the contribution of quasars [85] and dark-matter annihilation [125, 126].

Reionisation also provides a means to study the IGM, as the morphology of HI and HII regions is related to the temperature structure and density fluctuations of the IGM [127, 128].

Perturbations in the density of HI during the EoR are predicted to follow underlying perturbations in the cold dark matter, allowing measurements of fundamental cosmological parameters, such as the total mass density, Ω_{m}, the global baryonic density, Ω_{b}, and the scalar spectral index, n_{s}, to be made [129, 130, 131, 132]. Observations at multiple redshifts will allow a much greater amount of data pertaining to these cosmological parameters to be collected than with the 2-D surface of the CMB [85].

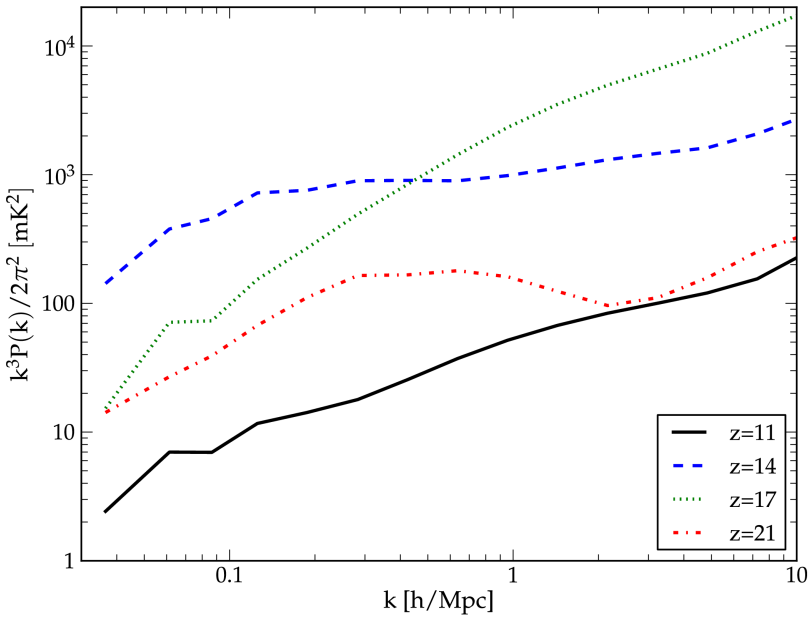

Figure 5.2: Simulated EoR power spectrum for a range of redshifts as a function of wavenumber, k. Figure from Santos et al. [133].

5.3 Power spectrum sensitivity

Measurements of the EoR power spectrum are an important task for SKA1-Low, and studies of the EoR are considered as the primary science project of SKA1-Low. It is for this reason that the telescope's specifications have changed significantly in light of the conclusions of Mellema et al. [85]. For instance, station sizes were reduced from 180 m to 35 m to provide greater fields of view, as it is claimed that EoR fluctuations and structures cannot be accurately recovered on scales larger than that of a sample beam. The frequency range was also changed, from $70 - 450$ MHz to $50 - 350$ MHz, to more accurately cover the redshift range of the 21-cm line from the Cosmic Dawn through the EoR.

It is also important to consider the distribution of baselines of the SKA1-Low stations, so the telescope is sensitive to the most interesting angular scales. Power spectra measurements typically require compact arrays sensitive to large angular scales, from tens of arcminutes to degrees, whereas tomographical imaging of the EoR requires

arcminute-scale resolutions and hence long baselines. A telescope array with a range of baselines will enable measurements of both to be accomplished.

Simulations have provided an insight into the expected EoR power spectrum. The values of the expected EoR power spectrum are given, amongst others, in Mellema et al. (which includes data from Santos et al. [133]). Other studies yielding expected values for the EoR signal include McQuinn et al. [134] and Furlanetto, Zaldarriaga and Hernquist [135]. Expected values for a range of redshifts, as given in [133], are plotted in Fig. 5.2.

The EoR power spectrum is conventionally given in units of $k^3P(k)/2\pi^2$, where k is the wavenumber in units of h/Mpc. This is the power spectrum in 3-D space where a flat power spectrum corresponds to equal powers in logarithmic intervals of $|k|$. Observations in a single redshift slice however give constraints on the 2-D power spectrum. To compare with predicted measured the 3-D power spectrum is here converted to 2-D as a function of the observed angular scale expressed in terms of the spherical harmonic multipole coefficient ℓ. The 3-D power spectrum in k-space can be converted into a 2-D power spectrum in ℓ-space by [136]

$$\Delta^2_{\rm 2D} = k_1^2 \int_{k_1}^{\infty} \Delta^2_{\rm 3D}(k) \frac{k^{-2}}{\sqrt{k^2 - k_1^2}} {\rm d}k \,, \tag{5.2}$$

where $\Delta^2_{\rm 3D}$ is the 3-D power spectrum $k^3P(k)/2\pi^2$, $\Delta^2_{\rm 2D}$ is the 2-D power spectrum $C(\ell)\,\ell(\ell+1)/2\pi$ and the integral is over k-space for each value in the EoR k-space signal k_1. k-space is converted to ℓ-space by

$$\ell = khD_{\rm A} \,, \tag{5.3}$$

where h is the dimensionless Hubble parameter and $D_{\rm A}$ is the angular diameter distance. ℓ relates to baseline lengths by

$$\ell = 2\pi\sqrt{u^2 + v^2} \,, \tag{5.4}$$

where u and v are the baseline lengths in the north-south and east-west directions, respectively.

The 2-D power spectrum is a slice through the 3-D power spectrum. Considering the EoR signal in a 2-D power spectrum has the advantage of illustrating the powers of its modes to angular scales on the sky. The distribution of baselines in SKA1-Low will dictate how sensitive it is to different angular scales on the sky, hence the ability of SKA1-Low to detect the EoR can be directly compared for different telescope models.

SKA1-Low must be capable of detecting this power spectrum; therefore the theoretical angular power spectrum sensitivity of any proposed telescope model should be checked against the simulated EoR angular power spectrum levels. The theoretical angular power spectrum sensitivity, ΔT^2, of a telescope model is given by

$$\Delta T^2 = \frac{\ell(\ell+1)}{2\pi} S^2_{\text{sample}} \left(\frac{\Delta\ell}{N_s}\right) \left(\frac{1}{\Omega_{\text{beam}}}\right) \left(\frac{\lambda^2}{2k_B}\right) , \tag{5.5}$$

where N_s is the number of baselines corresponding to each multipole bin, $\Delta\ell$ is the width of each multipole bin, λ is the wavelength and k_B is the Boltzmann constant. Ω_{beam} is the solid angle of the square primary beam, defined by

$$\Omega_{\text{beam}} = \int_0^{2\pi} \int_0^{\pi} P^2(\theta,\phi)\, \sin(\theta)\, \mathrm{d}\theta\, \mathrm{d}\phi\,, \tag{5.6}$$

where $P(\theta,\phi)$ is the power of the beam in the direction (θ,ϕ). In this study, the limits of the integral in Eq. (5.6) are taken to be the angles at which the beam power drops below 5% of the peak power. S_{sample} is the sensitivity of a single visibility from one baseline (composed of identical stations) and is given by [105]

$$S_{\text{sample}} = \frac{k_B}{\sqrt{\Delta\nu\, t_{\text{obs}}}} \left(\frac{A_{\text{eff}}}{T_{\text{sys}}}\right)^{-1} \frac{1}{\eta_s}\,, \tag{5.7}$$

where $\Delta\nu$ is the channel bandwidth, t_{obs} the observation length, A_{eff} the effective collecting area, set to $1.93^2\,\text{m}^2$ [83], T_{sys} the system temperature and η_s the system efficiency, assumed to be 0.9. T_{sys} can be modelled according to Eq. (3.3) in Chapter 3, which are repeated here,

$$T_{\text{sys}} = T_{\text{sky}} + T_{\text{rcvr}}\,, \tag{5.8}$$

where

$$T_{\rm rcvr} = 0.1\,T_{\rm sky} + 40{\rm K}\,, \tag{5.9}$$

where $T_{\rm sky}$ is the sky temperature and $T_{\rm rcvr}$ is the receiver temperature.

$T_{\rm sky}$ is determined by the mean temperature of a typical patch of sky the telescope would observe. The Haslam 408-MHz all-sky map [106, 107, 137] can be used to provide an estimation of $T_{\rm sky}$.

A power law can be assumed to convert the measured temperatures at 408 MHz to the other frequencies for the EoR power spectrum. This is given in Eq. (3.4), but is repeated here for convenience as Eq. (5.10),

$$T_\nu = T_{\nu_0}\left(\frac{\nu}{\nu_0}\right)^\beta\,. \tag{5.10}$$

Here β, the temperature spectral index of diffuse galactic emission, is equal to -2.55; this assumes a synchrotron flux density power law index of -0.55, which is equivalent to -2.55 in temperature units. The synchrotron power law spectral index, α, varies between approximately -0.55 and -0.7 for different sources and frequencies. ν_0 is the template map frequency (408 MHz) and the target frequency and T_{ν_0} and T_ν are the template map temperature at 408 MHz and the target frequency, respectively.

5.3.1 Point-source power spectrum

To ensure accurate measurements of the EoR power spectrum, the effect of foreground point sources must be accounted for. If the point-source power spectrum exceeds the EoR power spectrum, individual point sources must be removed from the telescope's visibilities.

An estimation of the point-source power spectrum must hence be determined. The combined power from point sources, dP, per flux bin, dS, per steradian, is the product of the number of sources per flux bin, $N(s)$, and the power per source, S^2, per unit steradian,

$$\mathrm{d}P = N(s)S^2\,\mathrm{d}S\,. \tag{5.11}$$

Intema et al. [138] determine the Euclidean-normalised differential source counts from a 153-MHz catalogue as

$$\frac{\mathrm{d}N}{\mathrm{d}S} S^{\frac{5}{2}} = A\, S^{0.91}\,, \tag{5.12}$$

where $A = 6000\,\mathrm{Jy}^{0.59}\,\mathrm{sr}^{-1}$. Euclidean-normalised source counts give the actual source count relative to the source count that would be present if the sources were uniformly distributed in a flat, non-evolving Euclidean universe. In the Euclidean case, the number of sources at a distance r in a volume element of thickness $\mathrm{d}r$ is proportional to $r^2\mathrm{d}r$ and the apparent flux of a source is proportional to r^{-2}, hence the differential source count $\mathrm{d}N/\mathrm{d}S$ is proportional to $S^{-5/2}$.

Substituting Eq. (5.12) into Eq. (5.11) yields

$$P = \int_{S_{\mathrm{low}}}^{S_{\mathrm{up}}} A\, S^{0.41}\, \mathrm{d}S\,, \tag{5.13}$$

where S_{low} and S_{up} are the lowest and greatest flux values of the point sources to be included in the point-source power spectrum. Substituting in a power law equivalent to Eq. (5.10) but with flux density with spectral index $\alpha = -0.7$ gives

$$P = \int_{S_{\mathrm{low}}}^{S_{\mathrm{up}}} A\, \left(\frac{\nu}{\nu_0}\right)^{-0.29} S_{\nu_0}^{0.41}\, \mathrm{d}S\,. \tag{5.14}$$

Integrating Eq. (5.14) leads to

$$P = B \left(\frac{\nu}{\nu_0}\right)^{-0.29} {S_{\mathrm{up}}}^{1.41}\,, \tag{5.15}$$

where $B = 6000/1.41\,\mathrm{Jy}^{0.59}\,\mathrm{sr}^{-1}$ is a constant and S_{low} is set to 0.

The total flux density per unit angle, $\Sigma S\,/\,\Omega$, is related to the average background temperature due to sources, T, and wavelength by the Rayleigh-Jeans approximation,

$$\frac{\Sigma S}{\Omega} = \frac{2k_B T}{\lambda^2}\,, \tag{5.16}$$

hence

$$T^2 = \left(\frac{\lambda^2}{2k_B}\right)^2 \left(\frac{\nu}{\nu_0}\right)^{-0.29} B\, {S_{\rm up}}^{1.41}\,. \tag{5.17}$$

The contribution of the point sources to the measured power spectrum can be reduced through two methods: decreasing the value of $S_{\rm up}$ and by spectral fitting. The EoR power spectrum is expected to fluctuate on small spectral scales due to the structure in the redshift direction, whilst foregrounds such as the point-source power spectrum will be spectrally smooth as they are generally composed of a small number of power-law components. Through spectral fitting of the smooth foreground, it should be possible to subtract it from the EoR power spectrum [139, 140, 141, 142].

Harker et al. [143] propose a non-parametric spectral fitting technique termed 'Wp smoothing' (from the German 'Wendepunkt', which translates as 'inflection point'). The method penalises variations in the fitting function's curvature and has shown fit accuracies of $\sim 1\%$ or better in simulations, a greater level of accuracy than alternative methods (for example [144, 145, 146, 147], although the complexity of the foreground model used in simulations influences the accuracy of a fit). Consequently, the residual point-source power spectrum level drops to $\sim 1\%$ of its initial value. Wp smoothing is, however, a relatively compute-intensive approach and hence is relatively slow to process large data cubes; Moore's law will hopefully help with this over the time it will take to build SKA1-Low.

Other foreground removal methods have been investigated with respect to EoR detection, which are also expected to be able to allow SKA1-Low to detect the EoR power spectrum [148]. One of the aims of the MWA precursor telescope is to be a test bed for spectral smoothing techniques [74].

The complementary approach to decreasing the power of the point-source power spectrum, decreasing $s_{\rm up}$, requires that every point source of flux greater than $s_{\rm up}$ must be accounted for and its presence removed from the telescope's visibilities. To do this requires knowledge of the source positions and brightnesses, as well as an accurate understanding of the telescope's signal processing.

To determine the source positions and brightnesses, the telescope must be able to detect it and isolate it from surrounding point sources. The smallest flux density a telescope can theoretically detect is given by its thermal noise limit. In practice, to reliably detect a point source its flux must be ~ 5 times brighter than the thermal noise limit, to provide a sufficient signal-to-noise ratio.

To separate the point-source contribution from the EoR signal the telescope must have baselines long enough to resolve the point sources out of the EoR. The telescope must also be able to separate point sources from one another, necessitating a resolution better than the typical source separation of the faintest detectable sources. Using the differential source counts of Intema et al. [138] the number of sources per steradian of flux equal to, or greater than, s_{up} is given by

$$N = 6000 \int_{s_{\mathrm{up}}}^{S_{\mathrm{max}}} s^{-1.59} \, \mathrm{d}S \,, \tag{5.18}$$

where S_{max} is the maximum flux considered in the source count. This is taken to be the flux at which there is only one source per steradian, as given by Eq. (5.12) ($\approx 250\,\mathrm{Jy}$).

In general, the mean separation of point sources, $\Delta\theta$, can be determined from the number of sources per steradian, N_{sources}, by

$$\Delta\theta = \sqrt{\frac{1}{N_{\mathrm{sources}}}} \,, \tag{5.19}$$

therefore the longest baseline length required to resolve the point sources, d_{max}, is

$$d_{\mathrm{max}} \approx \frac{\lambda}{\Delta\theta} \,. \tag{5.20}$$

The maximum baseline needed is hence dependent on the sensitivity of the telescope on the longer baselines. Any sources not resolved and removed from images will contribute to the residual point-source power spectrum, which will contaminate the EoR signal.

As well as point sources, there will be foreground contamination of the EoR through extended sources, such as free-free emission from HII regions in the ISM and IGM, and synchrotron radiation from cluster radio haloes and cluster radio relics [149]. These extended sources are not accounted for in this analysis, however spectral smoothing will help reduce the effect of extended sources.

5.4 Calculating the power spectra

5.4.1 Residual noise

To determine the residual noise levels of the telescope, a similar technique to that discussed in Chapter 3 is employed. Here the total residual noise is measured rather than separating the in-field and sidelobe residual noise values, hence the entire VLSS point source sky model is used with the exception of the six brightest sources. Measurements are simulated at 65, 79, 95, 110, 119 and 130 MHz, which correspond to redshifts 21, 17, 14, 12, 11 and 10 of the 21-cm line.

5.4.2 Thermal noise limit

The thermal noise limit of the telescope models is calculated using Eq. (5.7). To determine the thermal noise level of a set number of baselines of identical stations, Eq. (5.7) can be extended to [105]

$$S_{\text{sample}} = \frac{k_B}{\sqrt{\Delta\nu t_{\text{obs}}}} \left(\frac{A_{\text{eff}}}{T_{\text{sys}}}\right)^{-1} \frac{1}{\eta_s} \frac{1}{\sqrt{N_{\text{B}}(N_{\text{B}}-1)}}\,, \tag{5.21}$$

A bandwidth of 3 MHz, A_{eff} of $1.93^2\,\text{m}^2$ and an efficiency of 0.9 are assumed.

T_{sys} is determined from Eqs. (5.8) and (5.9) and T_{sky} is calculated using the Haslam sky map. A power law can be assumed to convert the measured temperatures at 408 MHz to the other frequencies for the EoR power spectrum, in accordance with Eq. (5.10), where S_{ν_0} and S_ν are the map's temperature at 408 MHz and the target frequency, respectively.

The value of N_{B} is determined for each telescope model by calculating the number of baselines of length greater than 1.6 km (corresponding to $\ell = 10\,000$). This length was chosen so as to ensure that only baselines on which the EoR signal is fully resolved out were counted for the purpose of detecting confusing sources, although the exact value selected is by no means definitive.

The thermal noise levels were compared to the point-source fluxes; the point sources should have a flux at least five times greater than the thermal noise limit to ensure detection.

5.4.3 EoR power spectra

To determine the angular power spectrum sensitivity of the telescope models, Eq. (5.5) is employed. The number of samples in each multipole bin is determined from the telescope model *uv*-coverage. Baselines are converted to multipole numbers using Eq. (5.4) and binned into bins of $\Delta\ell = 10$.

The beam solid angle, Ω_{beam}, is determined from the station's simulated primary beam pattern. Station beam simulations are obtained using the OSKAR-2 simulator. The total angular area covered by all pixels in the simulated primary beam within 5% of the peak power is used as the beam solid angle.

The theoretical EoR power spectrum plotted are from Fig. 5.2 here and Fig. 21 in [85], converted to 2-D power spectra using Eqs. (5.2) and (5.3).

5.5 The angular sensitivity of various telescope models

The EoR power spectra for $z = 10 - 20$ over a range of multipoles are plotted in Fig. 5.3 alongside the thermal noise sensitivity, point source and total residual noise power spectra of a range of telescopes models. The telescope models comprise discrete and cluster models, as shown in Figs. 4.7 to 4.9 and 4.13 to 4.17.

The plots in Fig. 5.3 contain several separate power spectra overlaid on the same plots. The EoR power spectrum is shown by the thick dashed lines. The thin solid lines depict the thermal sensitivity of the telescope model plotted; the telescope model is stated in the subfigure captions. The dotted lines show the expected total residual noise level from a 10% gain error. The dot-dashed lines show the expected point-source power spectrum. Different colours represent different observation frequencies.

5.5.1 Discrete models

Focussing initially on the power spectra of the discrete telescope models, shown in Figs. 5.3a to 5.3i, there is a clear decrease in the thermal sensitivity of the telescope models as the stations are apodized by greater amounts. Each Gaussian apodized telescope model has an approximately 13 times poorer thermal noise limit than its equivalent unapodized station, whereas a Hanning window apodized telescope decreases the sensitivity by approximately a factor of 1.7. This trend is not unexpected as apodization decreases the sensitivity of individual antennas by down-weighting their contribution to beamformed data and the Gaussian window function is an especially strong apodization window.

There is also a trend of increasing point-source power spectrum levels with increasing apodization. This is because the point-source power spectra are calculated from the thermal noise limit of the long baselines (i.e. baselines exceeding 1.6 km) of the telescope. It has been assumed that all the point sources above the long-baseline thermal noise limit in the observations are resolved and removed from visibilities; the remaining sources below the thermal noise limit contribute to the point-source power spectrum. No improvement by spectral smoothing has been accounted for in the figures, hence the residual point-source power spectra should be expected to decrease. The level of this decrease depends on the success of the spectral smoothing algorithm employed, but, as mentioned in Section 5.3.1, the Wp smoothing technique could decrease the point-source power spectrum by approximately two orders

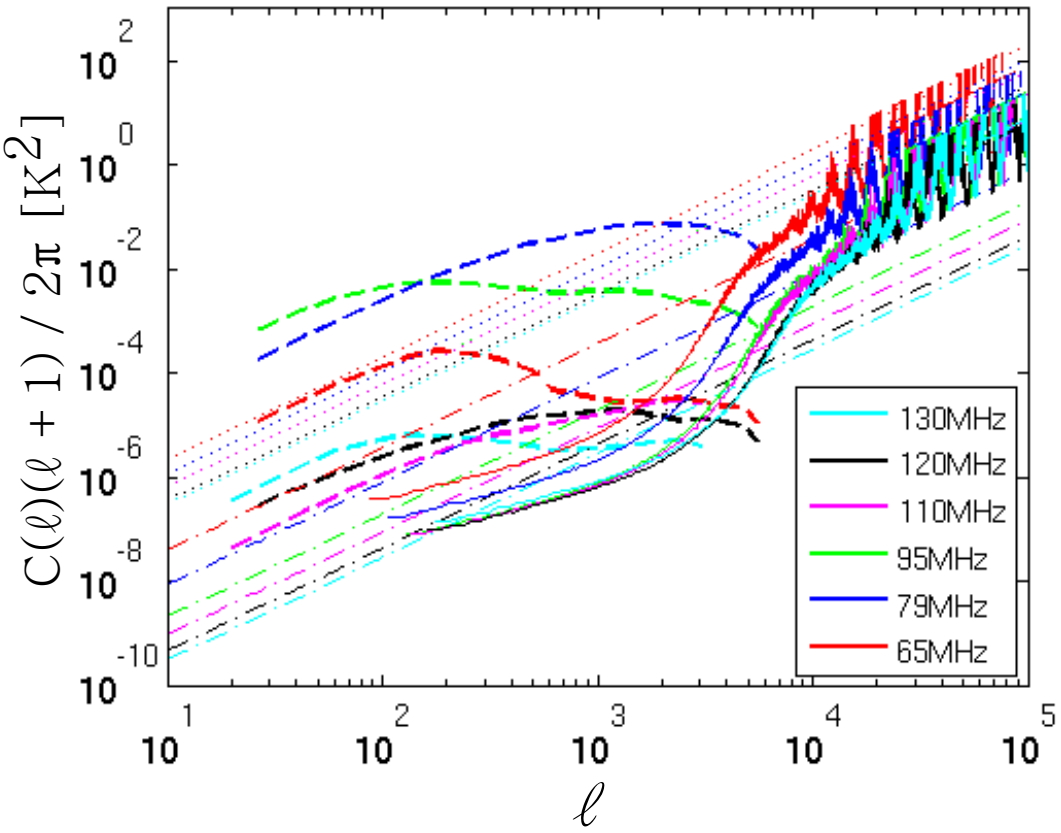

(a) 1024 × 256, no window

Figure 5.3: Power spectra plots showing the expected power level of the EoR (dashed lines), thermal sensitivity of the telescope model (solid lines), residual point-source power spectra after subtraction of detected individual sources (dot-dashed lines) and total residual noise levels from a gain error of 10% (doted lines) at different angular scales, defined by the multipole, ℓ. No benefits from spectral smoothing have been assumed for the point-source power spectra, but it could decrease the power by a factor of ~ 100. The different coloured lines show the power spectra at different frequencies: red 65 MHz ($z = 21$); blue 79 MHz ($z = 17$); green 95 MHz ($z = 14$), magenta 110 MHz ($z = 12$); black 119 MHz ($z = 11$); cyan 130 MHz ($z = 10$). Expected EoR power levels are from [85, 133]. The total residual noise levels are plotted for a 12-hour integration (beyond this residual noise levels may not decrease as baselines are repeated with earth rotation aperture synthesis) and 10 000-hour integration for the thermal noise level. Figs. 5.3a to 5.3f and 5.3h to 5.3n are on the following pages.

