Investigation of Horn Optimisation, Simulation and Measurement Techniques, with Application to Integrating Cavity Based Pixel Design

Presented by

Darragh McCarthy, B.Sc.

A thesis submitted for the degree of

Doctor of Philosophy

Terahertz Space Optics Group,
Department of Experimental Physics,
Maynooth University,
National University of Ireland Maynooth,
Maynooth,
County Kildare,
Ireland

September 2014

Research Supervisor

Dr. Neil Trappe, B.Sc., Ph.D.

Head of Department

Prof. J. A. Murphy, B.Sc., M.S., M.Sc., Ph.D.

Contents

List of Figures

Abstract

The main topic of this thesis is to examine the design and analysis of millimetre and sub-millimetre optical systems for use in future cosmic microwave background radiation and far infrared space based astronomy missions. Astronomical observations in these regions of the electromagnetic spectrum are vital for studying the formation and evolution of planets, stars, galaxies, and on an even larger scale the origin of the Universe, through an analysis of the Cosmic Microwave Background (CMB) radiation. In order to further our knowledge of these astronomical processes it will be necessary to use highly calibrated instruments with increased sensitivity, and the capability to measure the polarisation signal of the cosmic microwave background radiation. It will also be necessary to design instruments with higher angular resolution so as measurements can be made on a finer scale. This will allow our current understanding of the origin and evolution of our Universe to be better understood by placing further constraints on the various models that currently exist which serve to explain these processes, for example the various inflationary models. Significant efforts are being made to develop and characterise the technology required to realise these ambitions. Examples of missions that are currently being planned are COrE (Cosmic Origins Explorer) Plus and SPICA (SPace Infrared telescope for Cosmology and Astrophysics), which contains an instrument known as SAFARI, a far infrared Fourier transform spectrometer that makes use of a Mach-Zehnder optical configuration.

Such missions require high performance focal plane antennas, a position traditionally filled using corrugated horns. Future missions will require large arrays of high performance antennas, and the thermal and mechanical penalties associated with high mass corrugated horns (in addition to their cost) can become prohibitive. In this thesis, the design of an alternative smooth-walled profiled horn antenna will be examined. The resulting simple design is easier and cheaper to manufacture than its corrugated counterpart, and carries less mechanical penalties. The horn was designed using a genetic algorithm based optimisation process that was applied to an implementation of the mode matching technique (used to model the performance of horn antennas) that has been optimised for computational efficiency. A number of improvements in the execution of the mode matching code were made to increase its efficiency and are described, including singular value decomposition analysis and reducing the number of waveguide modes used in the analysis. The limitations of the approach are considered, including determining the ideal horn performance metric to optimise with respect to in order to maintain the required levels of performance in all of the important characteristics. A further goal of the optimisation process was to use only desktop computing power and to complete the process within a reasonable amount of time (within a day).

A further area analysed is the multi-moded detector pixels for use in the SAFARI instrument. Such pixels are currently being tested in the Space Research Organisation of the

Netherlands (SRON), Groningen. Models already exist to attempt to describe such pixels, however in order to attempt to improve the agreement between the measured and simulated results it is necessary to refine the model. Such pixels contain, for mechanical reasons, a section between the horn and cavity which can be considered to be a free space gap. This thesis considers the implementation of such a gap within the mode matching technique, and examines the impact of the gap on the measured pixel efficiency. This is critical for understanding the optical performance of the receiver system. Future missions will make use of such multi-moded systems as they maximise the amount of power available to the detector, a critical feature when low strength signals are being measured. Measurements of multi-moded horns are not well understood. Consideration is given in this thesis to a method that may be used in the future, with further development, to measure the performance of multi-moded horns. Such measurements will lead to a better understanding of multi-moded systems. A brief description of the background science is also described in the first chapter, along with the specific contributions of the author.

Acknowledgements

I would firstly like to sincerely thank my supervisor Dr. Neil Trappe for all of his help, guidance and support over the last three and a bit years (and in the years leading up to my time as a postgraduate), it can't have been easy at times. He invested countless hours in my work, but always had unlimited time and patience for both work and non work related chats, even when I was ranting about Declan Kidney. His enthusiasm and expertise meant that motivation was never an issue for me and he could always dig me out of a hole when I was stuck. I could not have asked for a better supervisor. I would also like to thank Professor Anthony Murphy, Dr. Marcin Gradziel, Dr. Créidhe O'Sullivan, Dr. Stephen Doherty and Dr. Massimo Candotti for the hours they all spent discussing my work with me and helping me. Thank you also to Mr.Ian McAuley, Mr. Enda McLoughlin and Mr. Niall Tynan for their help with the measurements side of my work. A special thanks to Professor Anthony Murphy for giving me the opportunity to join the Terahertz Space Optics group and to be a part of so many interesting projects, and also for the use of the facilities of the department.