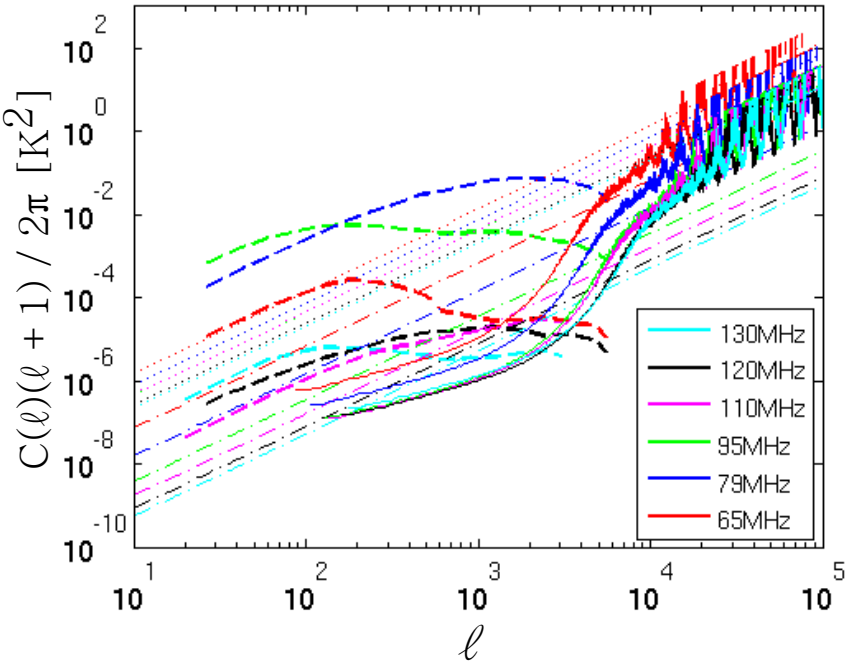

(b) 1024 × 256, Hanning window

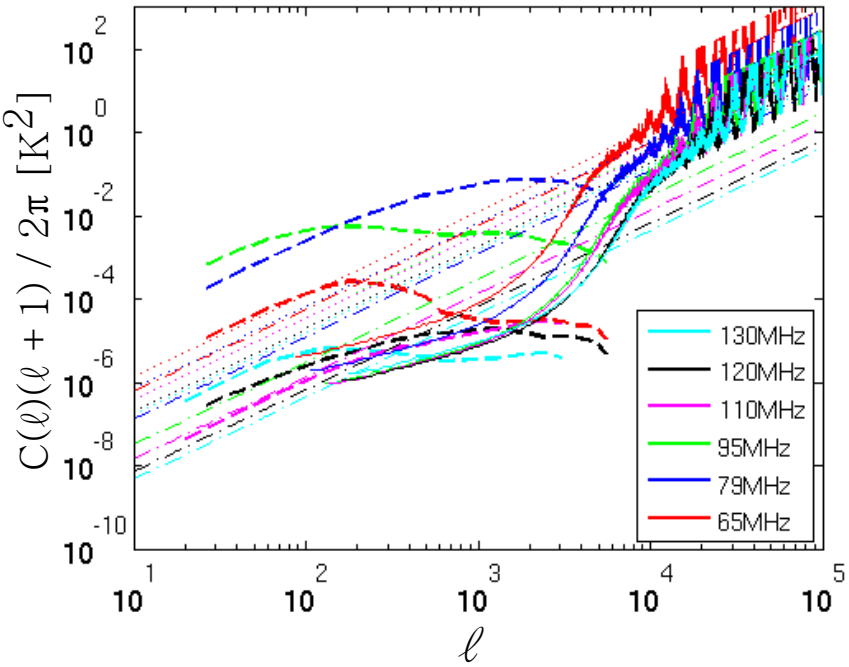

(c) 1024 × 256, Gaussian window

Figure 5.3: continued (2/8).

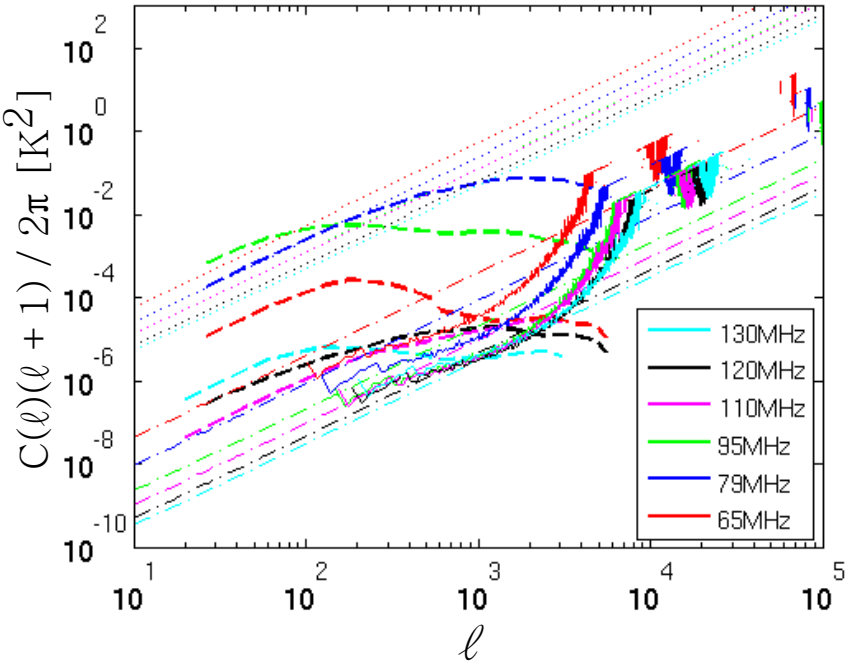

(d) 200 × 1250, no window

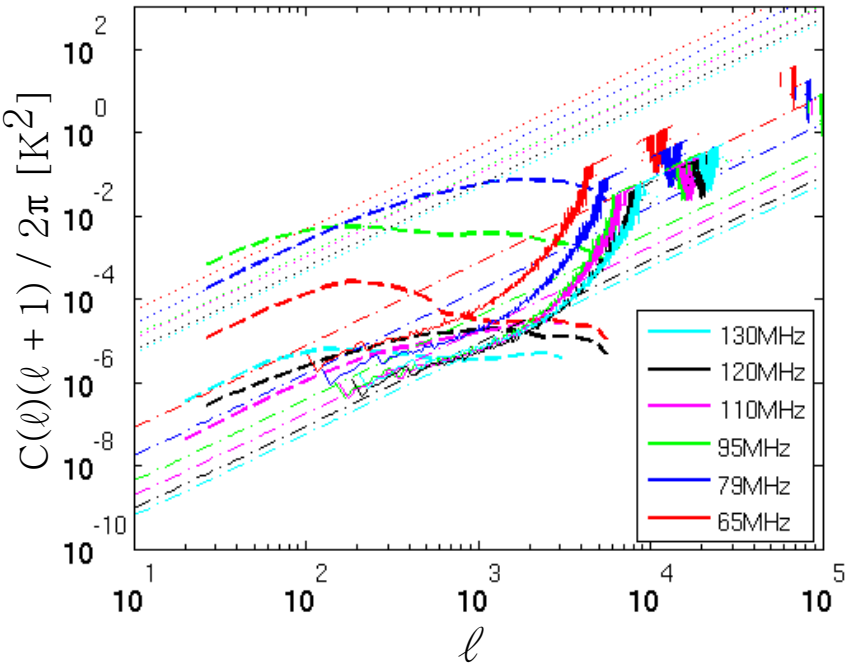

(e) 200 × 1250, Hanning window

Figure 5.3: continued (3/8).

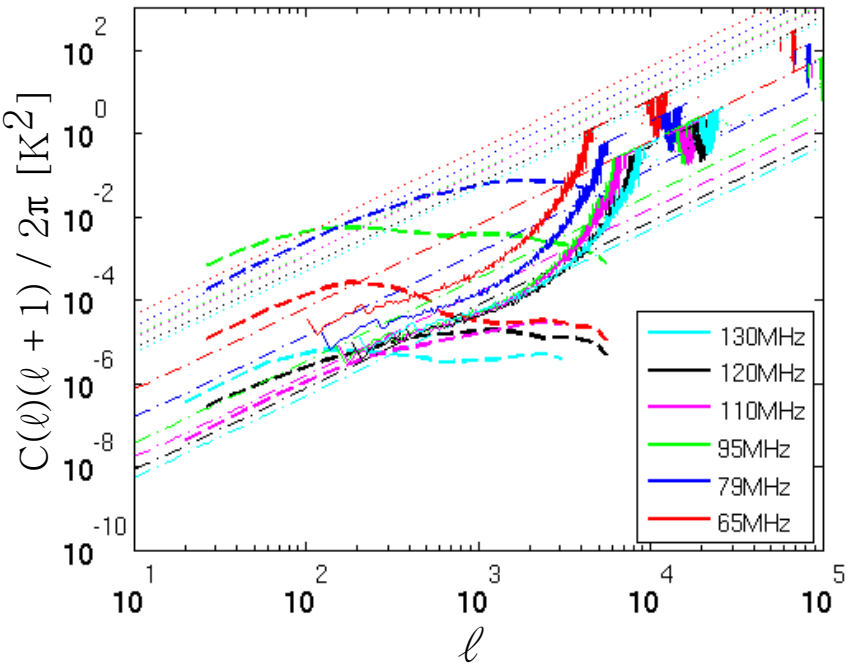

(f) 200 × 1250, Gaussian window

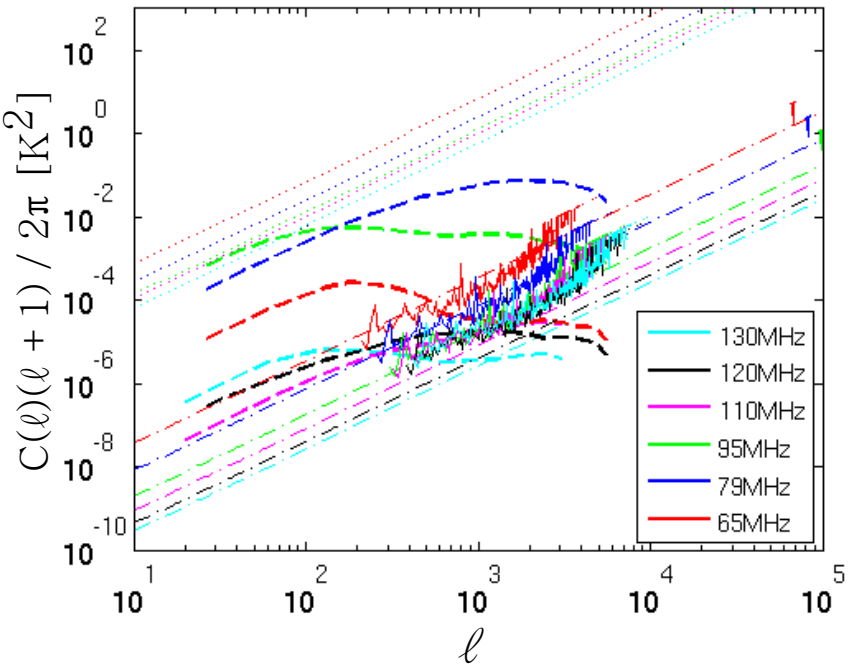

(g) 50 × 5000, no window

Figure 5.3: continued (4/8).

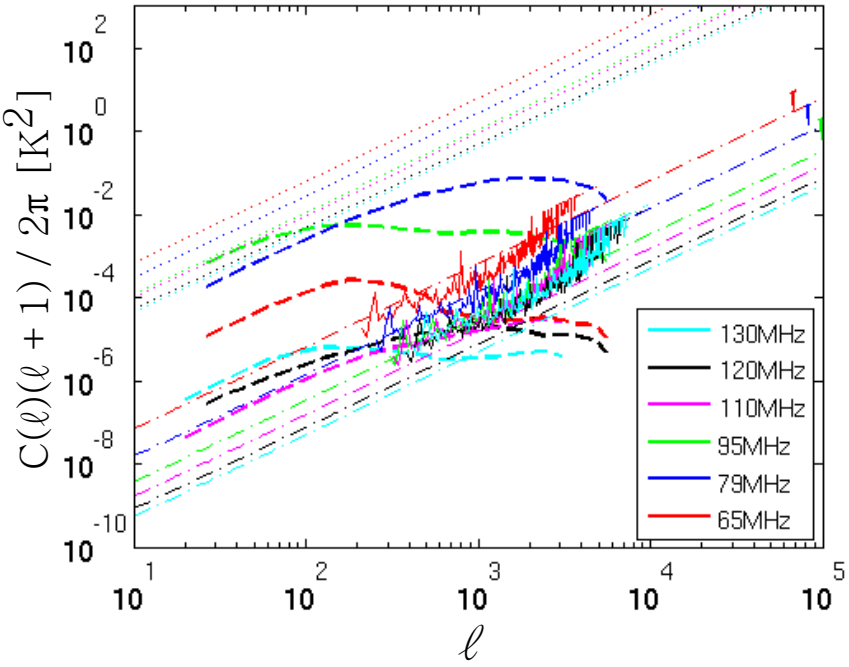

(h) 50×5000, Hanning window

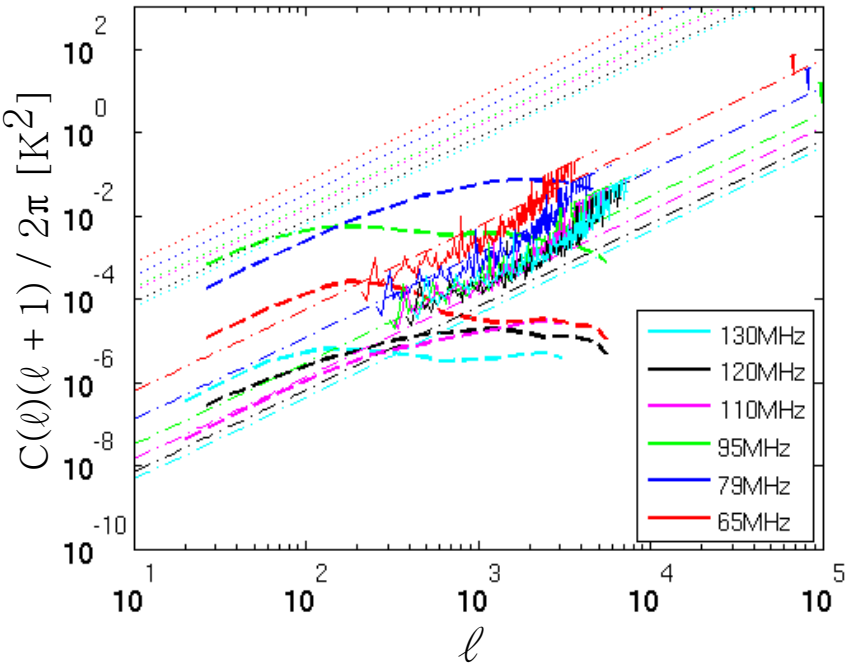

(i) 50×5000 Gaussian window

Figure 5.3: continued (5/8).

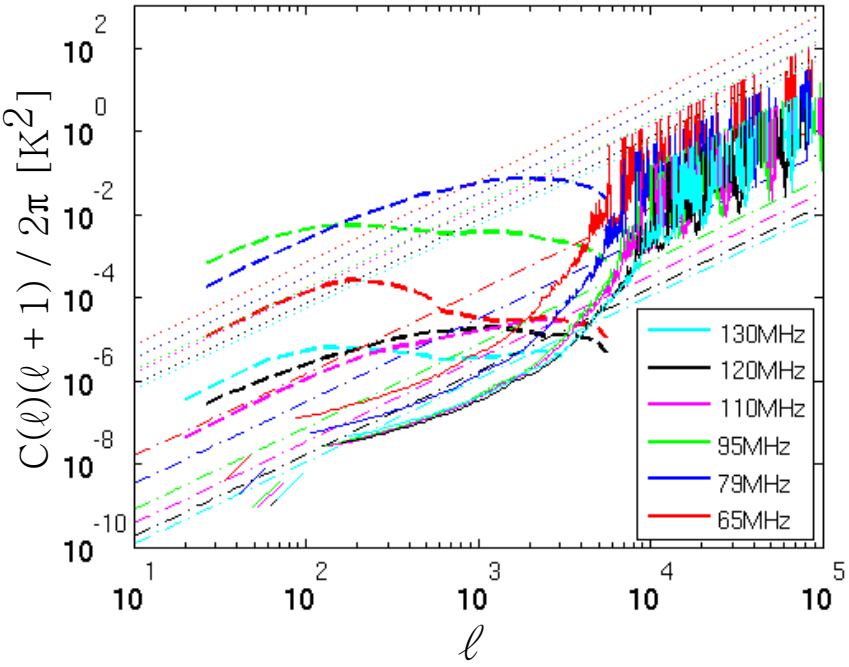

(j) Cluster model A, Hanning window

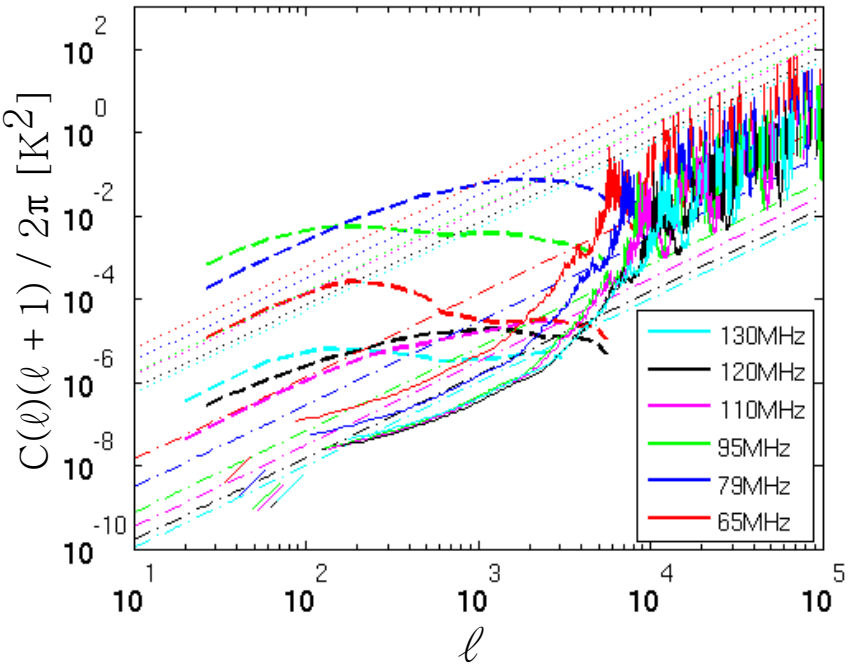

(k) Cluster model B, Hanning window

Figure 5.3: continued (6/8).

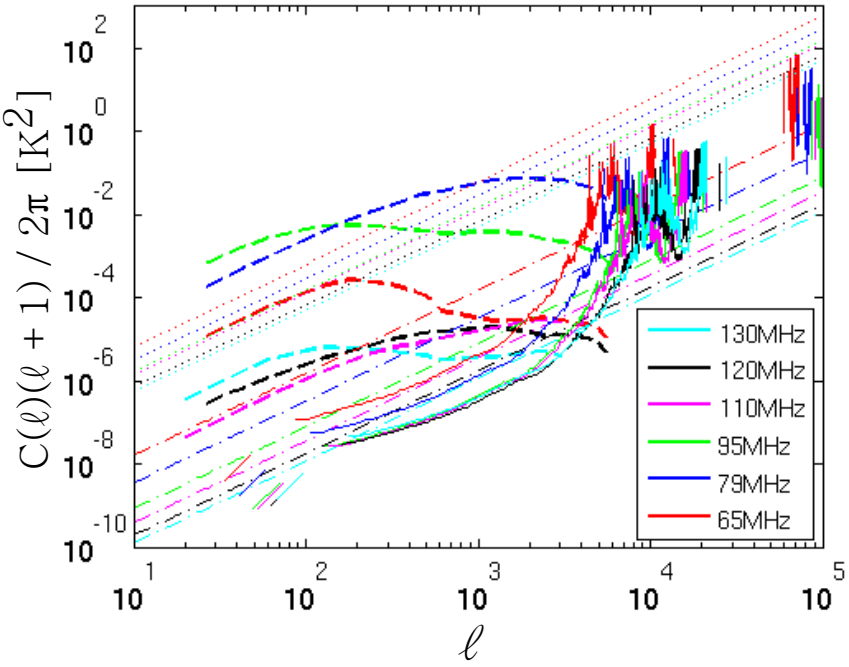

(l) Cluster model C, Hanning window

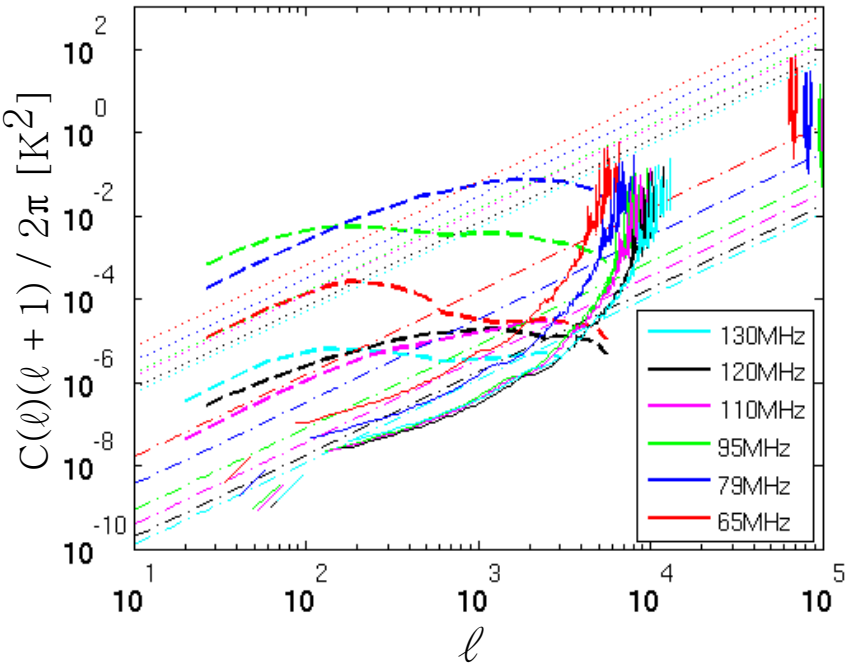

(m) Cluster model D, Hanning window

Figure 5.3: continued (7/8).

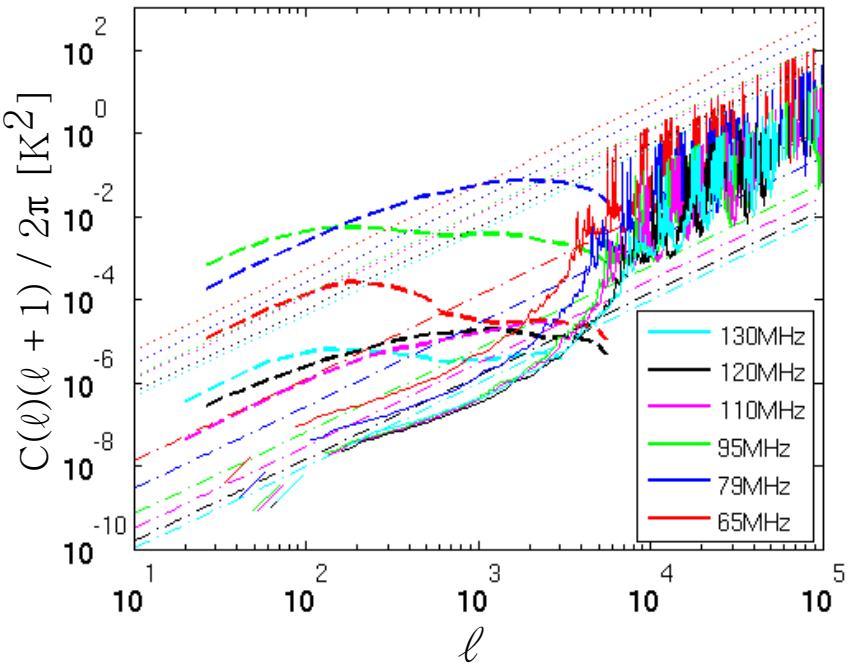

(n) Cluster model E, Hanning window

Figure 5.3: continued (8/8).

of magnitude.

5.5.1.1 Total residual noise

The total residual noise power spectra levels of the telescope models with no apodization are slightly greater than those with a Hanning window applied, which are in turn slightly greater than those with the Gaussian window. The exception to this is the 50×5000 model, where the Gaussian window results in an increase in the power spectra levels. The 50×5000 model aside, the general trend of decreasing residual noise with increasing apodization is expected, as near-in sidelobe levels are reduced with apodization.

What is surprising is that the RMS difference between the ideal and error-telescope images increase with apodization. This is shown in Fig. 5.4, where the Gaussian apodization telescopes have approximately $2 - 5$ times as high total residual noise levels as the Hanning window and no apodization cases. However, when the total

residual noise is plotted in the power spectra, it decreases with apodization. This is a result of the solid angle of the beam being a factor in the power spectra calculation (see Eq. (5.5)). The primary beam widths of the discrete telescope models across the frequency range investigated for different apodization windows are plotted in Fig. 5.5a and the beam solid angles of the primary beams are plotted in Fig. 5.5b. Stronger apodization clearly increases the beam width and solid angles, along with decreasing the effective number of antennas in a station.

Another possible factor for the increase in total residual noise is that the majority of the contribution to the total residual noise comes from sources in the main beam, as opposed to sources in the sidelobes. To test this, further simulations were run with sources within twice the HPBW of the main beam removed from the sky model. This hence determined the residual sidelobe noise, which is plotted with the EoR, point source and thermal sensitivity power spectra in Fig. 5.6 for the three discrete telescope models with no apodization window applied. The residual sidelobe noise levels are lower than the total residual noise levels by a factor of approximately 2.5, which is significant.

The effect of apodization on the shape of the primary beams is shown in Fig. 5.7. The primary beams of a station in each of the discrete telescope models for each of the three apodization windows are plotted. Increasing the apodization results in a broader main beam and reduces the number and prominence of near-in sidelobes; the remaining sidelobes are broader. Stronger apodization also increases the power of far-out sidelobes by small amounts.

This could explain the increase in total residual noise level with more apodization. The broader main beam has approximately the same power as the replaced sidelobes for most of its width, but doesn't have the low power levels of the troughs between sidelobes. This means sources around the centre of the primary beam direction will be detected with a greater flux density. As shown in Chapter 3, the greater the total flux density of sources in an observation, the greater the residual noise level that results.

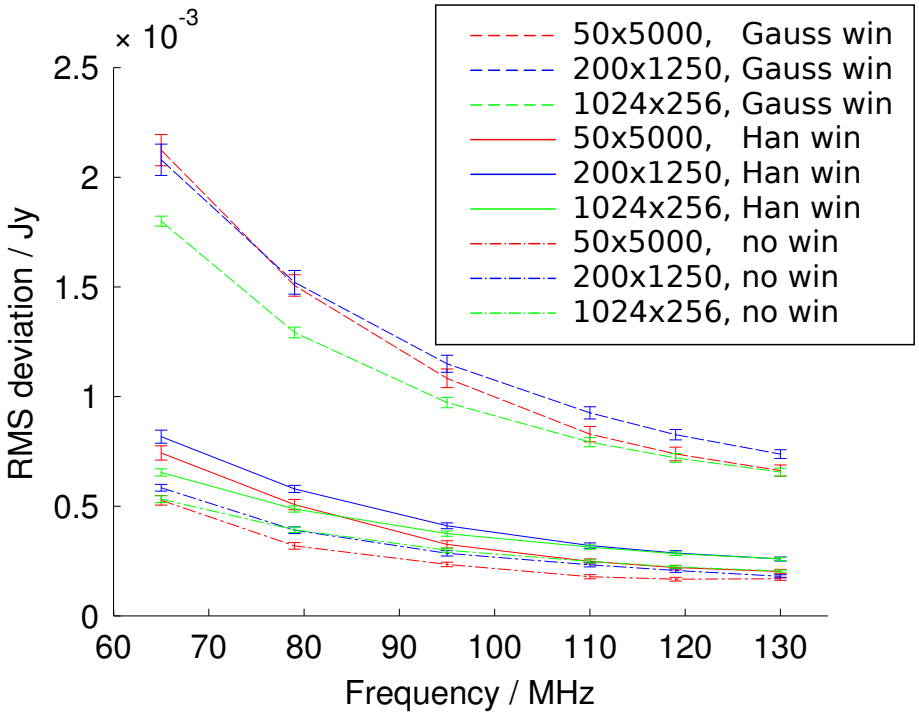

(a) Discrete telescope models

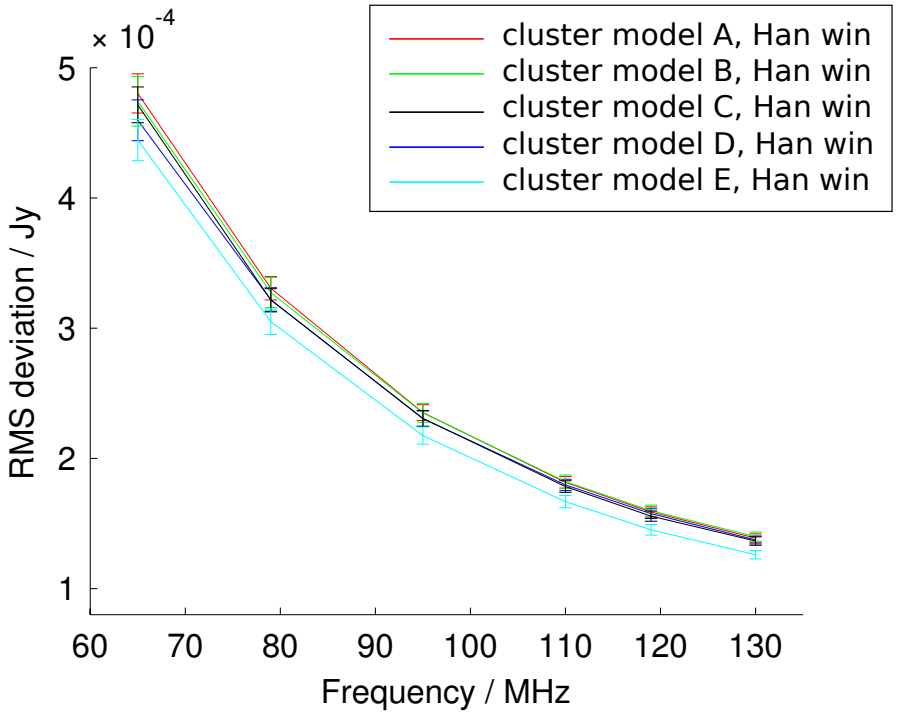

(b) Cluster telescope models

Figure 5.4: Total residual noise from simulated visibilities of various discrete and cluster telescope models. 10-s observations were simulated. For discrete telescope models, three apodization windows were used; Gaussian ('Gauss win' in legend, Hanning ('Han win') and none ('no win'). For cluster models a Hanning window was used.

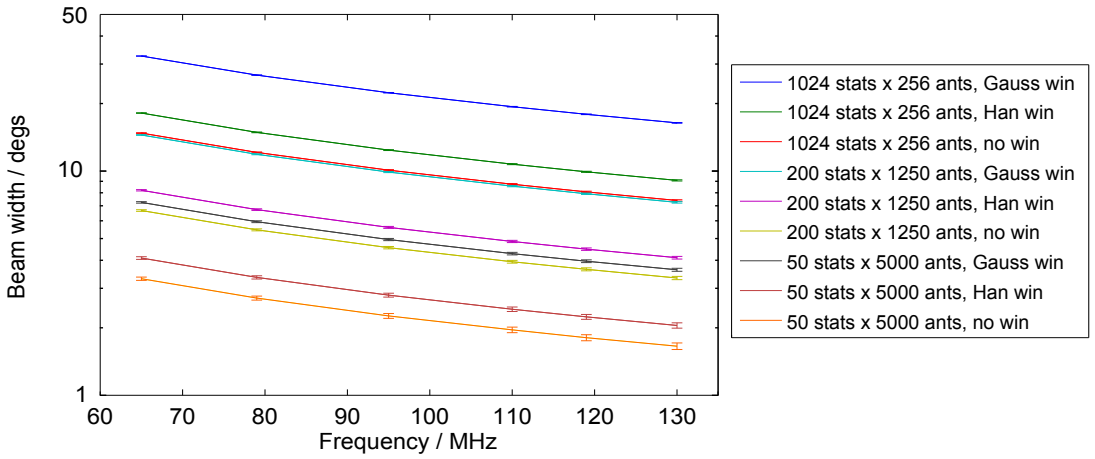

(a) Width of primary beam

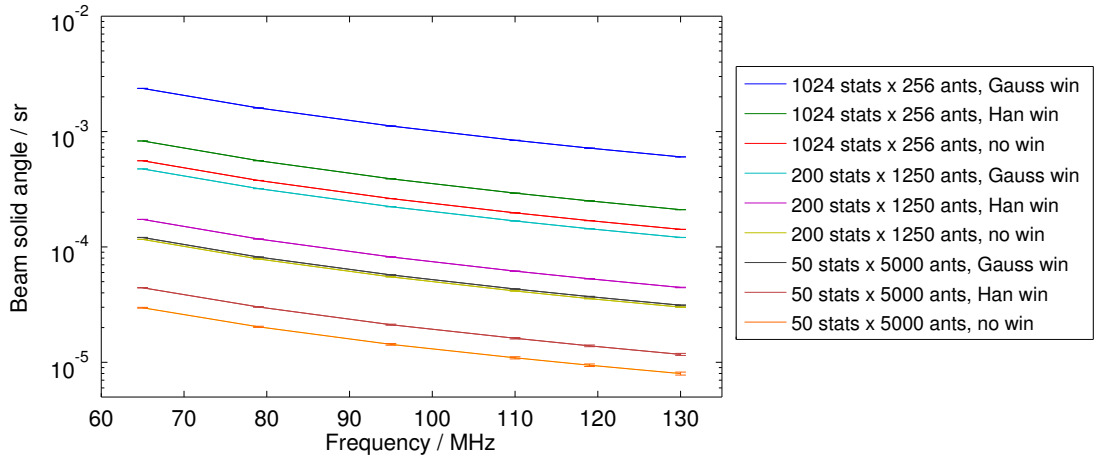

(b) Primary beam solid angle

Figure 5.5: Width of primary beam, and primary beam solid angle, as a function of frequency for different station sizes. The stations have three different apodization windows applied: Gaussian ('Gauss win'), Hanning ('Han win') or none ('no win').

Furthermore, the raised far-out sidelobe level, caused by the decreased sensitivity from apodization, means that the primary beam is more sensitive to far-out sidelobe sources.

The primary beam patterns in Fig. 5.7 also highlight that stations with greater numbers of antennas have more preferable primary beam patterns: their main beams are narrower and their mean sidelobe levels are lower. The lower total residual noise level of the 1024×256 discrete model, with small stations, compared to the 50×5000 discrete model, with large stations, hence suggests that uv-coverage is a more

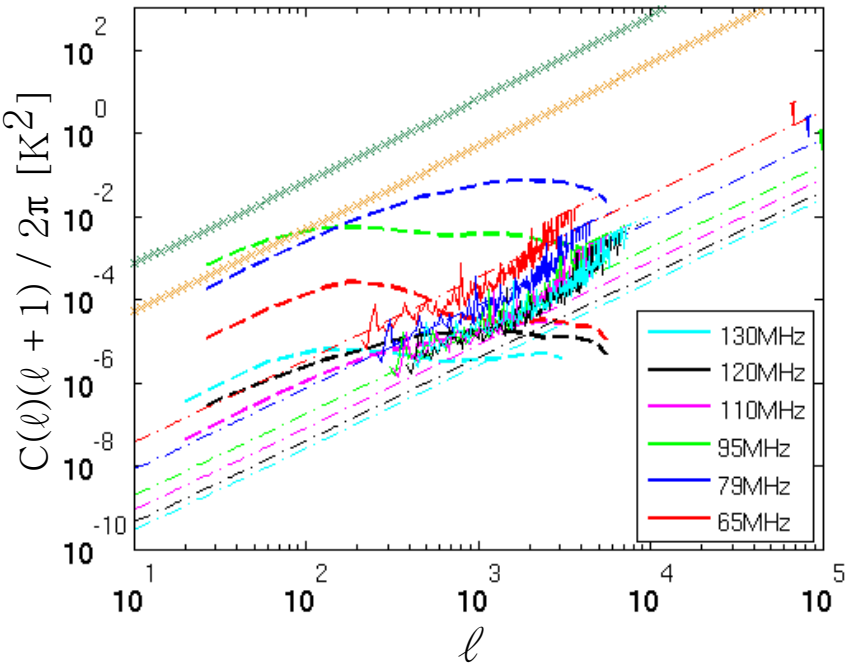

(a) 50×5000, no window

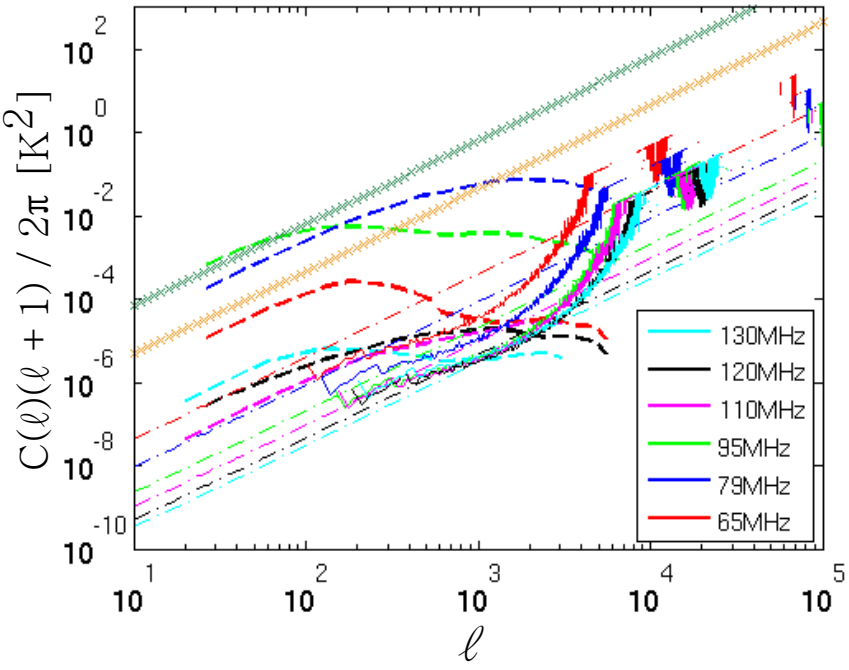

(b) 200×1250, no window

Figure 5.6: Power spectra plots of the discrete telescope models shown in Fig. 5.3, except the residual sidelobe noise from 10% gain error on antennas is plotted rather than the total residual noise. The residual sidelobe noise at 50 MHz (dark green) and 100 MHz (orange) are shown by the starred lines. Expected EoR power levels are from [85, 133]. Figure 5.6b is on the following page.

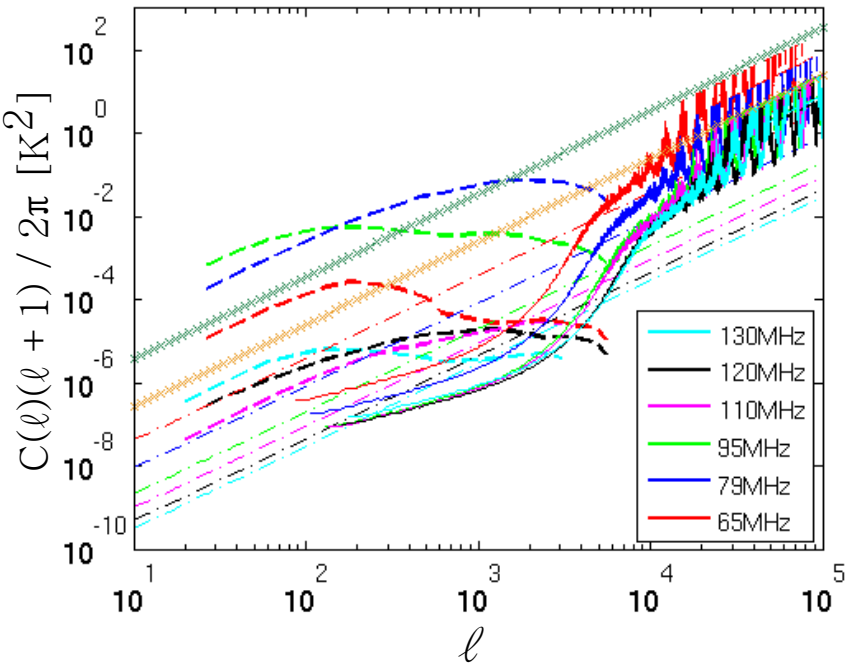

(c) 1024×256, no window

Figure 5.6: continued (2/2).

important factor in minimising the total residual noise level than the primary beam pattern.

5.5.1.2 Benefits of primary beams compared to *uv*-coverage

To test the relative effect of primary beam pattern compared to *uv*-coverage on the total residual noise, a range of discrete telescope models were created. Models to test the effect of the primary beam pattern have a fixed number of stations, 50, and have 256, 512, 1024, 2048 or 4096 antennas per station. Each model has the same, fixed, station positions. To ensure that sensitivity does not affect the results of simulations with these telescope models, the gains of the antennas are down-weighted relative to the total number of antennas in the telescope model. For instance, the 50 stations with 256 antennas per station model has a weighting of 1 for each antenna and the 50 stations with 512 antennas per station model is weighted at 0.5 for each antenna.