A huge thanks goes to all of the other members of the department, including Ms. Gráinne Roche for just about everything (!), Mr. John Kelly for sorting out the maze of technological issues I had, Mr. Ian McAuley for all of the help with the teaching (and terahertz!) side of things and to Mr. Derek Gleeson for all of his help and advice on work matters. Thanks also to Dr. Frank Mulligan, Dr. Michael Cawley and Dr. Peter van der Burgt for everything. As both an undergrad and a postgrad they were always there with help and advice. I would like to thank all of the postgrads in the department, who made the last three years very enjoyable. There was always a laugh and a distraction when needed. To the people that I shared offices with over the years - Paul, Stephen, Enda, Anthony and Niamh. I think we can all agree that we had some discussions that won't be forgotten about. Ever. It can all be traced back to Francis Mahon. John Paul, Bobby and Damo-need I say anything else?

I would like to thank Dr. Bruno Maffei and Dr. Fahri Ozturk of the Jodrell Bank Centre for Astrophysics at the University of Manchester, and Dr. Giampaolo Pisano of the University of Cardiff, with whom I collaborated on the portion of my thesis that deals with the off-axis illumination of horn antennas. They were of great help and always made the trips to Manchester great fun. I worked closely with Dr. Massimo Candotti on the Golay cell testbed that we developed in the department in this regard, so a huge thanks to him also for all of the help.

I would like to formally acknowledge and thank the Irish Research Council (formally the Irish Research Council for Science, Engineering and Technology) for generously funding my Ph.D. through the Embark Initiative. I would also like to acknowledge and thank the Royal Irish Academy for the travel grant that allowed our collaboration with the University of Manchester.

Finally, I would like to thank my family (Mum, Dad, Killian and Órla) for all of their support, encouragement and help throughout all of the years that I have spent in college.

Chapter 1

Introduction

1.1 Introduction

The Universe is widely accepted to have come into existence in the form that we know it, due to an event known as the Big Bang, that occurred approximately 14 billion years ago. At its beginning, the Universe was compressed into a very small, extremely dense space, at an extremely high temperature. The Big Bang represents the rapid expansion or 'explosion' of this space, and gave rise to various subatomic particles (which ultimately became matter and formed the constituents of the Universe), as well as space and time. The Big Bang left a footprint in the form of remnant radiation that permeates the entire Universe currently, and this is known as the Cosmic Microwave Background (CMB) radiation. An in-depth study of the CMB can yield an incredible amount of information about the Universe, from its origin and potential ultimate fate, to how it has evolved over the years and why the structures that we see came into being. This will be considered in more detail in section 1.2.

Various theoretical models currently exist which attempt to describe the observable Universe using difference mechanisms. The only conclusive way to determine the correct mechanism, based on current theories, is to map the CMB in extreme detail. This allows a careful examination of the polarisation and temperature deviations within the CMB, thus shedding light on what theoretical model is most likely correct. Such detailed mappings require extremely sensitive instruments, that can also measure polarisation accurately. These instruments will likely make use of arrays of detectors (fed for example by horn or lens antennas), in order to maximise the efficiency of the measurement process.

The terahertz space optics group in NUI Maynooth has a strong heritage in the development of such systems for use in CMB measurements, including Planck, specifically the High-Frequency Instrument (HFI). The Planck surveyor satellite was a European Space Agency (ESA) mission, with several significant scientific aims, including measuring the temperature and polarisation of the CMB with sensitivity, angular resolution and a frequency range that had never been achieved by previous missions. It operated with an angular resolution as low as

5 arcminutes, in both the low (30-70 GHz) and high (100-857 GHz) frequency ranges (spread across 9 bands), with the CMB peaking at 160.2 GHz, [2]. The development of such a high mission performance required a high degree of innovation in hardware and simulation, and was made possible by new technologies emerging in the areas of bolometric detection and cryogenics. Such high performance detectors, when coupled with active and passive cooling systems allowing these detectors to function at a temperature of 0.1°C above absolute zero, made it possible to achieve these sensitivity performance goals. Planck was launched on the 14^{th} of May 2009, and entered a Lissajous orbit around the second Lagrangian point. The mission came to end in October 2013, 4 years. 5 months and 9 days later. Several papers have been released to date by the Planck collaboration, showing preliminary, intermediate and advanced results. Such papers deal with the instrumentation aspect of the mission as well as the science results, for example [3], [4] and [5].