A comparison of the relative effect of *uv*-coverage is made by creating a range of

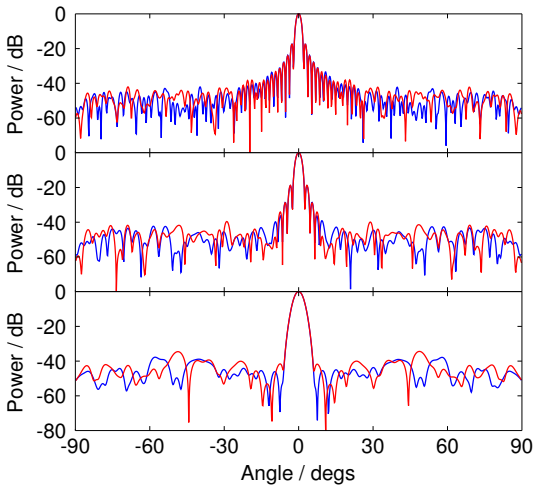

(a) 5000-antenna station

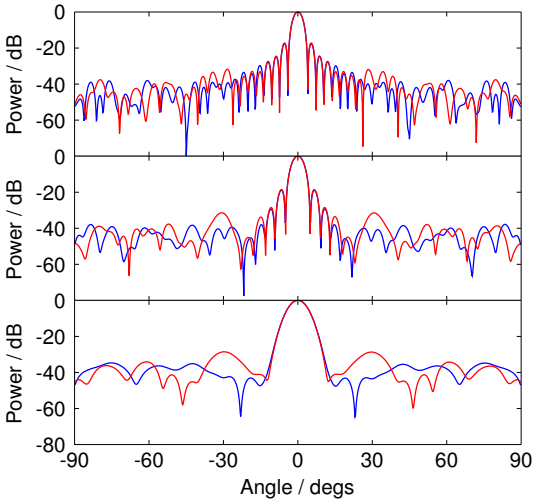

(b) 1250-antenna station

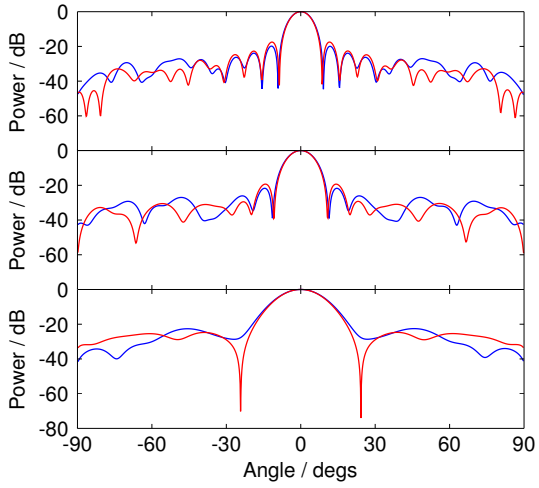

(c) 256-antenna station

Figure 5.7: One-dimensional cuts of the primary beams of a single station of each of the three discrete telescope models, with varying degrees of apodization. Other stations, in their respective telescope models, will have similar primary beams. Within each subfigure, the upper plot shows the primary beam with no apodization, the middle plot with a Hanning window described in Eq. (4.1) and the lower plot a Gaussian window illustrated in Fig. 4.3. Red and blue lines shows orthogonal cuts through the centre of the primary beam.

discrete telescope models with a fixed number of 256 antennas per station, but a range of numbers of stations: 50, 100, 200, 400 and 800. Each fixed number of antennas per station telescope model has an equal number of antennas to one of the fixed number of stations telescope models. Again, antennas are down-weighted to retain a constant sensitivity across all telescope models. Station locations are distributed according to the *Requirements specification*.

The total residual noise from telescopes with gain errors of 10% are compared for each telescope model. Observations are simulated at 110 MHz, with 10-s integrations, observing the zenith of the VLSS point source sky model.

Figure 5.8 shows that increasing the station sizes and increasing the number of stations both decrease the RMS deviation caused by errors in the telescopes. This is expected, due to the improving primary beam, and the improved *uv*-coverage yielding synthesised beams with lower sidelobes.

Figure 5.8 suggests that the RMS deviation decreases more rapidly when increasing the number of stations than when increasing the number of antennas per station. This suggests that narrowing the primary beam is not as important as increasing the density of the *uv*-coverage. However this is not consistent with the results in Fig. 5.4a, where a smaller number of larger stations yields a lower total residual noise than a larger number of smaller stations, when no apodization is applied. On the other hand, Fig. 5.4a suggests that for a Hanning window below ~ 80 MHz, or a Gaussian window between $65 - 130$ MHz, a large number of small stations is the better distribution of antennas.

Combining the data from both Fig. 5.4a and Fig. 5.8 suggests that there is little difference in the overall effect of RMS deviation from portioning the antennas into small or large stations, at least within the ranges investigated here. The differences shown in the plots are probably due to factors such as sky model, observing direction, station positions and antenna positions within each station; the latter two of which are randomly generated within the confines of the required radial distribution and

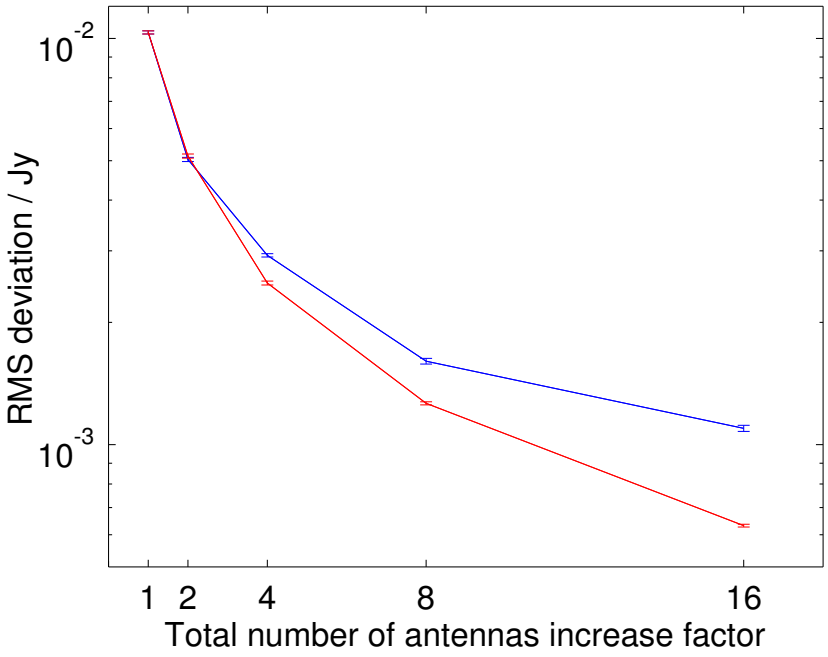

Figure 5.8: RMS deviation between ideal and error simulated images. Simulations were made with a range of telescope models with increasing numbers of antennas, either by increasing the number of antennas per station, or increasing the total number of stations. The RMS deviation for increasing the number of antennas per station by multiples of two from 256 to 4096, with 50 stations, is shown in blue. The RMS deviations for increasing the number of stations by multiples of two, from 50 to 800 stations with 256 antennas per station, are shown in red. Antennas were uniformly down-weighted to keep the total sensitivity of the telescope constant across all telescope models, however no apodization was applied. The simulated images were made at 110 MHz, observing the VLSS sky model at the zenith. Gain errors of 10% were applied to the antennas.

station and antenna sizes.

5.5.1.3 Sample variance

A further contribution to the noise level of measuring the EoR power spectrum is the sample variance. The sample variance is the error arising from the finite number of measurements from which estimates of the EoR signal can be made. It is determined by the expected EoR power divided by square root of the number of unique beam areas corresponding to the angular scale being measured in the field of view. As such, it is not dependent upon the distribution of baselines, assuming the uv-coverage is

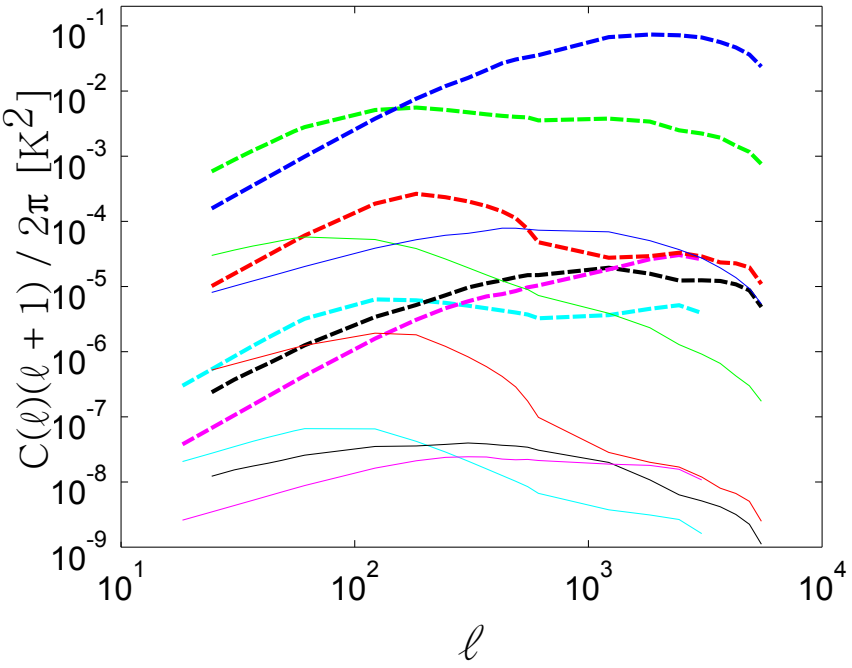

Figure 5.9: Sample variance of EoR power spectrum. A 25 deg^2 field of view is assumed. Dashed lines plot the EoR power spectrum at different frequencies, solid lines the corresponding sample variance. The different coloured lines show the power spectra at different frequencies: red 65 MHz ($z = 21$); blue 79 MHz ($z = 17$); green 95 MHz ($z = 14$), magenta 110 MHz ($z = 12$); black 119 MHz ($z = 11$); cyan 130 MHz ($z = 10$). Expected EoR power levels are from [85, 133].

approximately uniform.

The EoR power spectrum and sample variance are plotted in Fig. 5.9 for a single 25 deg^2 field of view. The sample variance will decrease by observing multiple 25 deg^2 fields of view.

5.5.1.4 Suitability for Epoch of Reionisation observations

The 1024 × 256 telescope models are sensitive to a wider range of angular scales than the other discrete models. This is due to both the small stations allowing shorter baselines and the large number of stations allowing more continuous coverage of baseline lengths out to the longest baselines. The thermal noise and total residual noise levels of the 1024 × 256 models are also significantly lower than the 200 × 1250 and 50 × 5000 models. These factors all support the current design of 1024 × 256

stations ahead of other discrete designs.

These calculations assume a 3-MHz bandwidth on observations. In these plots improvements in the thermal noise of the telescopes are directly proportional to the bandwidth: a bandwidth of 300 kHz would raise the thermal noise limit by a factor of ten. On the other hand, narrower bandwidths allow the evolution of the EoR signal to be more tightly constrained in time, as the bandwidth covers a smaller redshift range.

With the 10% gain errors assumed here, the total residual noise is the limiting factor to the telescope sensitivity. For the 1024×256 models the total residual noise level is greater than the EoR signal in most instances: the EoR signal is only detectable at $\ell \lesssim 10^3$ at 79 and 95 MHz (the exact limits vary with apodization scheme). The total residual noise is more restrictive for the other discrete models. It is hence imperative that the total residual noise is reduced significantly by constraining the phase and gain errors in SKA1-Low to significantly below the equivalent of the 10% gain error simulated here.

As outlined in Chapter 3, a gain error of 10% creates approximately the same residual sidelobe noise as a phase error of 6°. It is also shown that decreasing the gain or phase error by a factor of 10 results in an approximate decrease in the residual sidelobe noise or residual in-field noise by a factor of 10. If it is assumed that the total residual noise has the same relationship and that a telescope with a gain error of 10% and phase error of 6° has approximately double the total residual noise as a telescope with a gain error of 10% and no phase error, it is possible to estimate the upper limit on the gain and phase error required to prevent the total residual noise exceeding the thermal noise sensitivity of SKA1-Low.

A gain error of $\lesssim 1 \times 10^{-5}$ and a phase error of $\lesssim 3 \times 10^{-4}$ degs would be required for the 1024×256 model without apodization to bring the total residual noise level below the EoR signal for the ℓ range and frequency range over which the EoR signal is stronger than the thermal sensitivity. This is determined by comparing the EoR

signals to the total residual noise. The values are calculated from the factor by which the gain and phase errors must decrease to reduce the total residual noise level to below the thermal noise level of the telescope at all frequencies.

These numbers are exceptionally low and it is perhaps unrealistic to target SKA1-Low to be within these limits. It will be possible to detect the EoR signal in smaller parameter spaces with much larger gain and phase errors (Fig. 5.3 shows it is possible to detect at 79 MHz and 95 MHz with 10% gain error, for example). The EoR signal is most powerful at low ℓ; for the total residual noise to be less than the EoR signal at $\ell = 10^2$ at all frequencies requires a maximum gain error of $\lesssim 1 \times 10^{-3}$ and phase error of $\lesssim 0.05$ degs.

These calculations assume that the gain and phase error on each antenna is static. Over the course of a 10 000-hour integration these errors may vary; this variation would help decrease the total residual noise levels as the variation in error on each antenna is more likely to contribute to an overall mean error of zero. On the other hand, if the errors do indeed remain static for such extended periods of time, it may be possible to calibrate for them to some extent, again decreasing the residual noise error. Any calibration would need to be able to act in buffered real-time, as the data rate out of SKA1-Low will be too great to store. Without such improvements, SKA1-Low will be limited by its total residual noise level and not the thermal noise limit.

The level of the gain and phase errors will dictate the most suitable apodization scheme to use for SKA1-Low. Whilst the total residual noise is the noise floor, strong apodization, such as the Gaussian scheme, is most advantageous choice, as it lowers the total residual noise. On the other hand, should the thermal sensitivity or point-source power spectrum be the limiting factor, weaker apodization should be used. The Hanning window apodization may be the best option as it has the advantage of lowering the total residual noise whilst only mildly increasing the point source and thermal noise levels; in the Gaussian case these rise significantly.

5.5.2 Cluster models

The power spectra pertaining to the cluster telescope models are plotted in Figs. 5.3j to 5.3n. The cluster models outperform the 50×5000 and 200×1250 discrete models in terms of thermal noise limit, total residual noise and point source noise level. The five cluster models tested all give very similar sensitivities to one-another for the various power spectra, but cluster model A marginally gives the best ℓ coverage of thermal sensitivity below the EoR signal.

When assessing thermal noise sensitivity, the cluster models give greater coverage at low ℓ due to the intra-station baselines, which are shorter than all the 1024×256 models' baselines. The thermal noise limit of the cluster models also provide sensitivity below the EoR signal level up to a higher ℓ than the 1024×256 telescope models.

On the other hand, the total residual noise levels with the cluster models are greater than the 1024×256 models; for example the Hanning apodized cluster models' total residual noise are approximately four times greater than the Hanning apodized 1024×256 model. To reduce the total residual noise to below the EoR signal over the ℓ range at which the thermal sensitivity of the telescope is below the EoR signal ($\ell \sim 10^2 - 10^4$), over the redshift range investigated would require gain errors of $\lesssim 5 \times 10^{-6}$ and phase errors of $\lesssim 3 \times 10^{-4}$ degs. Alternatively, if the EoR signal is to be detected at its most powerful ℓ across the entire frequency/redshift range, gain errors of $\lesssim 1 \times 10^{-3}$ and phase errors of $\lesssim 0.07$ degs are required.

The point source noise level for the 1024×256 models is approximately $3 - 41$ times greater than for the cluster models. For both types of model the point-source power spectrum exceeds the thermal noise limit of the telescopes below $\ell \sim 5000$, although the precise level of the point-source power spectrum will depend on the accuracy of spectral smoothing and point source subtraction.

Incorporating all these factors, the optimum choice of telescope model depends on the total residual noise level of the telescope. If gain and phase errors are limited

to below $\lesssim 5 \times 10^{-6}$ and $\lesssim 3 \times 10^{-4}$ degs, respectively, then the cluster model design provides superior angular coverage, with cluster model A performing best. If gain and phase errors cannot be reduced then the 1024×256 model is the better option, with its lower total residual noise level for given gain and phase errors.

Circular stations were used for these calculations and simulations. If elliptical stations were incorporated into these simulations, and the benefits of maintaining a symmetric beam at relatively low observation altitudes accounted for, it may be that cluster models outperform the 1024×256 discrete model with greater gain and/or phase errors.

5.5.3 Maximum baseline length required

The levels of the point-source power spectra in Fig. 5.3 are dependent upon the assumption that the telescope can resolve all the point sources down to five times thermal noise limit of the telescope's long baselines. The better the thermal noise limit of the telescope's long baselines, the greater the number of point sources it will be sensitive enough to detect. By resolving these point sources, their effects can be removed from EoR detections.

By combining the thermal noise limit of the long baselines of each telescope model and Eqs. (5.18) to (5.20) it is possible to determine the maximum baseline length required by each telescope model. These are given in Table 5.1.

The maximum baseline required depends heavily on the telescope model. The better a telescope's thermal noise sensitivity at long baselines, and hence its ability to detect weak point sources, the longer the maximum baseline is required to be, to allow the greater number of point sources to be resolved. The maximum baseline could be decreased for telescope models, but it would not minimise the confusion noise caused by point sources. It is dependent upon the performance of spectral smoothing techniques like Wp smoothing as to how much point source confusion noise can be tolerated by a telescope model before it detracts from the ability to detect the EoR.

Table 5.1: Maximum baseline required, as a function of frequency, for different telescope models to allow point sources to be resolved down to five times the thermal noise limit of the telescope's long baselines.

		Max baseline required, in km, at frequency:					
Model	apodization	65 MHz	79 MHz	95 MHz	110 MHz	119 MHz	130 MHz
50×5000	None	12.2	15.3	19.0	22.7	29.4	36.3
50×5000	Hanning	10.7	13.4	16.6	19.9	25.7	31.7
50×5000	Gaussian	6.8	8.5	10.6	12.6	16.4	20.2
200×1250	None	17.8	22.3	27.7	33.1	42.9	52.9
200×1250	Hanning	15.6	19.6	24.2	29.0	37.6	46.3
200×1250	Gaussian	9.9	12.5	15.4	18.5	23.9	29.5
1024×256	None	28.8	36.3	44.9	53.7	69.6	85.8
1024×256	Hanning	25.4	32.0	39.5	47.3	61.3	75.6
1024×256	Gaussian	16.2	20.4	25.3	30.3	39.2	48.3
Cluster model A	Hanning	29.1	36.6	45.3	54.2	70.3	86.6
Cluster model B	Hanning	25.9	32.6	40.3	48.2	62.5	77.0
Cluster model C	Hanning	25.1	31.6	39.2	46.9	60.7	74.8
Cluster model D	Hanning	25.1	31.6	39.1	46.8	60.6	74.7
Cluster model E	Hanning	26.2	33.0	40.8	48.8	63.2	77.9

Table 5.1 assumes that the differential source counts given in Eq. (5.18) are accurate at very low fluxes. As no telescope has previously been capable of detecting sources at these fluxes it is impossible to assess whether this is indeed the case. If the differential source counts flatten out at low fluxes, for instance if there are a large number of weak galaxy sources, the maximum baseline required could be substantially greater.

5.6 Alternative radial distribution of stations

The power spectra of the 1024×256 discrete telescope models without apodization and with a Hanning window perform best of the telescope models tested. However it may be possible to improve upon this by changing the radial distribution of stations.

The 1024×256 discrete telescope models conformed to the *Requirements specification* that 50% of stations should be within 600 m of the telescope centre and 75% of the stations should be within 1 km. This results in a compact telescope core and a relatively sparse number of mid- and long-range baselines.

It is shown in Fig. 5.3a that the thermal noise sensitivity of the *Requirements*

specification telescope model is best at large angular scales. This is also where the power of the EoR power spectra is greatest. It could be argued that the core is hence over-populated and some short baselines should be sacrificed for longer ones. This could be achieved by locating some core stations at a greater radial distance from the core. This assumes that the noise floor is not limited by the residual noise or point sources.

A further telescope model was created with 1024 stations of 256 antennas each. The radial distances of the stations were distributed so as to be scale-invariant. That is, radial distances up to 50 km were split into logarithmically spaced bins and an equal number of stations placed in each bin. There were limits to this distribution, as only a finite number of stations can physically be located at small radial distances. This telescope model is displayed in Fig. 5.10, alongside its synthesised beam and the density of its *uv*-coverage. Its synthesised beam has an RMS deviation from a Gaussian fit of 5.4×10^{-3} and a mean near-in sidelobe level of ~ -25 dB, compared to ~ -30dB for the 1024×256 model and approximately the same as the cluster model (although the near-in sidelobe levels of the latter are more variable).

Figure 5.11 shows the expected performance of this telescope model. The logarithmic-radial distribution of the stations is shown in the structure of the thermal noise limit. The log-radially binned telescope model gives a thermal noise level below the EoR signal level for a smaller range of ℓ than 1024×256 and cluster models, as there is a reduction in the number of baselines corresponding to $\ell \sim (1-5) \times 10^3$.

This implies that this logarithmically distributed telescope model would not improve upon the 1024×256 or cluster model A designs.

Figure 5.3a does, however, suggest that improvements to the ℓ at which the EoR and thermal noise power spectra cross-over point could be made by moving some stations away from the core to mid-ranges baselines corresponding to $\ell \sim (1-5) \times 10^3$. At 65 MHz this is equivalent to baselines of $\sim 200 - 800$ m.

Moving stations away from the core could be more expensive to implement as

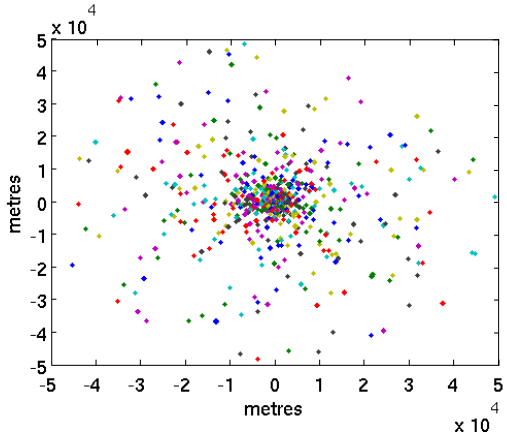

(a) Telescope layout of logarithmically distributed model. Each coloured circle shows the position of a station.

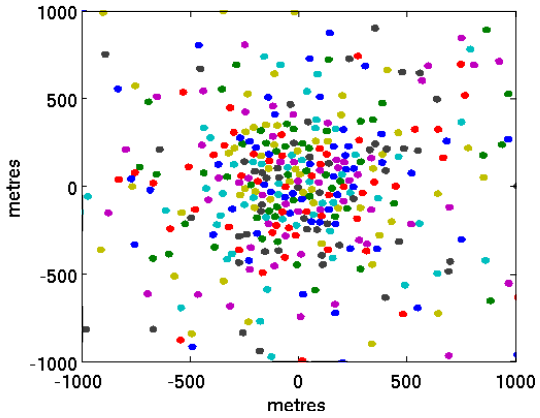

(b) Central area of Fig. 5.10a

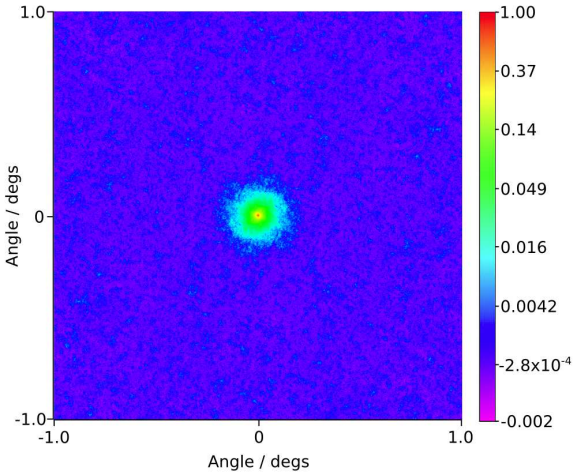

(c) Synthesised beam at 110 MHz

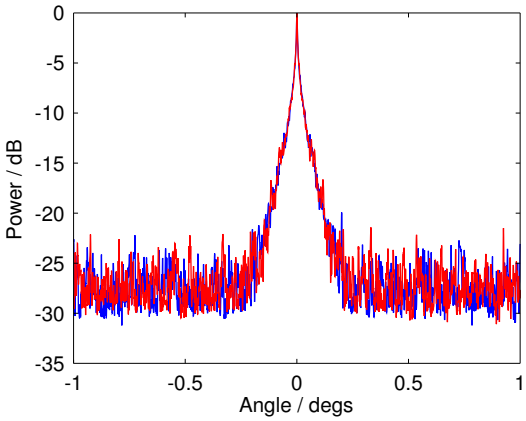

(d) Orthogonal one-dimensional cross-sections of synthesised beam at 110 MHz

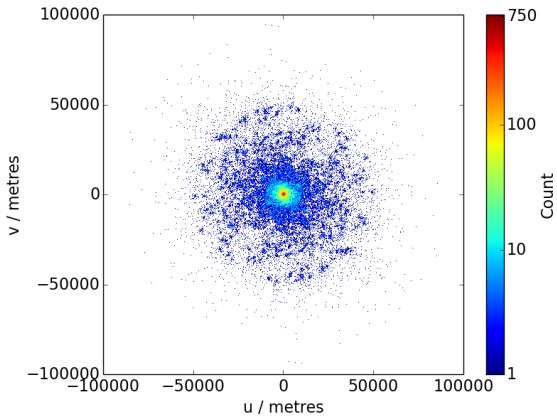

(e) *uv*-coverage of telescope model shown in Fig. 5.10a

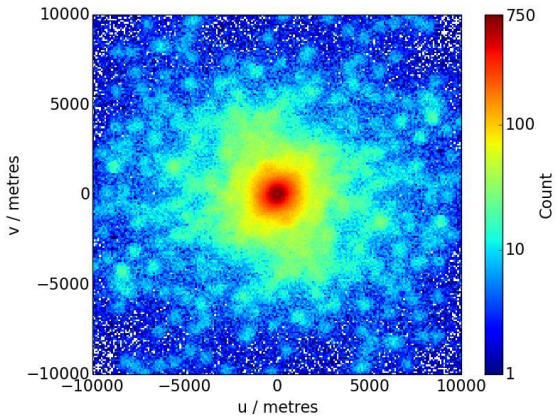

(f) Central area of Fig. 5.10e

Figure 5.10: Telescope layout, synthesised beam and density of *uv*-coverage of logarithmically distributed telescope model.

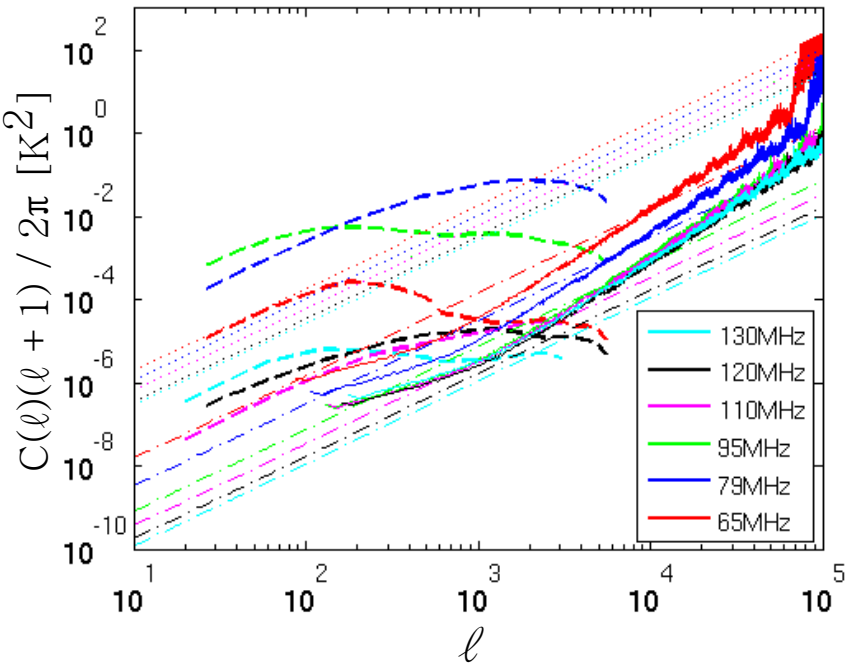

Figure 5.11: Power spectra plots as with Fig. 5.3, except with the logarithmically distributed telescope model shown in Fig. 5.10. Expected EoR power levels are from [85, 133].

the data links between the stations and central processor would need to be longer. The *Requirements specification* currently states that antenna data will be largely processed at a central location. Costs could be saved by beamforming antenna data at the stations; this would reduce the bandwidth of the data transported to the central location housing the correlator. A concern is that increasing the electronics at the antennas could increase the RFI contaminating the data. Some investigations are underway by Peralex Electronics Ltd. into the feasibility of economically shielding the electronics from this RFI [150].

Any significant change in the radial distribution of antennas would need to be constrained to within the land available at the SKA1-Low site, Boolardy station in Western Australia. This would principally limit the angular distribution of large baselines. Conditions of the land at the site (for instance, roads, rocky outcrops, flood zones, emergency runways and heritage areas) will also put localised constraints on

the positions of aperture array, as they have with the MWA, which is located on the same site [24]. An algorithm such as that discussed by Beardsley et al. [151] could be used to optimise placement within the constraints of the physical environment.

5.7 Summary

In conclusion, the current SKA1-Low telescope layout, as given in the *Requirements specification*, is the optimum discrete design for detecting the EoR signal, but cluster model A gives a greater range of angular scales at which the thermal sensitivity of telescope is below the EoR signal level. However, the total residual noise level is an issue and is predicted to be the greatest contribution to the noise level. Attention should be given to minimising the gain and phase errors in the signal processing chain. If gain and phase errors cannot be reduced to extremely low values (gain error $\lesssim 5 \times 10^{-6}$, phase error $\lesssim 3 \times 10^{-4}$ degs) then a strong apodization window should be applied to the current SKA1-Low telescope layout to minimise the impact of total residual noise on EoR observations.

The high frequency end of the EoR observing band ($\gtrsim 110\,\mathrm{MHz}$) will be the most difficult to detect due to the relatively low EoR power spectra at this frequency, relatively high total residual noise and relatively poor thermal noise limit of the telescope. Conversely, it will be easiest to detect the EoR at approximately 79-$-$ 95 MHz.

Apodizing stations in the 1024×256 discrete telescope model (the *Requirements specification* design) will decrease the total residual noise level of EoR observations, which may be the strongest noise contributor. However apodization decreases the thermal noise sensitivity of the telescope.

Apodizing stations too strongly, as in the Gaussian window case, can be detrimental to telescope performance as the thermal noise limit of the telescope increases. The optimal apodization window may be have a tapering somewhere between a Hanning and Gaussian window.

Point sources need to be accounted for in EoR detections, but between resolving them with long baselines and the benefits of spectral smoothing, the confusion noise caused by point sources should be below the thermal noise limit of the telescope.

The maximum baseline length required to resolve point sources depends on the telescope model's ability to detect weak point sources, but is found to be $\sim 86\,\mathrm{km}$ for the current design. This maximum baseline could be decreased at the cost of increasing the point source confusion noise. The ability to do this without impairing telescope performance depends on the success of spectral smoothing techniques.

Moving some stations from the core to mid-range baselines ($\sim 200 - 800\,\mathrm{m}$) might increase the range of angular scales over which the EoR is detectable by SKA1-Low.

Cluster telescope models perform similarly to the current *Requirements specification* design, however their near-in sidelobe level. They could potentially perform better than the *Requirements specification* design when observing at low elevations if elliptical stations were implemented.

Chapter 6

Beamforming hardware for SKA1-Low

6.1 Background

Most early radio telescopes used analogue signal processing, in which astronomical signals are detected at antennas and amplified at the 'front-end' signal processing, before being beamformed and correlated in the 'back-end'. Analogue signal processing acts continuously on signals.

More recently, digital signal processing (DSP) has come to the fore in radio astronomy. Digital signal processing provides greater accuracy, as signals do not degrade along the signal processing chain as readily. Higher rates of data can be processed for the same cost than with analogue processing. It allows for fewer moving parts, hence improving reliability, as the geometrical delays do not require physical changes in lengths of wire. Several digital processing technologies are also reprogrammable, allowing beamformers to be adapted to the science goal of an observation (bandwidth could be sacrificed to produce more beams, for instance). Consequently, several older interferometers, such as the VLA [152], MERLIN (Multi-Element Radio Linked Interferometer Network) [153] and the Northern Cross at Medicina [154] have recently upgraded their back-ends to make greater use of digital signal processing. Current generation telescopes, such as LOFAR and MWA, use a mixture of analogue and digital processing. These telescopes analogue-beamform the antennas in their tiles

before digitally beamforming their tiles [155, 156].

It is preferable to implement the SKA1-Low beamformer with all-digital beamforming. However this is only possible if cost-effective digital processors with sufficient processing speeds and input/output (I/O) data rates are available. Digital processing technology is rapidly evolving; an overview of the current major processors used in radio astronomy is presented below.

This chapter outlines the various popular hardware options available for digital signal processing in radio astronomy. The operations of field programmable gate arrays (FPGAs) and the implementations of beamformers on FPGAs are discussed. Finally, calculations are made as to the number of FPGAs required to operate different beamformer specifications in SKA1-Low.

6.2 Processors

There is a range of options for digital signal processing in radio astronomy, which include CPUs (Central Processing Units), GPUs (Graphics Processing Units), ASICs (Application Specific Integrated Circuits) and FPGAs, each with their own advantages. The choice of which type of digital signal processor to use is not always obvious and is dependent upon the requirement of the application and the current state and cost of each technology, which can changes rapidly.

6.2.1 CPUs

CPUs are the type of digital signal processors most familiar to astronomers, as these are the standard processors used in computers. They allow rapid development of code, as they are flexible, reprogrammable and are programmed by languages used in a wide range of applications. They can have large memories and are relatively cheap, allowing the possibility of frequent upgrades without excessive cost.

Their key drawbacks for radio astronomy digital signal processing are that their signal I/O is typically lower than alternative options and they can require substan-

tially more power to run. Furthermore, CPUs are not optimised for highly parallelised operations with relatively simple algorithms, which are the type of operations used in beamformers and correlators.

LOFAR uses a CPU-based signal processor for its beamformer, an IBM Blue Gene/P supercomputer [157]. The relatively low bandwidth processed with LOFAR at any one time makes it more amenable to CPU processing.

6.2.2 GPUs

Market demand for improved graphical displays, such as in video games and animation, has driven the development of GPUs. GPUs specialise in performing parallelisable vector arithmetic, a valuable feature that has recently begun to be widely harnessed in radio astronomy.

GPUs have higher memory bandwidths, and can perform an order of magnitude more floating point operations per unit cost and per energy cost, than CPUs [158]. Like CPUs, GPUs are cheap, allowing hardware upgrades at low cost and provide a quick and flexible development environment. GPUs are also capable of performing complex algorithms. On the other hand, GPUs have a greater power requirement than FPGAs and ASICs [159]. GPUs also require a developer proficient in a GPU programming language, such as CUDA or OpenCL, which most astronomers are not familiar with and hence could slow down development. Whilst CUDA is specific to Nvidia GPUs, OpenCL is targeted at a broad range of device types, including CPUs, GPUs and FPGAs.

GPUs are currently used in many radio astronomy projects, such as the LEDA correlator for the Owens Valley LWA [79], the MWA [24, 160], the Advanced Radio Transient Event Monitor and Identification System (ARTEMIS) for LOFAR [161] and the Versatile Green Bank Telescope Astronomical Spectrometer (VEGAS) [162].

6.2.3 ASICs

ASICs, as the name implies, are integrated circuits designed to perform an exact task. As a result of their fully customised nature, ASICs are highly efficient in power consumption, have an extremely low unit cost and offer the highest operation and I/O performance levels. In the past ASICs have proved the only digital processor capable of correlating major interferometers, such as the VLA and the Very Long Baseline Array (VLBA) [163, 164]. The advantages that full customisability bring also have a cost: most significantly the non-recurring engineering cost involved in designing and setting up the fabrication of the ASIC [165]. This cost makes ASICs uneconomical for most applications unless a very large number of units are required; this crossover point between ASICs and other processors is constantly shifting and is application-dependent, but can be considered to be of the order one million ASICs when compared to FPGAs. ASICs are designed in a hardware development language (such as VHDL or Verilog). These are dissimilar to most computer programming languages and hence can require a non-trivial time investment to learn. Unlike CPUs, GPUs and FPGAs, ASICs cannot be reprogrammed; this is a disadvantage for SKA1-Low as the ability to remotely change the digital signal processing of the SKA is likely to prove valuable over the lifetime of the telescope.

6.2.4 FPGAs

FPGAs are integrated circuits with programmable connections between the on-chip logic blocks and I/O. This allows generic FPGAs to be programmed to do a wide range of tasks without necessitating a change in hardware. Typically they have faster design cycles than ASICs, as designs can be easily tested. Like ASICs, FPGAs can only process simpler algorithms than CPUs and GPUs, have smaller memory sizes and require a hardware development language to program. However FPGAs can process high data rates, perform high levels of operations and have lower power requirements

than CPUs and GPUs. However their unit costs are higher than GPUs and CPUs, making upgrades less economical [166].

FPGAs have recently found widespread usage in radio astronomy. For instance, a significant fraction of the recently upgraded Very Large Array's (VLA's) correlator, as well as the LWA correlator [79] use the technology. FPGAs are also used in signal processing on PAPER [167], Parkes [168], the C-Band All Sky Survey (C-BASS) telescopes [169] and the MWA [24], amongst others. The rapid development of FPGA capabilities has been encouraged by collaborations such as CASPER (Collaboration for Astronomy Signal Processing and Electronics Research[1]) which develop FPGA platforms and libraries of common features found in radio telescopes and allow FPGAs to be programmed with higher level tools than VHDL and Verilog.

The current state of FPGAs compared to the other processing technologies means they are generally considered to be the optimal medium for SKA1-Low. The low power consumption of FPGAs is one feature that makes them stand out ahead of GPUs, which also have many advocates in radio astronomy. Power consumption is an important consideration, as the scale of the SKA will mean its power costs are a significant overhead, especially as the Australian site is not yet connected to the Australian national grid.

6.2.5 Summary of processors

Each of the four processors outlined is optimum for different project specifications. ASICs have the highest I/O, highest operation rate and are the most power efficient, however their non-recurring engineering costs make them prohibitive for all but the largest product runs, which does not include SKA1-Low. They are also unable to be reprogrammed. FPGAs also have high I/O and operation rates and low power requirements, although not to the extent of ASICs. They are much cheaper for small production runs however. GPUs allow more complicated algorithms, are highly

[1] https://casper.berkeley.edu

parallel and are cheaper than FPGAs, at the expense of lower I/O and operation rates. CPUs run the most complex algorithms and have the cheapest unit cost, but have very low I/O and operation rates with high power consumption.

The progress made by GPUs in terms of power-efficiency in recent years suggests that GPUs may be more viable for future large-scale radio telescopes. Some current generation telescopes, such as the MWA [24] and the LWA, [79] use a mixture of both FPGAs and GPUs.

6.3 FPGA operation

An FPGA is primarily composed of an array of configurable logic blocks. Each logic block can optimally configured to provide a range of logical functions. Memory blocks are also embedded in the FPGA, along with I/O pins and dedicated multiply-accumulate operation blocks known as DSP slices (or DSP blocks).

When a program is synthesised for use on a target FPGA, an iterative process occurs which maps the requisite logic onto the logic cells and attempts to route the signals along the reconfigurable connections. One of the key problems the synthesiser must overcome is meeting the timing constraints; when signals are routed across an FPGA they must be kept synchronised with respect to one another. If signal A is separated into signals B and C at one logic block and, after performing independent operations, signals B and C are to be summed at another logic block, signals B and C must take the same time to travel from the first logic block to the second (within the tolerance of the clock cycle).

Timing differences can be overcome by changing the blocks used on the FPGA fabric to those located in different positions or by adding delays to intentionally retard signal propagation, in a process known as pipelining, which uses memory on the FPGA to buffer signals. The greater the utilisation percentage of blocks on an FPGA, and the greater the FPGA's clock speed, the harder it is to meet these timing constraints.

Each generation of FPGAs has a greater number of blocks than the previous, but improved semiconductor device fabrication technology allows the distance between blocks to decrease, hence there is not necessarily an increase in signal propagation time across the FPGA.

The clock is a regular pulse fed to each component of the FPGA design, providing accurate timing on when each component is to perform its operation. It is distributed along dedicated clock global routing lines on the FPGA fabric, which minimise the difference in arrival time of the clock across the FPGA.

6.3.1 Channelisation

Beamforming in aperture arrays can be implemented in either the time or frequency domains. This chapter will focus on the latter: phase-shift beamforming. This is the current proposed implementation for SKA1-Low [39]. Beamforming requires specified delays, but it is easier to implement phase-shifts. To accurately perform phase-shift beamforming the antenna signals must be divided into narrow frequency channels, as delays are frequency dependent. The narrower the frequency channel, the more accurate the beamforming operation; errors in the beamforming occur most prominently at the edges of the frequency channels, as phase-shift values are calculated for the centre of the band.

Channelisation is implemented by means of a Fourier transform. In digital signal processing signals are discrete, hence one option is to use a discrete Fourier transform. The discrete Fourier transform is generally not practical to implement in digital signal processing, unless the efficient Fast Fourier Transform (FFT) algorithm is utilised. Although originally discovered by Carl Gauss in the early 19th century [170], the algorithm was popularised after it was independently rediscovered by Cooley and Tukey in 1965 [171]. The complex discrete Fourier transform requires $\mathcal{O}(N^2)$ multiplications, whereas the FFT requires $\mathcal{O}(\frac{N}{2}\log_2 N)$. The FFT is mathematically identical to the discrete Fourier transform.

The finite length of a Fourier transform is equivalent to putting a rectangular window function around an infinitely long time signal. Hence the Fourier transform of a sinusoidal time signal is equivalent to the convolution of the Fourier transform of the time signal and the Fourier transform of a window function; the former is a delta function and the latter a sinc function. This means the single frequency sine wave is present to some degree in every frequency bin. This is called spectral leakage. This can be especially problematic with RFI, which, as it is typically orders of magnitude stronger than astronomical signals, can necessitate data from a large number of frequency bins being discarded.

A second issue is that the amplitude response across the principal frequency bin is non-uniform. Consequently, two sine waves with different frequencies and of equal amplitude in the time-domain may not produce signals of equal amplitude in the frequency domain.

The Polyphase Filter Bank (PFB) technique is an approach to reduce the impact of these effects, albeit at the cost of increased computational complexity. A PFB flattens the frequency response of each channel and suppresses its sidelobes [172, 173], and consists of the combination of an FFT with a finite impulse response (FIR) filter.

If an N-point Fourier transform is desired, a PFB requires a time-domain input of length PN, where P is the number of taps in the FIR filter. The greater the number of taps, the better the improvement in the shape of the channel band. This time-domain signal is weighted by a sinc function by multiplying point-by-point. Just as the rectangular window function in a FFT results in the ringing sinc function in the Fourier transform domain, the sinc function now results in a rectangular function to sharply cut off the frequency bin responses. Once weighted, the signal is separated in to P signals of length N and the signals summed point-by-point. The application of a regular FFT will now result in sharper, flatter frequency bin responses, as shown in Fig. 6.1.

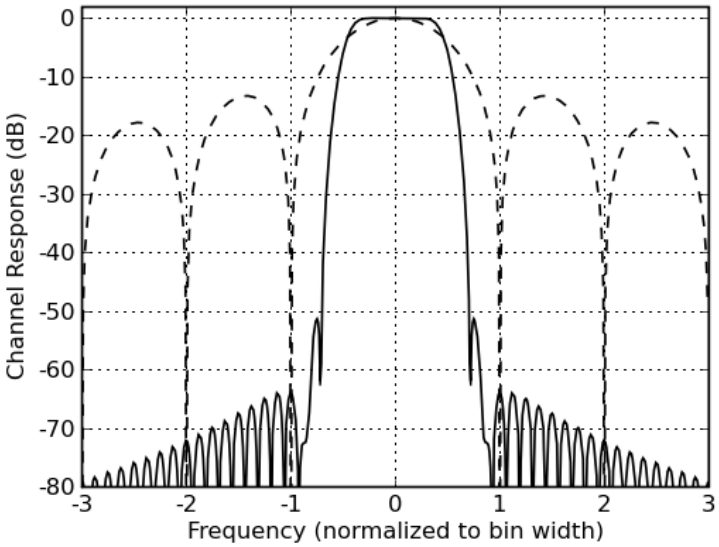

Figure 6.1: The frequency channel response of an FFT (dashed line) and a Polyphase Filter Bank (PFB) FFT (solid line), with an 8-tap filter is used. The PFB-FFT flattens the main channel response and significantly attenuates the response to other frequency bins, compared to the FFT alone. Image from [167].

6.3.2 Phase-shift beamforming operation

Once channelised, phase-shift beamforming is implemented by the multiplication of

$$\mathbf{a} \cdot \mathbf{w} = \mathbf{b}\,, \tag{6.1}$$

where **a** is a vector that gives the voltage detected at each antennas, for every frequency channel, for each polarization and at every clock sample, **b** is a vector that gives the voltages in each beam created for every frequency channel at each time sample and **w** is the weighting matrix, with dimensions (**a**, **b**). The entries of **w** are complex and depend on the antenna layout and the desired observing directions.