It is clear from this, that in order for future missions to measure the CMB more accurately in order to gain a better understanding of the Universe, similarly significant advances will be required in terms of the technology used. A major component in determining the performance of these systems for space based CMB missions are the feed horns used to couple to the radiation, and so a large part of this thesis will be concerned with the design and optimisation of feed horn antennas in order to satisfy the requirements for future space based CMB missions.

The CMB will now be discussed in some detail, in order to provide some context for the necessity of its measurement.

1.2 The Cosmic Microwave Background Radiation

George Gamow proposed the idea that the Universe as we know it was created as a result of a large explosion (The Big Bang), and that the elements that are seen today were produced immediately after the explosion (within the first few minutes), at a time when the temperature and density of matter were both sufficiently high so as to allow subatomic particles to fuse [6,7], although it is now known that only hydrogen and helium were formed by the Big Bang, with the other heavy elements formed some time later by the nuclear fusion stellar process [8]. Despite this, the idea that Gamow put forward was the first one that introduced the idea of a Hot Big Bang that gave rise to a thermal radiation field with a wavelength dependent intensity distribution. When the Universe was hot and dense, as it was at this point, matter and radiation were closely coupled. This means that emitted photons were absorbed and re-emitted due to Thomson scattering, and the fact that the mean free path of the electrons was so short meant that this occurred with a frequency such as to give an optically thick plasma, with the Universe being in thermal equilibrium. Such a radiation field has a blackbody distribution. A further result from the group of George Gamow was that with the expansion of the Universe, the temperature would drop and the photons would be redshifted, whilst

retaining their blackbody distribution. The group postulated that these photons, which make up the CMB, would permeate the Universe and would have a present day temperature of approximately 5 K. The origins of the CMB, along with its measurable properties, will now be discussed in more detail.

1.2.1 The Origins of the Cosmic Microwave Background Radiation

Immediately after the Big Bang occurred, the universe was very hot and dense, and permeated with a plasma consisting of neutral hydrogen and ionised hydrogen (protons and electrons). This state of equilibrium was dynamic as opposed to static, that is to say that the Universe did not contain a static sea of free nuclei and ionised electrons, but rather a continuous mechanism existed whereby electrons would bind with free nuclei to form neutral hydrogen but would immediately be freed when the hydrogen collided with a photon. The net result of this is a sea of free nuclei and electrons. The ionisation energy of hydrogen, the energy required to remove the electron from the atom, is 13.6 eV. At this point in time the Universe was sufficiently hot for the photons that were present to have this energy, and so the hydrogen could be kept in a fully ionised state.

This was possible due to the mean free path of the electrons. Electrons and photons interact via Thomson scattering, the mean free path of the photons (δ), relative to the electrons, can be expressed according to

$$\delta = \frac{1}{n_e \sigma_e}, \tag{1.1}$$

where σ_e is the Thomson scattering cross-section for an electron and n_e is the electron number density. Since the Universe was densely populated with ionised electrons and the Thomson interaction between electrons and photons is strong (i.e. σ_e is high), then this gives rise to a small value for δ. This physically means that photons did not travel very far without colliding with an electron, which enables the process outlined above. Due to the fact that matter and radiation co-existed and interacted strongly, matter and radiation were said to be coupled, with the resulting sea of free atoms and ionised electrons behaving like an ionised plasma. Due to the fact that photons could not travel far, the Universe was also opaque.

As the Universe expanded it cooled, with the cooling resulting in a drop in the energy of the photons. Fewer and fewer photons were able to ionise hydrogen, and so the density of neutral hydrogen began to rise. The Thomson scattering cross-section of neutral hydrogen is much smaller than that of a free electron, and so the mean free path of the photons increased as a result of this, to the extent that they were able to travel without further collision, and still do to this day. The process in which this neutral Hydrogen was formed is known as recombination. The Universe was now transparent, as photons were able to travel. Since matter and radiation were no longer coupled, this process is also known as decoupling. The oldest photons that are observed today (the CMB photons) were last scattered at decoupling, giving rise to a surface

that we can see which is known as the surface of last scattering. In order to calculate the equivalent temperature of the CMB photons today, it is necessary to calculate the time when decoupling occurred.

Decoupling occurred when the photon energy dropped below the ionisation energy of Hydrogen. After the Big Bang, the photon-to-baryon ratio was extremely high ($\approx 10^9$). Of the 10^9 photons available, if even one is of an energy higher than the ionisation of Hydrogen then there would still be one photon per atom of Hydrogen and so complete decoupling would not occur. This means that when determining at what point the photon energy dropped below the ionisation energy, it is insufficient to consider the average energy of the photon distribution. Statistical mechanics must be used instead.

For a given energy I, the fraction of photons in a blackbody distribution with an energy exceeding I is given by the Boltzmann factor

$$\frac{1}{\exp\left(\frac{I}{k_B T}\right) - 1}, \tag{1.2}$$

where k_B is the Boltzmann constant, and T is the temperature of the distribution in Kelvin. For matter and radiation to remain coupled, it can be assumed that only one ionising photon is required per atom. Photons in excess with this will interact with the free electrons. If I is the ionisation energy, then the fraction of photons which are ionising is given by

$$\frac{1}{\exp\left(\frac{I}{k_B T}\right) - 1} \approx \exp\left(\frac{-I}{k_B T}\right) \approx 10^{-9}, \tag{1.3}$$

assuming a temperature of $\approx$ 50,000 K and that $\exp\left(\frac{I}{k_B T}\right) >> 1$. A further factor that must be considered is the presence of bound excited states in hydrogen. The first such state occurs at 10.2 eV. If an electron gets excited by a photon of this energy, then it can be ionised by further excitation of a photon of energy 10.2 or 13.6 eV. Thus, if the photon-to-baryon ratio is assumed to have a value of 10^{-9} at decoupling, and I is taken to be 10.2 eV, then the temperature at decoupling can readily be found to be

$$T_{dec} \approx \frac{10.2 \text{ eV}}{k_B \ln(10^9)} \approx 5700 \text{ K}. \tag{1.4}$$

A more detailed calculation gives the exact answer as being 3000-4000 K. This is equivalent to a redshift of approximately 1000. Using that

$$R = \frac{1}{z+1}, \tag{1.5}$$

this gives $R_{dec} \approx 1000$, which, assuming that $t = 0$ at the Big Bang, corresponds to decoupling

occurring approximately 380,000 yr after the Big Bang. Assuming that t_0, the current time, $\approx$ 14 byr, then by calculating the value of the scale factor for today, the temperature of the CMB photons should be $\approx$ 3 K.

1.3 Measurable Properties of the CMB: Anisotropies

NASA launched the **CO**smic **B**ackground **E**xplorer (COBE) satellite on November 18$^{\text{th}}$, 1989. COBE was designed to measure the infrared and CMB radiation and was comprised of three instruments;

- **Diffuse InfraRed Experiment**(DIRBE), the purpose of which was to locate and measure the CMB radiation
- **Differential Microwave Radiometers** (DMR), the purpose of which was to develop a map of the CMB radiation
- **Far-InfraRed Absolute Spectrophotometer** (FIRAS), the purpose of which was to measure the spectrum of the CMB radiation relative to that of a perfect blackbody

The research group headed by George Smoot worked extensively on results from COBE, [9,10]. Using FIRAS, it was shown that the CMB spectrum is described extremely accurately by that of a blackbody of temperature 2.728 K, to within a precision of 0.3 %. Therefore, the CMB should be noted to be uniform. There are reasons however, to expect deviations both above and below this value of temperature, known as hotspots and cold spots respectively. The fluctuations are known as anisotropies in the CMB, and they have been found using the COBE DMR. There are several sources of these anisotropies, which will now be examined.

1.3.1 The Dipole Anisotropy

The biggest contributor to the CMB anisotropy is that due to a Doppler shift caused by the peculiar velocity of the Milky Way through space. If a blackbody is located such that it is moving in the direction of motion of the galaxy, then its measured frequency will be higher than its rest frequency. Using Wien's law, a change in wavelength, and therefore frequency, can be related to a change in temperature such that $\lambda_{\text{max}} \propto \frac{1}{T}$. This means that if the frequency of the blackbody appears higher, then the wavelength will appear lower, and so the blackbody will actually appear to be hotter than it actually is. If the relative motion of the blackbody is in the opposite direction then it will appear colder.