For a 2-D planar array

$$\mathbf{w} = \mathrm{e}^{-ik\mathbf{r}_i \cdot \hat{\mathbf{r}}_0}\,, \tag{6.2}$$

where

$$k = \frac{2\pi}{\lambda}\,, \tag{6.3}$$

is the wavenumber, $\mathbf{r}_i$ is the position of the i^{th} antenna (relative to the centre of the coordinate system), given by

$$\mathbf{r}_i = \mathbf{x}_i\,\hat{x} + \mathbf{y}_i\,\hat{y} + \mathbf{z}_i\,\hat{z}\,, \tag{6.4}$$

and $\hat{r}$ is a unit vector in the observing direction (θ, ϕ),

$$\hat{\mathbf{r}}_0 = \sin(\theta)\cos(\phi)\,\hat{x} + \sin(\theta)\sin(\phi)\,\hat{y} + \cos(\theta)\,\hat{z}\,. \tag{6.5}$$

If equal amplitudes are applied to the entries of $\mathbf{w}$ the result is a flat illumination. By modifying the amplitudes the beam shape can be varied. This can be used to generate nulls in the directions of RFI sources as well as point the beam

6.3.3 Multiply-Accumulate Operations

Calculating Eq. (6.1), or any other matrix multiplication, requires repeated multiplications and summations. This is a base unit of operation in digital signal processing, known aptly as a multiply-accumulate (MAC). Most recent FPGAs contain DSP blocks, which are blocks optimised to perform MAC operations. The 'Xtreme DSP 48E' slice for Xilinx's Virtex-5 FPGAs, for instance, contains a 25×18-bit multiplier and a 3-input 48-bit adder. Other DSP slices also contain pre-adders. It is possible to implement MACs in the logic fabric of an FPGA, external to DSP blocks, however DSP blocks allow for greater clock speeds and reduced power consumption, both of which are important for an application such as the SKA. Fabric implemented MAC operations also require substantially more on-chip resources that could be used elsewhere in the design.

The number of MACs per second, $N_{\text{MAC s}^{-1}}$, an FPGA is capable of performing is given by

$$N_{\text{MACs s}^{-1}} = N_{\text{DSP}}\,S_{\text{CLK}}\,, \tag{6.6}$$

where N_{DSP} is the number of DSP slices on the FPGA and S_{CLK} is the clock speed. In reality, this can be considered to be an upper limit on $N_{\text{MAC s}^{-1}}$, as it is near

impossible to fully utilise components on an FPGA, especially at high clock speeds. This is related to the time constraint problem; the greater the utilisation of a type of block on the FPGA, the less flexibility the synthesiser has to find a routing that will meet all timing demands.

The number of MAC operations required for a generic vector-by-matrix multiplication is given by the product of the matrix dimensions. However as **a** and **w** are complex numbers, it is the number of complex MACs (cMACs) that must be calculated by the FPGA which is pertinent to Eq. (6.1).

There is a range of options for implementing a cMAC operation. The simplest approach consists of the straight forward mathematical operation to get the real, P_{Re}, and imaginary, P_{Im}, products:

$$\begin{aligned} P_{\mathrm{Co}} &= (a_{\mathrm{Re}} + ia_{\mathrm{Im}})(b_{\mathrm{Re}} + ib_{\mathrm{Im}})\,, \\ P_{\mathrm{Co}} &= a_{\mathrm{Re}}b_{\mathrm{Re}} + ia_{\mathrm{Re}}b_{\mathrm{Im}} + ia_{\mathrm{Im}} + b_{\mathrm{Re}} + i^2 a_{\mathrm{Im}}b_{\mathrm{Im}}\,, \\ P_{\mathrm{Re}} &= a_{\mathrm{Re}}b_{\mathrm{Re}} - a_{\mathrm{Im}}b_{\mathrm{Im}}\,, \\ P_{\mathrm{Im}} &= a_{\mathrm{Re}}b_{\mathrm{Im}} + a_{\mathrm{Im}}b_{\mathrm{Re}}\,, \end{aligned} \tag{6.7}$$

where P_{Co} is the complex product. This requires four MACs. Eq. (6.7) can be rearranged to require only three unique multiplications,

$$\begin{aligned} P_{\mathrm{Re}} &= (a_{\mathrm{Re}} - a_{\mathrm{Im}})b_{\mathrm{Im}} + (b_{\mathrm{Re}} - b_{\mathrm{Im}})a_{\mathrm{Re}}\,, \\ P_{\mathrm{Re}} &= (a_{\mathrm{Re}} - a_{\mathrm{Im}})b_{\mathrm{Im}} + (b_{\mathrm{Re}} + b_{\mathrm{Im}})a_{\mathrm{Im}}\,. \end{aligned} \tag{6.8}$$

If the DSP slices on the FPGA have a pre-adder, a cMAC can hence be implemented with three DSP slices [174]. This is illustrated in Fig. 6.2.

A cMAC can be implemented on single DSP slice if bit depth is sacrificed. 3-bit cMACs are possible on a single DSP slice with an 18×18-bit multiplier [167] and Hickish discusses a method of completing a 4-bit cMAC on a single DSP slice with an additional subtraction operation that is implemented in the fabric of the FPGA [175]. Hickish's approach uses unsigned 4-bit complex numbers which are zero-padded to allow the multiplication of different 4-bit numbers to be contained within a single 36-bit number.

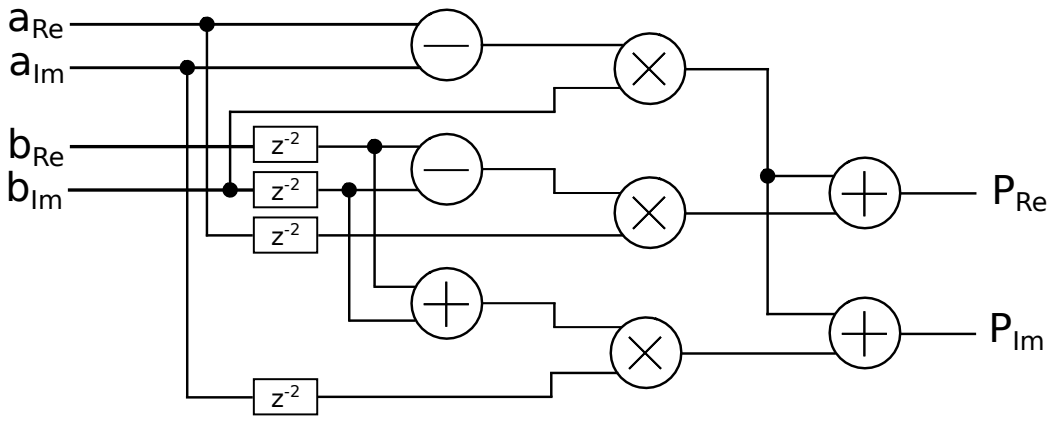

Figure 6.2: Implementation of a Complex-Multiplication operation with three multiplications, as described by Eq. (6.8). Image from [174].

6.4 Beamforming approaches

The current design for SKA1-Low states that it will comprise 262 144 antennas grouped into 256 geographically localised stations. Each antenna is set to sample 8-bits at $\sim 1\,\mathrm{GSa\,s^{-1}}$ per polarization, hence the data rate that the first beamformer stage will be $\sim 16\,\mathrm{GSa\,s^{-1}}$, a significant amount. It would be cost prohibitive to connect every one of the antennas to one another for the purpose of processing their signals. The solution is to use a tiered approach where the signals from small groups of antennas are beamformed before being passed further along the signal processing chain [39].

6.4.1 Hierarchical beamforming

One approach of making this task more manageable is to use hierarchical beamforming to quickly reduce the combined data rate. This can be achieved by initially beamforming the signals of a small number of antennas. This small group of antennas is known as a tile and typically comprises a low power-of-2 number of antennas, such as 16 or 32; the power-of-2 will assist efficient signal processing. As the antennas comprising a tile are typically close together relative to the scale of a station, they create a (relatively) wide angle tile beam on the sky when beamformed.

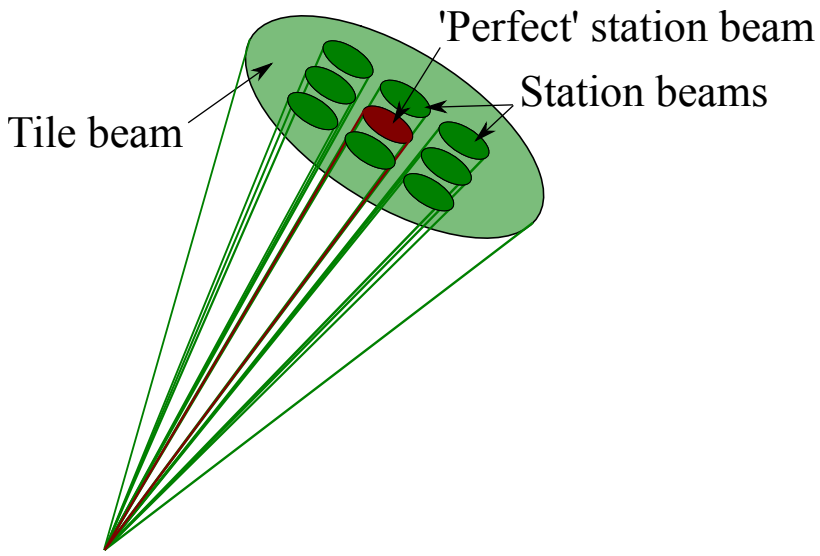

Figure 6.3: Station beams within a larger tile beam, as occurs with hierarchical beamforming. Only the central station beam is looking in precisely the same direction as the tile beam; the other station beams are degraded dependent on their offset from the tile beam direction. Image design from [39].

The beams from each tile in a station are then beamformed together to create one or more station beams, as illustrated in Fig. 6.3. A station beam formed in the direction of the tile beam's phase centre will perform well, however the drawback to hierarchical beamforming arises when station beams are steered in a direction offset from a tile beam's phase centre. As a consequence of creating a beam in a direction for which previous time-delays/phase-shifts have not been set for, the resultant beam has more powerful sidelobes. The greater the offset between the tile beam and station beam phase centres, the greater the impact of these effects [39].

An illustration of the system architecture for this a hierarchical beamformer is shown in Fig. 6.4.

6.4.2 Partial beamforming

An alternative multilevel beamforming approach is known as partial beamforming. In this method, instead of creating a single wide-angle tile beam at the first beamforming level, one beam is created in each tile for every station-level beam desired. Each tile-level partial beam is steered to point in the direction of a station beam's phase

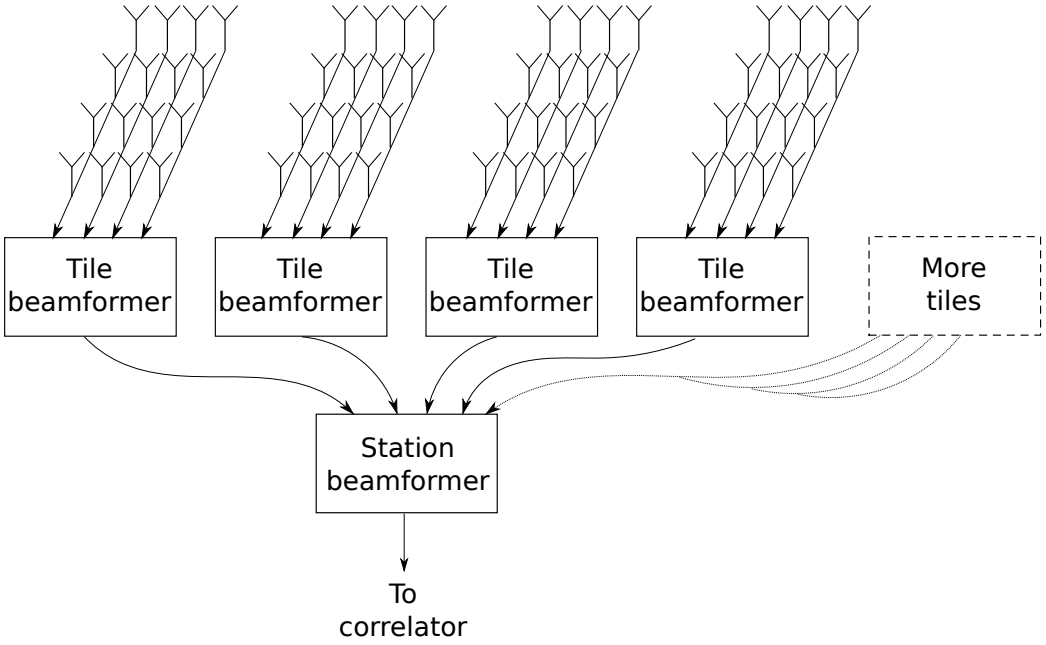

Figure 6.4: Illustration of an example beamformer architecture. The signals from 16 antennas are beamformed in each tile beamformer, before the tile signals are beamformed at the station beamformer. The tile beamformers also contain the channelisers.

centre. Hence when the station beams are formed from the tile beams, the tile beams will have the correct time-delays/phase-shifts, so the gain and sidelobe problems of hierarchical beamforming are prevented. Beams from each tile (all with the same phase centre) are summed together at the station-level to create the station beam. Partial beamforming has been successfully implemented in small scale projects such as 2-PAD [93, 176].

The disadvantage to partial beamforming is that it does not reduce the data rate through the signal processors as quickly as hierarchical beamformers. In fact, should the number of station beams required exceed the number of antennas inputting to the tile-level beamformer, a greater data rate will be output from the tile-level beamformer than is input to it. Practically speaking, this sets a limit on the minimum size of tile that must be produced, such that the number of antennas per tile significantly exceeds the number of beams.

Figure 6.5 illustrates the difference between partial and hierarchical beamforming

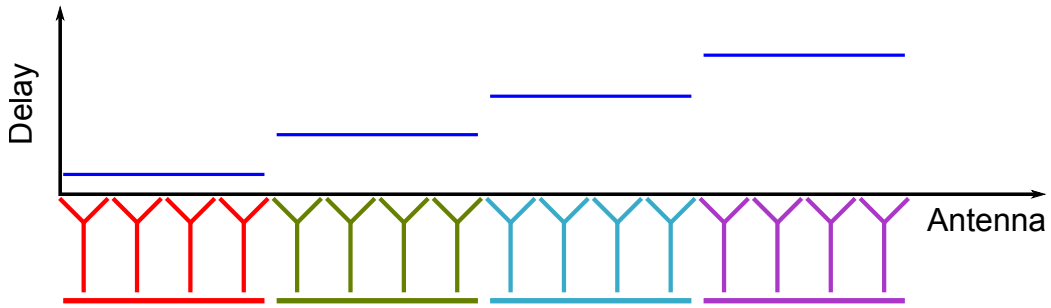

(a) Hierarchical beamforming. Tiles are grouped by colour and underlined.

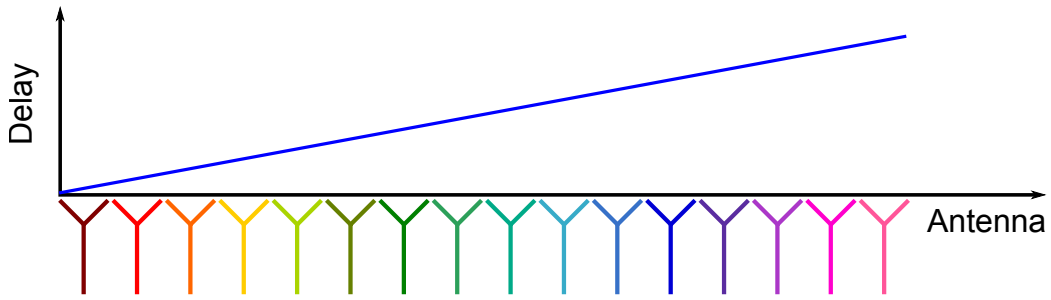

(b) Partial beamforming

Figure 6.5: A comparison of hierarchical and partial beamforming. In partial beamforming the station-level beamforming delay is tuned to each antenna's specific location. In hierarchical beamforming small collections of antennas are grouped tiles. The antennas are initially beamformed at the tile stage, before the tile beams are beamformed at the station-level. Whilst reducing the data rate that must be processed, the delays at the station-level beamformer are tuned to the mean location the antennas in each tile, which increases the sidelobes in the beam.

in the phases of antenna signals at the station-level beamformer.

6.4.3 FFT beamforming

A third approach is to dispense with conventional beamforming and take a 2-D FFT of the signal from each antenna. This takes advantage of the computational efficiencies of an FFT [177]. FFT beamforming allows the whole sky to be imaged simultaneously and the result is equivalent to performing partial beamforming over the whole hemisphere with a regularly spaced grid of beams. A 2-D FFT requires samples to be on a grid, so either antennas must be placed with a gridded layout, leading to both redundant baselines and grating lobes, or samples must be computationally interpolated

on to a grid, which can be computationally intensive, especially for a large number of antennas, which makes it unlikely to be implemented in SKA1-Low. Telescopes such as the MWA [24] and LWA [79] have digital signal processors allowing FFT beamforming. A detailed comparison between FFT and conventional beamforming can be found in Hickish 2013 [175].

6.5 Beamformer resource cost

There are many possible implementations of the SKA1-Low beamformer, each with their own costs. It is useful to be able to calculate how many FPGAs would be required to implement a beamformer for any given telescope specification.

6.5.1 Resource use equations

The resources used on an FPGA for a beamformer comprise principally the DSP slices for MACs and transceivers for high-speed data input and output. There is also a memory requirement to store the beamforming weights use in Eq. (6.1), which can be held in Block Random-Access Memory (BRAM) and Distributed Random-Access Memory (DRAM).

The number of cMACs per second required to channelise data for the beamformer with a PFB-FFT, $N_{\text{cMACs, ch}}$, is given by [175],

$$N_{\text{cMACs, ch}} = \Delta\nu_{\text{ch}} N_{\text{A}} N_{\text{P}} (\log_2(N_{\text{c}}) + N_{\text{taps}}) \,, \tag{6.9}$$

where $\Delta\nu_{\text{ch}}$ is the bandwidth of the channeliser, N_{A} the number of antennas, N_{c} the number of frequency channels the signal is split into and N_{taps} is the number of taps in the PFB-FFT. The bandwidth and number of channels at the channeliser can be greater than in the rest of the DSP chain. This is because a wider bandwidth is sampled to reduce aliasing; unwanted channels are discarded after the channeliser.

The memory required on an FPGA for storing the beamformer weights, M_{weights}, is given by,

$$M_{\text{weights}} = N_{\text{A}} N_{\text{P}} N_{\text{cu}} N_{\text{bitsw}} \,, \tag{6.10}$$

where N_{cu} is the number of channels used in the beamformer after channelisation. N_{bitsw} is the number of bits required to store the beamforming weights. Memory is stored in the BRAM and DRAM on Xilinx FPGAs. BRAMs are dedicated memory blocks on the FPGA, whereas DRAMs can be flexibly made from logic cells in the FPGA fabric to create smaller memory storage.

The required data rate out of an FPGA, D_{ro}, is determined by,

$$D_{\mathrm{ro}} = R\, N_{\mathrm{bits}} N_{\mathrm{P}} N_{\mathrm{B}}\,, \tag{6.11}$$

where R is the sampling rate of the FPGA, N_{bits} is the bit depth of the antenna signals and N_{B} is the number of beams created from the input signals. The maximum data rate an FPGA can actually output is limited by the number, N_{Tx} and speed, S_{Tx} of the transceivers, and how many of those transceivers are required for data input to the FPGA.

6.5.2 Beamformer specifications

The *Requirements specification* of SKA1-Low dictates that a bandwidth of 300 MHz is to be processed simultaneously, with at least 8-bit precision. There will be 256 antennas per station, with two polarizations sampled.

Further assumptions must be made to calculate the required number of FPGAs. For this purpose, it is assumed that a PFB-FFT will be utilised in the channeliser; the *Requirements specification* of ≤ -30 dB spectral leakage between adjacent channels and ≤ -60 dB between non-adjacent channels will be more readily achieved with the strong out-of-band suppression of spectral windows of a PFB-FFT. It is also assumed that the channeliser will operate from $0 - 500$ MHz and once channelised, the data outside the $50 - 350$ MHz range will be discarded. It is further assumed that the DSP blocks within the FPGA will operate at 300 MHz; the maximum clock speed of the FPGAs used in these calculations can be greater than this, but it can be difficult to meet timing constraints at these high speeds. Finally, it is taken that each cMAC

Table 6.1: Specifications of three FPGAs. The Virtex UltraScale specifications are from [178], the Virtex-7 specifications from [179] and the Virtex-6 specifications from [180].

FPGA	$\mathbf{N_{Tx}}$	$\mathbf{S_{Tx}/Gbs^{-1}}$	$\mathbf{N_{DSP}}$	**BRAM/Mb**	**DRAM / Mb**
Virtex UltraScale+: XCVU13P	128	32.75	11904	526.5	46.4
Virtex-7: XC7VX690T	80	13.1	3600	52.920	10.888
Virtex-6: XC6VSX475T	36	6.6	2016	38.304	7.640

required 3 MACs, corresponding to Eq. (6.8), as cMACs performed on a single DSP block would not retain the bit depth required.

Code using Eqs. (6.6) and (6.9) to (6.11) and the specifications of an FPGA and a beamformer determines how many of that type of FPGA would be required to perform the channelising and beamforming operations. The code calculates how many FPGAs are required to perform the tile beamforming and the station beamforming. It also determines whether the limiting factor for a given FPGA and beamformer is the number of MAC operations required, the I/O data rate or the memory on the FPGA.

6.5.3 FPGAs for calculations

Three FPGAs were trialled for beamformer implementations of different SKA1-Low specifications. The first is the Xilinx UltraScale+ XCVU13P, one of Xilinx's latest generation of FPGAs. The XCVU13P has the greatest number of DSP slices of Xilinx's FPGAs and the highest number of high-speed transceivers. The second FPGA considered is the Xilinx Virtex-7 XC7VX690T, a high-end FPGA from the previous FPGA generation. The final option is the Xilinx Virtex-6 XC6VSX475T, an older generation FPGA. This FPGA is well used in the radio astronomy instrumentation community, as it is the FPGA used in the CASPER Roach II boards which are used on telescopes such as the LWA [79]. The pertinent specifications of the FPGAs are listed in Table 6.1.

The advance in performance between each generation of FPGA is evident in Table 6.1. However, the unit cost of the Virtex UltraScale will be significantly more than the Virtex-7 (when the price is released), which in turn is significantly more expensive

than the Virtex-6. On the other hand, the manufacturer claims that there is a power reduction per unit performance for each successive FPGA generation [181, 182], which would reduce costs.