By considering the relativistic Doppler shift of a blackbody, Peebles and Wilkinson were able to derive an expression for the effect of the motion of our galaxy on the observed frequency of the CMB , [11]

$$\nu_{\text{obs}} = \nu_{\text{rest}} \frac{\sqrt{1 - \frac{u^2}{c^2}}}{1 + \left(\frac{u}{c}\right) \cos\theta}, \tag{1.6}$$

where c is the speed of light. The observer is immersed in the CMB, the blackbody, and so u is the velocity of the observer and θ is the angle between the observer's line of sight and their velocity, i.e. the detector angle. By considering changes in the photon propagation angle and the detector solid angle, they were able to show that the observed temperature of the blackbody has a dependence on the detector angle, or the line of sight of the observer, given by

$$T(\theta) = \frac{\sqrt{1-\beta^2}}{1-\beta\cos\theta} T_0, \tag{1.7}$$

which can be written as

$$T(\theta) = \left(1-\beta^2\right)^{\frac{1}{2}} (1-\beta\cos\theta)^{-1} T_0. \tag{1.8}$$

Expanding equation 1.8 to the second order in β reveals that $T(\theta)$ is not a pure dipole.

$$\begin{aligned} T(\theta) &= \left(1-\frac{1}{2}\beta^2+\dots\right)\left(1+\beta\cos\theta+\beta^2\cos^2\theta+\dots\right) T_0, \\ T(\theta) &= T_0\left(1+\beta\cos\theta+\frac{\beta^2}{2}\cos 2\theta+\dots\right). \end{aligned} \tag{1.9}$$

Equation 1.9 is interpreted as the CMB having a uniform background temperature of T_0 (=2.728 K) that has a dipole anisotropy term ($\cos\theta$) superimposed on it that has amplitude βT_0. The next term superimposed upon the background is a quadrupole term ($\cos 2\theta$), that is smaller again by a factor of $\frac{\beta}{2}$. This quadrupole term is due to relativistic light aberration. Due to the fact that an observer at the surface of last scattering, where the CMB originated, would not observe such aberrations, an observer here would see a true dipole anisotropy, as all terms beyond the dipole would disappear.

The other higher order multipoles present which also contribute to the observed temperature will not be considered. The dipole term dominates the other terms and, as mentioned earlier, has been measured by the COBE mission. The experimentally obtained value for the temperature fluctuation is [12]

$$\frac{\Delta T}{T} = \frac{\pm 3.5 \text{ mK}}{2.728 \text{ K}} = \pm 1.3 \times 10^{-3}.$$

The dipole anisotropy is the largest contributor of all the CMB anisotropies, exceeding the next largest by two orders of magnitude. It is taken to be due to the motion of the Solar System relative to the CMB field.

1.3.2 The Primordial Anisotropies

The Universe is not homogeneous, but rather contains structure against a smooth background. The structures can vary in size scale from dynamically aligned kiloparsec sized galaxies, to sheets of galaxies that extend for over 100 megaparsecs or more. This structure is thought to

have originated due to gravitational instabilities acting on small density perturbations (or seeds) that were present in the Universe and then subjected to inflation. The resulting inhomogeneities in the distribution of matter at the surface of last scattering can be related to changes in temperature, which can then be measured as temperature anisotropies in the CMB radiation.

This relationship can be explained by considering the Sachs-Wolfe Effect, which gives that photons undergo a gravitational redshift as they climb out of gravitational potential wells that surround density fluctuations. Since $\rho \propto T^3$, the relationship between density and temperature perturbations can be approximated in the first order to

$$\frac{\delta\rho}{\rho} = 3\frac{\delta T}{T}. \tag{1.10}$$

Equation 1.10 implies that regions of enhanced density, where the fluctuation gives rise to a higher local density than the background, should appear cooler by an amount δT, and that regions of diminished density should appear warmer. Measurements made of the CMB by COBE in the campaign outlined earlier have uncovered primordial anisotropies that are of the order of $\frac{\delta T}{T} \approx 10^{-5}$. Since these fluctuations in temperature can be measured and are related to fluctuations in the density of matter in the early Universe, they can therefore be used to deduce information about the initial conditions that were required for the formation of the structures that we see in the Universe, from galaxies and galactic clusters, to the large voids that dominate the Universe.

1.3.3 Polarisation Anisotropies

During inflation, the Universe expanded at an incredible rate within a very small period of time. If inflation did in fact occur, then this rapid expansion would have left 'ripples' in the space-time fabric, known as gravity waves. The presence of these gravity waves should result in a distinctive contribution to the polarisation pattern of the photons that were emitted from the surface of last scattering, i.e. the CMB. Therefore, if this contribution to the polarisation pattern of the CMB, the polarisation anisotropy, can be measured then this is evidence for inflation having occurred, [13].