There is an alternative supplier of FPGAs, Altera, but only Xilinx FPGAs are considered in this analysis. The latest three generations of Altera FPGAs can be considered to be approximately equivalent to these three generations of Xilinx FPGAs.

Calculations are made both assuming the potential for 100% block utilisation of the FPGA and a more realistic assumption of a maximum 75% utilisation of DSP blocks and memory for channelising and beamforming operations (not accounting for memory used in pipelining). A fraction of the transceiver data rate will be required to input and output overheads such as packetisation, hence calculations are made with 100% and 90% transceiver data rates available for the beamformer. The utilisation rates of realistic FPGAs are approximately guided by the utilisation efficiency of the LEDA correlator [79], but the actual values achieved for SKA1-Low may be markedly different.

Calculations are also made to determine the number of FPGAs required if 10 beams are created per station instead of the single beam currently specified for SKA1-Low.

Furthermore, the number of channels in the beamformer and the bit depth of the antenna signals were varied. The number of frequency channels to create in the channeliser is varied between 2^{10} and 2^{16}. For comparison, according to Eq. (6.10), the Random-Access Memory (RAM) available on a Virtex-7, at 75% utilisation, is capable of storing 16-bit beamforming weights of 16 antennas for 2^{15} frequency channels.

The bit depth of antenna signals was considered at 8-bits and 16-bits. 8-bits is the minimum bit depth of the signals according to the *Requirements specification*. The maximum multiplication possible on a single DSP slice is 25×18-bits, so 16-bits is considered a reasonable upper limit.

Only factors of 256 (the number of antennas per station) were considered for

the number of antennas per tile, to ensure the same firmware and hardware could be replicated for each tile when deploying SKA1-Low. Variations in firmware and hardware would increase costs and complexity.

6.5.4 Number of FPGAs required

As shown in Fig. 6.6, substantially more Virtex-6 FPGAs would be required than Virtex-7s, but Virtex UltraScales unsurprisingly need the fewest FPGAs, with some telescope specifications requiring only five per station.

A primary cause of the differences is the number of transceivers on the respective FPGAs. A single antenna will output $1\,\mathrm{GSa\,s^{-1}}$ per polarization. If this is sampled at 8-bits, a data rate of 8 Gbps is created, per polarization. The transceivers on the Virtex-6s are only capable of processing 6.6 Gbps, hence more than one transceiver would be required to process each polarization of each antenna. This limits the Virtex-6s to only being able to beamform a maximum of 8 antennas' worth of data in a tile, whereas the UltraScale can beamform 64 antennas in some instances at 100% utilisation. Demultiplexing the signal between the antennas and the Virtex-6s, to enable signals to be split between multiple transceivers, would also require extra hardware. Likewise, the speeds of the UltraScale transceivers, 32.75 Gbps, means they can process data from multiple polarization signals simultaneously, such that 64-antenna tiles could be created, but additional hardware would be needed to multiplex these signals together.

The number of transceivers is also limiting by the FPGAs' ability to output beamformed data, especially when ten beams are created. In these instances, antenna signals would need to be duplicated, with extra hardware, to be fed into multiple FPGAs.

$N_{\mathrm{MACs\,s^{-1}}}$ is a further limiting factor in how many antennas' signals could be channelised and beamformed at tiles, especially with the UltraScale FPGAs.

Varying the number of channels also affects how many FPGAs are required, but

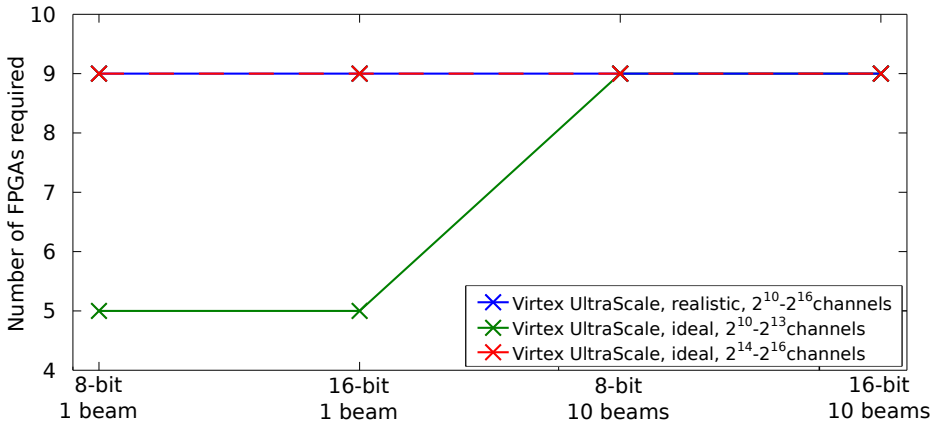

(a) Virtex UltraScale+ XCVU13P

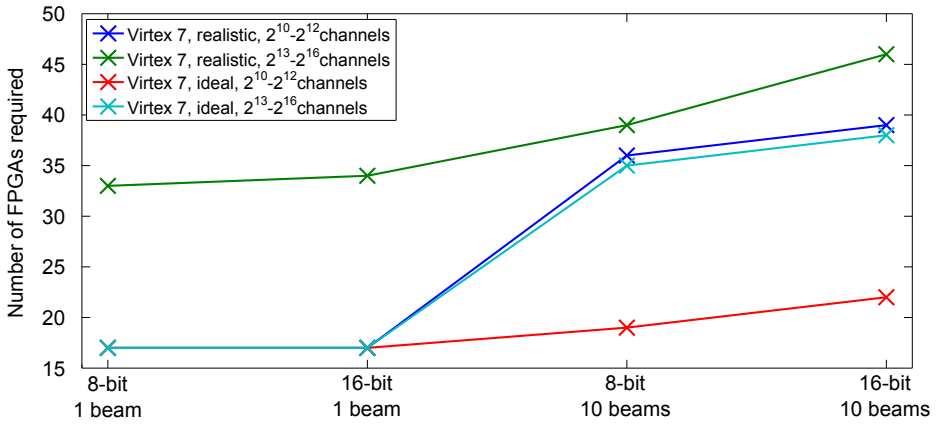

(b) Virtex-7 XC7VX690T

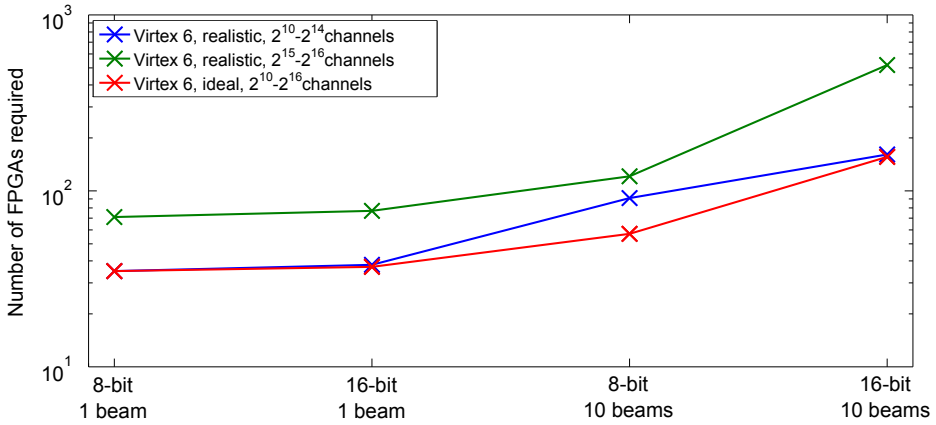

(c) Virtex-6 XC6VSX475T

Figure 6.6: Number of FPGAs required to beamform a 256-antenna SKA1-Low station. The number of frequency channels is varied along with the bit depth of the antenna signals. 'Ideal' FPGAs are considered as FPGAs from which 100% utilisation of resources can be achieved, 'realistic' FPGAs utilise a maximum of 75% of DSP blocks and memory and allocate 10% of transceiver bandwidth to overheads. A clock speed of 300 MHz is assumed.

only at one specific incremental point for each FPGA. For instance, with 'realistic' Virtex-7s there is no difference in FPGA numbers between $2^{10} - 2^{12}$ channels and $2^{13} - 2^{16}$ channels, but $2^{13} - 2^{16}$ channels require more than $2^{10} - 2^{12}$ channels.

These results suggest that for a 1 or 10 beam SKA1-Low, an 8-bit or 16-bit beamformer could be made with a single Virtex UltraScale (or alternative FPGA with equivalent specifications) per 32-antenna tile and a single Virtex UltraScale as the station-level beamformer: a total of 9 FPGAs per station. Although the unit cost of a Virtex UltraScale will exceed that of a Virtex-7 or Virtex-6, the reduction in FPGAs required, together with the cost savings of fewer connecting cables and (manufacturer claimed) low power recommends that Virtex UltraScale FPGAs be used. With fewer FPGAs it may also prove easier trace hardware faults, although the replacement costs would be higher.

Constructing the beamformer with fewer FPGAs also increases the potential versatility of the beamformer. For example, bandwidth could be sacrificed to allow a greater number of beams to be processed; a large number of FPGAs each connected to fewer antennas potentially places restrictions on this flexibility.

As noted above, it is possible to channelise and beamform 64-antenna tiles with Virtex UltraScale FPGAs in some circumstances, hence requiring only 5 FPGAs per station. Specifically these are with 100% utilisation of the FPGAs, a maximum of 2^{13} frequency channels and a single beam; 8 or 16-bit signals are both possible.

Further calculations find that if 79% of DSP blocks are utilised, it would be possible to channelise and beamform 64-antenna tiles with 2^{10} frequency channels and 8 or 16-bit signals. Alternatively a faster clock speed would allow for a greater $N_{\mathrm{MACs\,s^{-1}}}$ and hence achieve the same results. Another option is to implement some MACs outside of DSP blocks, but this would be difficult at high clock speeds. Ergo, if UltraScale (or equivalent) FPGAs are used, careful consideration should be made as to whether it is possible to process 64-antenna tiles, as this would be a significant cost saving.

This analysis has only considered three FPGAs with significant differences in specifications. Further investigations should be made of high-performance FPGAs to determine which specific FPGA is the most appropriate choice for SKA1-Low.

Design work for the SKA1-Low beamformer is underway with the latest-generation FPGAs from Xilinx and Altera [150].

6.6 Summary

In this chapter a range of digital signal processor hardware options are outlined. FPGAs are suggested as the best signal processing hardware for SKA1-Low due to their high data rates, low power requirements and moderate non-recurring engineering costs. The beamforming operation is outlined and its application to FPGAs discussed.

It was calculated that the beamformer for SKA1-Low could be implemented with nine Virtex UltraScale+ XCVU13P FPGAs, for one beam and $\leq 2^{13}$ channels in the beamformer, assuming 75% utilisation efficiencies of the FPGA, 90% utilisation of the high-speed transceivers on the FPGA and a clock speed of 300 MHz. This could be reduced to five FPGAs if 79% utilisation efficiency was achieved and $\leq 2^{10}$ channels used. A greater number of less powerful FPGAs could also be used, but with additional costs to power requirement, rack space and connecting cables. The latest generation of FPGAs reduces the station beamformer task from a technically challenging design problem to a far more tractable one.

Chapter 7

Conclusion

7.1 Thesis summary

This thesis has investigated SKA1-Low, the low-frequency aperture array of the Square Kilometre Array (SKA).

It began with an introduction to the field of radio astronomy and aperture array telescopes, in which the operation of a radio interferometer and beamformer were explained. It discussed the SKA, with a focus on its low-frequency aperture array and its pathfinder and precursor telescopes, and touched on their key science cases.

Chapter two investigated the station beam of low-frequency aperture arrays. It outlined how antenna positioning affects the station beam and simulated the effect of the number of antennas, phase errors and broken antennas on station beams. It was shown that phase errors should be kept within $\sim 3.5°$ and $\sim 6°$ for a station with 16 and ≥ 128 antennas, respectively, to keep beam sensitivity and pointing offset to within an uncertainty of 1%. The number of broken antennas also affected these parameters, especially for small numbers of broken antennas.

Chapter three analysed station beams simulating how phase and gain errors on the signal processing path would impact SKA1-Low's imaging. It was shown how this could limit the minimum noise level achievable with the telescope. It also highlighted that the large beam sizes at low frequencies will lead to higher residual noise levels at lower frequencies, where residual noise is the noise introduced to images from gain

and phase errors in the signal processing. It was found that the systematic noise limit due to these effects on SKA1-Low images will range between $\sim 10 - 0.3\,\mu$Jy and will be reached after integrations of $\sim 100 - 100\,000$ hours, for observations from $50 - 600$ MHz.

Chapter four discussed the flexible station concept and its potential benefits. A range of telescope models conforming to the SKA specifications of number of antennas and their radial distribution were created, in both traditional 'discrete' station designs and the flexible concept. The synthesised beams and the *uv*-coverage were analysed. Synthesised beams from station models with 1024 stations with 256 antennas per station had the most Gaussian-shaped beams, owing to the large number of stations. Flexible starfish-shaped telescope models created beams with the least Gaussian beams and were considered unsuitable for SKA1-Low.

These telescope models were then assessed on their performance for measuring the EoR power spectra and how detrimental effects of the residual noise and foreground point source would impact observations. It was determined that residual noise in images could significantly restrict SKA1-Low's ability to detect the EoR signal. Unless SKA1-Low has extremely low gain and phase errors (gain error $\lesssim 5 \times 10^{-6}$, phase error $\lesssim 3 \times 10^{-5}$ degs) a strong apodization window should be applied to the current SKA1-Low telescope layout to minimise the impact of total residual noise on EoR observations. If these gain and phase errors can be achieved, a flexible cluster telescope model, in which antennas are shared between stations, would be a superior telescope layout owing to its better thermal noise sensitivity.

Foreground point sources could also hinder EoR observations, and accurate spectral smoothing and point source subtraction techniques will be required to prevent this. The maximum baseline length required for each telescope model to resolve the point sources to be subtracted from images was estimated to be ~ 86 km for the current SKA1-Low design.

It was determined that the EoR will be easiest to detect between $79 - 95$ MHz and

most difficult above 110 MHz. Discarding the radial distribution specification of 50% stations within 600 m and 75% stations with 1 km of the core and instead moving some stations from the core to mid-range baselines ($\sim 200 - 800$ m), may increase the range of angular scales over which the EoR is detectable by SKA1-Low.

Chapter six discussed the digital signal processing operation for beamforming an aperture array and the hardware options. It calculated the number of FPGAs required to implement a beamformer on an SKA1-Low station of 256 dual-polarization antennas with different values for the number of beamforming channels and number of beams with different FPGAs. Using FPGAs of specifications comparable to the Virtex UltraScale+ XCVU13P FPGAs results in a system that should be able to beamform a single beam for a station with five or nine FPGAs, depending on the utilisation rate of the FPGA and the number of channels.

7.2 Future work

There are several strands of this work that would be interesting to develop further. The flexible stations concept has much further scope to explore, for instance evaluating the potential benefits of elliptical stations for observations away from the zenith. The dense core of SKA1-Low, which currently is specified to contain 50% of antennas within a radius of 600 m, gives rise to the possibility of implementing a flexible station capability in the core without significantly affecting the positions of antennas; the potential benefits of this should be assessed.

Implementing flexible stations in any form would change the beamformer requirements and a study on how to best implement the digital signal processing and the costs involved would be required. If these were to prove too high, technological progress with Moore's law could facilitate its introduction at some point in the future.

The work in Chapter 5 suggests that although the radial distribution of antennas for SKA1-Low is effective it may not be optimal of measuring the Epoch of Reionisation power spectrum. Alternative radial distributions should be considered, although

the implications for other science cases should also be considered.

Finally, the recent rebaselining of the SKA to halve the number of antennas and stations in SKA1-Low must be considered when interpreting the findings here.

7.3 Concluding remarks

The SKA has already led a charge in radio astronomy technological development. It has become a focal point of international astronomy. The vast increases in sensitivity it promises will undoubtedly lead to new discoveries. The technology that will make the SKA possible also opens the possibility of moving away from traditional aperture array designs. This opportunity should be explored. The unprecedented scale, however, could also cause problems. Efforts must be made to constrain noise-contributing factors so as to allow this prodigious instrument to deliver far-reaching insights.

Bibliography

[1] K. G. Jansky, Nature **132**, 66 (1933).

[2] K. G. Jansky, Proc. Institute of Radio Engineers **21**, 1387 (1933).

[3] G. Reber, ApJ **91**, 621 (1940).

[4] G. Reber, ApJ **100**, 279 (1944).

[5] M. Goss, in Resolving The Sky - Radio Interferometry: Past, Present and Future (2012), p. 1.

[6] J. G. Bolton and O. B. Slee, Australian Journal of Physics **6**, 420 (1953).

[7] H. C. van de Hulst, Ned. Tijd. Natuurkunde **11**, 210 (1945) (Translation: W.T. Sullivan, III, ed., Classics in Radio Astronomy (D. Reidel, Dordrecht, 1982)).

[8] H. I. Ewen and E. M. Purcell, Nature **168**, 356 (951).

[9] C. A. Muller and J. H. Oort, Nature **168**, 357 (1951).

[10] J. L. Pawsey, Nature **168**, 358 (1951).

[11] W. N. Christiansen and J. V. Hindman, Australian Journal of Scientific Research A Physical Sciences **5**, 437 (1952).

[12] M. Ryle, Nature **194**, 517 (1962).

[13] M. Ryle, B. Elsmore, and A. C. Neville, Nature **205**, 1259 (1965).

[14] A. A. Penzias and R. W. Wilson, ApJ **142**, 419 (1965).

[15] A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins, Nature **217**, 709 (1968).

[16] M. Ryle, Nature **239**, 435 (1972).

[17] J. W. M. Baars, J. F. van der Brugge, J. L. Casse, J. P. Hamaker, L. H. Sondaar, J. J. Visser, and K. J. Wellington, IEEE Proceedings **61**, 1258 (1973).

[18] A. R. Thompson, B. G. Clark, C. M. Wade, and P. J. Napier, ApJS **44**, 151 (1980).

[19] R. H. Frater, J. W. Brooks, and J. B. Whiteoak, Journal of Electrical and Electronics Engineering Australia **12**, 103 (1992).

[20] G. Swarup, Current Science **60**, 106 (1991).

[21] R. T. Schilizzi, A. van Ardenne, G. K. Miley, B. Baud, L. Baath, B. O. Ronnang, and I. I. K. Pauliny-Toth, A&A **77**, 1 (1979).

[22] J. Jonas, in From Planets to Dark Energy: the Modern Radio Universe (2007), p. 7.

[23] D. R. DeBoer, R. G. Gough, J. D. Bunton, T. J. Cornwell, R. J. Beresford, S. Johnston, I. J. Feain, A. E. Schinckel, C. A. Jackson, M. J. Kesteven, et al., IEEE Proceedings **97**, 1507 (2009).

[24] S. J. Tingay, R. Goeke, J. D. Bowman, D. Emrich, S. M. Ord, D. A. Mitchell, M. F. Morales, T. Booler, B. Crosse, R. B. Wayth, et al., PASA **30**, 7 (2013).

[25] M. P. van Haarlem, M. W. Wise, A. W. Gunst, G. Heald, J. P. McKean, J. W. T. Hessels, A. G. de Bruyn, R. Nijboer, J. Swinbank, R. Fallows, et al., A&A **556**, A2 (2013).

[26] R. M. Prestage, K. T. Constantikes, T. R. Hunter, L. J. King, R. J. Lacasse, F. J. Lockman, and R. D. Norrod, IEEE Proceedings **97**, 1382 (2009).

[27] R. Wielebinski, N. Junkes, and B. H. Grahl, Journal of Astronomical History and Heritage **14**, 3 (2011).

[28] D. B. Campbell, M. M. Davis, P. F. Goldsmith, and P. J. Perillat, in American Astronomical Society Meeting Abstracts (1997), vol. 29 of Bulletin of the American Astronomical Society, p. 1305.

[29] R. Nan, D. Li, C. Jin, Q. Wang, L. Zhu, W. Zhu, H. Zhang, Y. Yue, and L. Qian, International Journal of Modern Physics D **20**, 989 (2011).

[30] B. Y. Mills, A. G. Little, K. V. Sheridan, and O. B. Slee, Proc. IRE **46**, 67 (1958).

[31] M. Ryle, Electrical Engineers, Journal of the Institution of **6**, 14 (1960).

[32] B. Y. Mills and O. B. Slee, Australian Journal of Physics **10**, 162 (1957).

[33] B. Y. Mills, O. B. Slee, and E. R. Hill, Australian Journal of Physics **11**, 360 (1958).

[34] B. Y. Mills, O. B. Slee, and E. R. Hill, Australian Journal of Physics **13**, 676 (1960).

[35] B. Y. Mills, O. B. Slee, and E. R. Hill, Australian Journal of Physics **14**, 497 (1961).

[36] J. F. R. Gower, P. F. Scott, and D. Wills, MmRAS **71**, 49 (1967).

[37] H. J. Visser, Array and Phased Array Antenna Basics (John Wiley and Sons, Ltd, Chichester, England, 2005).

[38] A. R. Thompson, J. M. Moran, and G. W. Swenson, Jr., Interferometry and Synthesis in Radio Astronomy (Wiley, New York, 2001), 2nd ed.

[39] M. Jones, K. Zarb Adami, S. Salvini, A. Faulkner, V. Khlebnikov, and C. Shenton, SKA Station Beamformer Concept Description, Online, accessed Aug 2013. Available at: http://www.skatelescope.org/public/2011-04_Signal_Processing_CoDR_Documents/03i-wp2-040.120.010-td-001-1-StationBFM-concept-description.pdf (2011).

[40] C. E. Shannon, IEEE Proceedings **37**, 10 (1949).

[41] R. Braun and W. van Cappellen, SKA Memo Series **87** (2006).

[42] W. A. van Cappellen, S. J. Wijnholds, and J. D. Bregman, in Radar Conference, 2006. EuRAD 2006. 3rd European (2006), pp. 76–79.

[43] T. J. Cornwell, K. Golap, and S. Bhatnagar, IEEE Journal of Selected Topics in Signal Processing **2**, 647 (2008).

[44] J. P. Hamaker, J. D. Bregman, and R. J. Sault, A&AS **117**, 137 (1996).

[45] J. P. Hamaker, A&AS **143**, 515 (2000).

[46] O. M. Smirnov, A&A **527**, A106 (2011).

[47] O. M. Smirnov, A&A **527**, A107 (2011).

[48] O. M. Smirnov, A&A **527**, A108 (2011).

[49] O. M. Smirnov, A&A **531**, A159 (2011).

[50] F. Dulwich, B. J. Mort, and S. Salvini, OSKAR Theory of Operation, Online, accessed Aug 2013. Available at www.oerc.ox.ac.uk/~ska/oskar2/OSKAR-Theory.pdf (2012).

[51] R. J. Mailloux, Phased Array Antenna Handbook (Artech House, London, 2005), 2nd ed.

[52] B. F. Burke and F. Graham-Smith, An Introduction to Radio Astronomy (Cambridge University Press, 2010).

[53] Y. V. Pavlenko, I. B. Vavilova, and T. Kostiuk, Organizations and Strategies in Astronomy, Vol. 6 **343**, 71 (2006).

[54] J. E. Noordam, R. Braun, and A. G. de Bruyn, Netherlands Foundation for Radio Astronomy Note nr 585 (1991).

[55] P. N. Wilkinson, in Radio Interferometry: Theory, Techniques and Applications, edited by T. J. Cornwell and R. A. Perley (IAU Colloqium 131, 1991), vol. 19 of ASP Conference Series, pp. 428–432.

[56] W. van Driel, T. Gergely, H. Liszt, and M. Ohishi, Expert Panel on Radio Quiet Zone and RFI Regulation Report, Online, accessed Jul 2015. Available at: `https://www.skatelescope.org/wp-content/uploads/2012/06/79_Report_of_The_Expert_Panel_on_Radio_Quiet_Zone_and_RFI_Regulation.pdf` (2011).

[57] A. R. Offringa, R. B. Wayth, N. Hurley-Walker, D. L. Kaplan, N. Barry, A. P. Beardsley, M. E. Bell, G. Bernardi, J. D. Bowman, F. Briggs, et al., PASA **32**, 8 (2015).

[58] J. M. Moran, S. Ananthakrishnan, J. W. M. Baars, J. Bell Burnell, W. N. Brouw, I. Corbett, J. Crocker, T. Garvin, S. Michalowski, E. R. Seaquist, et al., Report and recommendation of the SKA Site Advisory Committee (SSAC), Online, accessed Apr 2015. Available at `https://www.skatelescope.org/site-documentation` (2012).

[59] K. Rohlfs and T. L. Wilson, Tools of Radio Astronomy (Springer, 2006), 4th ed.

[60] P. E. Dewdney, W. Turner, R. Millenaar, R. McCool, J. Lazio, and T. J. Cornwell, SKA1 System Baseline Design, Online, accessed Aug 2013. Available at: http://www.skatelescope.org/wp-content/uploads/2013/03/SKA-TEL-SKO-DD-001-1_BaselineDesign1.pdf (2013).

[61] W. Turner, T. J. Cornwell, A. McPherson, and P. Diamond, SKA Phase 1 System (Level 1) Requirements Specification, Rev 6 (2015).

[62] C. Carilli and S. Rawlings, Science with the Square Kilometre Array (Elsevier, 2004).

[63] T. L. Bourke, R. Braun, R. Fender, F. Govoni, J. Green, M. Hoare, M. Jarvis, M. Johnston-Hollitt, E. Keane, L. Koopmans, et al., eds., Advancing Astrophysics with the Square Kilometre Array, Giardini Naxos, Italy (Proceedings of Science, 9-13 Jun 2014).

[64] C. L. Carilli, in Advancing Astrophysics with the Square Kilometre Array (AASKA14) (2014), p. 171.

[65] R. Braun, T. Bourke, J. A. Green, E. Keane, and J. Wagg, in Advancing Astrophysics with the Square Kilometre Array (AASKA14) (2015), p. 174.

[66] T. Kitching, D. Bacon, M. Brown, P. Bull, J. McEwen, M. Oguri, R. Scaramella, K. Takahashi, K. Wu, and D. Yamauchi, in Advancing Astrophysics with the Square Kilometre Array (AASKA14) (2015), p. 146.

[67] R. Maartens, F. B. Abdalla, M. Jarvis, and M. G. Santos, in Advancing Astrophysics with the Square Kilometre Array (AASKA14) (2015), p. 16.

[68] L. Staveley-Smith and T. Oosterloo, in Advancing Astrophysics with the Square Kilometre Array (AASKA14) (2015), p. 167.

[69] I. Prandoni and N. Seymour, in Advancing Astrophysics with the Square Kilometre Array (AASKA14) (2015), p. 67.

[70] L. Testi, L. Perez, I. Jimenez-Serra, M. G. Hoare, A. Boley, T. L. Bourke, J. R. Brucato, P. Caselli, C. J. Chandler, A. Isella, et al., in Advancing Astrophysics with the Square Kilometre Array (AASKA14) (2015), p. 117.

[71] A. Siemion, J. Benford, J. Cheng-Jin, J. Chennamangalam, J. M. Cordes, H. D. E. Falcke, S. T. Garrington, M. A. Garrett, L. Gurvits, M. Hoare, et al., in Advancing Astrophysics with the Square Kilometre Array (AASKA14) (2015), p. 116.

[72] S. A. Torchinsky, in Wide Field Astronomy & Technology for the Square Kilometre Array (2009), p. 4.

[73] R. Fender, A. Stewart, J. P. Macquart, I. Donnarumma, T. Murphy, A. Deller, Z. Paragi, and S. Chatterjee, in Advancing Astrophysics with the Square Kilometre Array (AASKA14) (2015), p. 51.

[74] J. D. Bowman, I. Cairns, D. L. Kaplan, T. Murphy, D. Oberoi, L. Staveley-Smith, W. Arcus, D. G. Barnes, G. Bernardi, F. H. Briggs, et al., PASA **30**, 31 (2013).

[75] A. R. Parsons, D. C. Backer, G. S. Foster, M. C. H. Wright, R. F. Bradley, N. E. Gugliucci, C. R. Parashare, E. E. Benoit, J. E. Aguirre, D. C. Jacobs, et al., AJ **139**, 1468 (2010).

[76] Z. S. Ali, A. R. Parsons, H. Zheng, J. C. Pober, A. Liu, J. E. Aguirre, R. F. Bradley, G. Bernardi, C. L. Carilli, C. Cheng, et al., ArXiv e-prints (2015), 1502.06016.

[77] D. C. Jacobs, J. C. Pober, A. R. Parsons, J. E. Aguirre, Z. S. Ali, J. Bowman, R. F. Bradley, C. L. Carilli, D. R. DeBoer, M. R. Dexter, et al., ApJ **801**, 51 (2015).

[78] G. B. Taylor, S. W. Ellingson, N. E. Kassim, J. Craig, J. Dowell, C. N. Wolfe, J. Hartman, G. Bernardi, T. Clarke, A. Cohen, et al., Journal of Astronomical Instrumentation **1**, 50004 (2012).

[79] J. Kocz, L. J. Greenhill, B. R. Barsdell, D. Price, G. Bernardi, S. Bourke, M. A. Clark, J. Craig, M. Dexter, J. Dowell, et al., Journal of Astronomical Instrumentation **4**, 50003 (2015).

[80] S. W. Ellingson, T. E. Clarke, A. Cohen, J. Craig, N. E. Kassim, Y. Pihlstrom, L. J. Rickard, and G. B. Taylor, IEEE Proceedings **97**, 1421 (2009).

[81] L. J. Greenhill and G. Bernardi, ArXiv e-prints (2012), 1201.1700.

[82] G. Paciga, T.-C. Chang, Y. Gupta, R. Nityanada, J. Odegova, U.-L. Pen, J. B. Peterson, J. Roy, and K. Sigurdson, MNRAS **413**, 1174 (2011).

[83] E. de Lera Acedo, B. Fiorelli, S. Wijnholds, N. Razavi-Ghods, A. Sutinjo, C. Jackson, F. Perini, P. Bolli, G. Virone, and K. Grainge, Low Frequency Aperture Array Substantiation of Level 1 Antenna/Station Requirements, AADC Consortium (2014).

[84] C. Carilli and S. Rawlings, eds., Science with the Square Kilometre Array (Elsevier, 2004).

[85] G. Mellema, L. V. E. Koopmans, F. A. Abdalla, G. Bernardi, B. Ciardi, S. Daiboo, A. G. de Bruyn, K. K. Datta, H. Falcke, A. Ferrara, et al., Experimental Astronomy **36**, 235 (2013).

[86] D. C. Price, J. Hickish, D. Sinclair, and M. E. Jones, SKA Memo Series **150** (2013).

[87] R. Braun, A&A **551**, A91 (2013), 1208.6352.

[88] J. Lazio, ed., The Square Kilometre Array Design Reference Mission: SKA Phase 1 (2011).

[89] K. Grainge, SKA Memo Series **151** (2014).

[90] K. Carver, W. Cooper, and W. Stutzman, IEEE Trans. Antennas and Propagation **21**, 199 (1973).

[91] S. Schediwy, D. Price, F. Dulwich, and B. Mort, in Phased Array Systems and Technology (ARRAY), 2010 IEEE International Symposium on (2010), pp. 256–260.

[92] B. J. Mort, F. Dulwich, S. Salvini, K. Z. Adami, and M. E. Jones, in Phased Array Systems and Technology (ARRAY), 2010 IEEE International Symposium on (IEEE, 2010), pp. 690–694.

[93] R. Armstrong, J. Hickish, K. Zarb Adami, and M. E. Jones, in Wide Field Astronomy and Technology for the Square Kilometre Array, edited by S. A. Torchinsky, A. van Ardenne, T. van den Brink-Havinga, A. J. J. van Es, and A. Faulkner (2009), p. 273.

[94] O. M. Smirnov, A&A **527**, A106 (2011), 1101.1764.

[95] P. E. Dewdney, J.-G. bij de Vaate, K. Cloete, A. Gunst, D. Hall, R. McCool, N. Noddis, and W. Turner, SKA Memo Series **130** (2010).

[96] W. Turner, T. J. Cornwell, A. McPherson, and P. Diamond, SKA Phase 1 System (Level 1) Requirements Specification, Rev 5 (2014).

[97] A. J. Faulkner, Low Frequency Aperture Array Technical Description, Online, accessed Jan 2014. Available at `https://www.skatelescope.org/wp-content/uploads/2013/09/AADC-TEL.LFAA_.SE_.MGT-AADC-PL-002-Technical-Description.pdf` (2013).

[98] A. Faulkner, Interfaces, Presentation, Aperture Array Design and Construction (AADC) consortium meeting, ASTRON (2013).

[99] W. Turner, T. J. Cornwell, P. E. Dewdney, and R. McCool, SKA Phase 1 System (Level 1) Requirements Specification, Rev 1 (2013).

[100] N. Razavi-Ghods, B. Mort, F. Dulwich, E. de Lera Acedo, K. Grainge, R. Bolton, M. Jones, and D. Sinclair, Station Response and Imaging Performance: LFAA 100 MHz to 600 MHz, Online, accessed Jul 2015. Available at: `https://skaoffice.atlassian.net/wiki/download/attachments/22183971/Appendix+C+-+Station+Response+and+Imaging+Performace+of+LFAA+for+100+to+600MHz.pdf` (2014).

[101] R. C. Hansen, Phased Array Antennas (John Wiley and Sons, Ltd, Chichester, England, 2009).

[102] W. Turner, T. J. Cornwell, and R. McCool, SKA Phase 1 System (Level 1) Requirements Specification (2013).

[103] E. de Lera Acedo, in Electromagnetics in Advanced Applications (ICEAA), 2012 International Conference on (2012), pp. 353–356.

[104] E. de Lera Acedo, in AADC (Aperture Array Design Consortium) Meeting (ASTRON, Dwingeloo, 2013).

[105] J. M. Wrobel and R. C. Walker, in Synthesis Imaging in Radio Astronomy II, Astronomical Society of the Pacific (ASP Conference Series, vol. 180, 1999).

[106] C. G. T. Haslam, U. Klein, C. J. Salter, H. Stoffel, W. E. Wilson, M. N. Cleary, D. J. Cooke, and P. Thomasson, A&A **100**, 209 (1981).

[107] C. G. T. Haslam, C. J. Salter, H. Stoffel, and W. E. Wilson, A&AS **47**, 1 (1982).

[108] B. Mort and F. Dulwich, Private Communication.

[109] U. J. Schwarz, A&A **65**, 345 (1978).

[110] U. J. Schwarz, in IAU Colloq. 49: Image Formation from Coherence Functions in Astronomy, edited by C. van Schooneveld (1979), vol. 76 of Astrophysics and Space Science Library, p. 261.

[111] N. Jackson, Lect. Notes. Phys. **742**, 193 (2008).

[112] R. N. Bracewell and J. A. Roberts, Australian Journal of Physics **7**, 615 (1954).

[113] J. A. Högbom, A&AS **15**, 417 (1974).

[114] B. G. Clark, A&A **89**, 377 (1980).

[115] F. R. Schwab, AJ **89**, 1076 (1984).

[116] J. G. Ables, A&AS **15**, 383 (1974).

[117] J. E. B. Ponsonby, MNRAS **163**, 369 (1973).

[118] S. F. Gull and G. J. Daniell, Nature **272**, 686 (1978).

[119] T. J. Cornwell, R. Braun, and D. S. Briggs, in Synthesis Imaging in Radio Astronomy II (Astronomical Society of the Pacific, 1999), ASP Conference Series, vol. 180.

[120] J. J. Condon, W. D. Cotton, E. W. Greisen, Q. F. Yin, R. A. Perley, G. B. Taylor, and J. J. Broderick, AJ **115**, 1693 (1998).

[121] P. Madau, A. Meiksin, and M. J. Rees, ApJ **475**, 429 (1997).

[122] L. Koopmans, J. Pritchard, G. Mellema, J. Aguirre, K. Ahn, R. Barkana, I. van Bemmel, G. Bernardi, A. Bonaldi, F. Briggs, et al., in Advancing Astrophysics with the Square Kilometre Array (AASKA14) (2015), p. 1.

[123] J. R. Pritchard and A. Loeb, Reports on Progress in Physics **75** (2012).

[124] C. L. Carilli, S. Furlanetto, F. Briggs, M. Jarvis, S. Rawlings, and H. Falcke, New Astronomy Reviews **48**, 1029 (2004).

[125] M. Taoso, G. Bertone, G. Meynet, and S. Ekström, Phys. Rev. D **78**, 123510 (2008).

[126] E. Ripamonti, F. Iocco, A. Ferrara, R. Schneider, A. Bressan, and P. Marigo, MNRAS **406**, 2605 (2010).

[127] T. Theuns, J. Schaye, S. Zaroubi, T.-S. Kim, P. Tzanavaris, and B. Carswell, ApJLett **567**, L103 (2002).

[128] L. Hui and Z. Haiman, ApJ **596**, 9 (2003).

[129] P. Tozzi, P. Madau, A. Meiksin, and M. J. Rees, ApJ **528**, 597 (2000).

[130] M. Zaldarriaga, S. R. Furlanetto, and L. Hernquist, ApJ **608**, 622 (2004).

[131] J. D. Bowman, M. F. Morales, and J. N. Hewitt, ApJ **661**, 1 (2007).

[132] M. McQuinn, O. Zahn, M. Zaldarriaga, L. Hernquist, and S. R. Furlanetto, ApJ **653**, 815 (2006).

[133] M. G. Santos, L. Ferramacho, M. B. Silva, A. Amblard, and A. Cooray, MNRAS **406**, 2421 (2010).

[134] M. McQuinn, A. Lidz, O. Zahn, S. Dutta, L. Hernquist, and M. Zaldarriaga, MNRAS **377**, 1043 (2007).

[135] S. R. Furlanetto, M. Zaldarriaga, and L. Hernquist, ApJ **613**, 1 (2004).

[136] J. A. Peacock, Cosmological Physics (Cambridge University Press, 1999).

[137] M. Remazeilles, C. Dickinson, A. J. Banday, M.-A. Bigot-Sazy, and T. Ghosh, ArXiv e-prints (2014), 1411.3628.

[138] H. T. Intema, R. J. van Weeren, H. J. A. Röttgering, and D. V. Lal, A&A (2011).

[139] P. A. Shaver, R. A. Windhorst, P. Madau, and A. G. de Bruyn, A&A (1999).

[140] T. Di Matteo, R. Perna, T. Abel, and M. J. Rees, ApJ (2002).

[141] S. P. Oh and K. J. Mack, MNRAS (2003).

[142] S. R. Furlanetto, M. Zaldarriaga, and L. Hernquist, ApJ (2004).

[143] G. Harker, S. Zaroubi, G. Bernardi, M. A. Brentjens, A. G. de Bruyn, B. Ciardi, V. Jelić, L. V. E. Koopmans, P. Labropoulos, G. Mellema, et al., MNRAS (2009).

[144] M. G. Santos, A. Cooray, and L. Knox, ApJ (2005).

[145] X. Wang, M. Tegmark, M. G. Santos, and L. Knox, ApJ2006 (2006).

[146] J. D. Bowman, M. F. Morales, and J. N. Hewitt, ApJ (2009).

[147] A. Bonaldi and M. L. Brown, MNRAS **447**, 1973 (2015).

[148] E. Chapman, A. Bonaldi, G. Harker, V. Jelić, F. B. Abdalla, G. Bernardi, J. Bobin, F. Dulwich, B. Mort, M. Santos, et al., ArXiv e-prints (2015), 1501.04429.

[149] T. Di Matteo, B. Ciardi, and F. Miniati, MNRAS **355**, 1053 (2004).

[150] M. E. Jones, Private Communciation.

[151] A. P. Beardsley, B. J. Hazelton, M. F. Morales, R. J. Capallo, R. Goeke, D. Emrich, C. J. Lonsdale, W. Arcus, D. Barnes, G. Bernardi, et al., MNRAS **425**, 1781 (2012).

[152] R. A. Perley, C. J. Chandler, B. J. Butler, and J. M. Wrobel, ApJLett **739** (2011).

[153] S. T. Garrington, B. Anderson, C. Baines, J. A. Battilana, M. N. Bentley, D. Brown, P. Burgess, P. J. Diamond, G. J. Kitching, R. McCool, et al., in Ground-based Telescopes, edited by J. M. Oschmann, Jr. (2004), vol. 5489 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, pp. 332–343.

[154] G. S. Foster, Ph.D. thesis, University of Oxford (2013).

[155] M. de Vos, A. W. Gunst, and R. Nijboer, IEEE Proceedings **97**, 1431 (2009).

[156] C. J. Lonsdale, R. J. Cappallo, M. F. Morales, F. H. Briggs, L. Benkevitch, J. D. Bowman, J. D. Bunton, S. Burns, B. E. Corey, L. Desouza, et al., IEEE Proceedings **97**, 1497 (2009).

[157] J. D. Mol and J. Romein, J. W., in Euro-Par 2011 Parallel Processing, edited by E. Jeannot, R. Namyst, and J. Roman (Springer Berlin Heidelberg, 2011), vol. 68535, pp. 328–339.

[158] A. Sclocco, Master's thesis, Vrije Universiteit Amsterdam (2011).

[159] J. Fowers, G. Brown, P. Cooke, and G. Stitt, in FPGA '12 Proceedings of the ACM/SIGDA International Symposium on Field Programmable Gate Arrays (2012).

[160] S. Ord, L. Greenhill, R. Wayth, D. Mitchell, K. Dale, H. Pfister, and R. G. Edgar, ArXiv e-prints (2009), 0902.0915.

[161] M. Serylak, A. Karastergiou, C. Williams, W. Armour, M. Giles, and LOFAR Pulsar Working Group, in IAU Symposium, edited by J. van Leeuwen (2013), vol. 291, pp. 492–494.

[162] J. Chennamangalam, S. Scott, G. Jones, H. Chen, J. Ford, A. Kepley, D. R. Lorimer, J. Nie, R. Prestage, D. A. Roshi, et al., PASA **31**, 48 (2014).

[163] P. J. Napier, A. R. Thompson, and R. D. Ekers, IEEE Proceedings **71**, 1295 (1983).

[164] R. Escoffier, VLBA Correlator Memo **87** (1987).

[165] I. Kuon and J. Rose, IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems **26**, 203 (2007).

[166] S. M. Trimberger, Proceedings of the IEEE **103**, 318 (2015).

[167] A. Parsons, D. Backer, A. Siemion, H. Chen, D. Werthimer, P. Droz, T. Filiba, J. Manley, P. McMahon, A. Parsa, et al., PASP **120**, 1207 (2008).

[168] D. Price, Ph.D. thesis, University of Oxford (2012).

[169] C. J. Copley, Ph.D. thesis, University of Oxford (2013).

[170] H. H. Goldstine, A History of Numerical Analysis from the 16th through 19th Century (Springer-Verlag, New York, 1977).

[171] J. W. Cooley and J. W. Tukey, Mathematics of Computation **19** (1965).

[172] M. G. Bellanger, G. Bonnerot, and M. Coudreuse, IEEE Trans. Acoustics, Speech and Signal Processing **24**, 109 (1976).

[173] J. Chennamangalam, The Polyphase Filter Bank Technique, CASPER memo 41 (2011).

[174] Xilinx Inc., Spartan-6 FPGA Spartan-6 DSP48A1 Slice User Guide, Onine, accessed Aug 2013. Available at `http://www.xilinx.com/support/documentation/user_guides/ug389.pdf` (2009).

[175] J. Hickish, Ph.D. thesis, University of Oxford (2013).

[176] R. P. Armstrong, Ph.D. thesis, University of Oxford (2011).

[177] M. Tegmark and M. Zaldarriaga, Phys. Rev. D **79** (2009).

[178] Xilinx Inc., UltraScale Architecture and Product Overview, Online, accessed Apr 2015. Available at `http://www.xilinx.com/support/documentation/data_sheets/ds890-ultrascale-overview.pdf` (2015).

[179] Xilinx Inc., 7 Series FPGAs Overview, Online, accessed Mar 2015. Available at `http://www.xilinx.com/support/documentation/data_sheets/ds180_7Series_Overview.pdf` (2014).

[180] Xilinx Inc., Virtex-6 Family Overview, Online, accessed Mar 2015. Available at `http://www.xilinx.com/support/documentation/data_sheets/ds150.pdf` (2012).

[181] Xilinx Inc., Virtex-7 Product Brief, Online, accessed Apr 2015. Available at `http://www.xilinx.com/publications/prod_mktg/Virtex7-Product-Brief.pdf` (2014).

[182] S. Leibson and N. Mehta, Xilinx UltraScale: The Next-Generation Architecture for Your Next-Generation Architecture, Xilinx White Paper (2014).