The underlying physics that connects these anisotropies to the initial conditions of the Universe, and ultimately the evolution of the Universe, relies mainly on linear perturbation theory and so is well understood. An analysis of the various CMB anisotropies can therefore reveal, with a high degree of confidence, the initial fluctuations and physical parameters that govern how the Universe came into being and how it came to be what we see today. Many theories exist that explain this evolution, however they do so by means of different mechanisms. The information gleaned from studying the anisotropies provides a way to place constraints on these theories, and so to ultimately move towards finding the correct theory that governs the evolution of the Universe from its beginnings to what we see today.

To understand the origins of the polarisation anisotropy of the CMB, one must begin by considering the scattering of the CMB photons from free electrons, or Thomson scattering. The incident radiation can be thought of as inducing oscillations in the scattering electron that are in the direction parallel to the plane of polarisation of the incoming radiation. If the radiation is moving in a direction that is perpendicular to the plane of polarisation, then these oscillations will therefore occur in a left to right direction, as viewed from the incoming radiation, as in figure 1.1. The intensity of the scattered radiation therefore peaks in a direction that is normal to the incident radiation, with the plane of polarisation of the scattered radiation being parallel to that of the incoming radiation.

For a given radiation field, there will be photons incident on the scattering electron from orthogonal directions, relative to each other, as shown in figure 1.1. When these both undergo

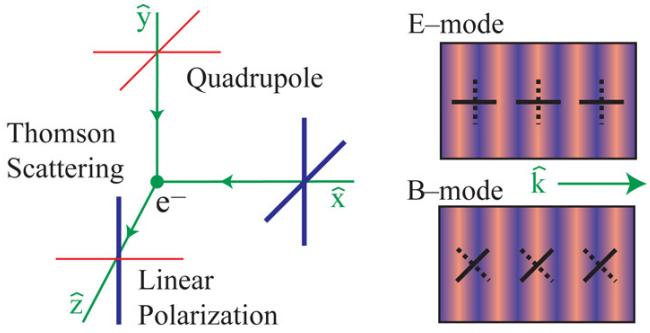

Figure 1.1: Thompson scattering of photons with a quadrupole temperature distribution. Taken from ned.ipac.caltech.edu/level5/Sept05/Hu/Hu3.html on August 27$^{\text{th}}$ 2014.

scattering from the free electron, the orthogonal components of the polarisation will align along the scattered direction such that in each case, the component of the incident photon polarisation that was parallel to the scattered direction no longer feature, as these point along the direction of travel of the scattered photons. The remaining orthogonal components of each photon therefore add to define the plane of polarisation of the scattered photons. Thus, for incident directions differing by 90°, if the radiation field is isotropic then the components of the polarisation of each photon are equal, and so there is no net polarisation for the scattered photons. If however the incident radiation has a quadrupole anisotropy (with the intensity peaking at 90°, hence the term quadrupole), then a net linear polarisation results. Clearly, if there is rapid Thomson scattering, then the number of photons with random directions will act so as to destroy any net quadrupole anisotropy and resulting polarisation features. Interpreting the CMB polarisation therefore amounts to attempting to deciphering the quadrupole temperature fluctuations at the surface of last scattering. The relationship between the quadrupole and temperature is now examined.

At decoupling, there was a temperature inhomogeneity present. When recombination had occurred and time had progressed such that regions of different temperature (hot and cold

relative to each other) could interact in terms of photons from these regions being able to scatter off the same free electron, this resulted in the scattered radiation being polarised. This was due to the fact that with the different temperatures between the regions, the intensities of the photons were different, and so as outlined above, this gave rise to a net polarisation when a photon from a warm region was incident upon a free electron normal to a photon from a cold region. It was only possible for photons from different regions to diffuse to a point where they were able to mix when the plasma that permeated space was optically thin enough, which occurred during recombination. It was also only possible for these scattering events to occur when there were still free electrons left, as these were used up during recombination. For these reasons, it was only possible for polarisation to occur during a very short period of time towards the end of recombination. This accounts for the fact that only a small fraction of the CMB radiation is actually polarised. The polarisation that is observed is not uniform. The variation is based on the fact that electrons at different locations produce polarisations of different orientations and magnitudes, however since the scattered photons were not able to diffuse too far, the polarisation is found to not vary by very much when large angular scales are considered.

The quadrupole temperature anisotropy could also be caused by other mechanisms, namely three types of perturbations which will now be considered, [13].

1. **Scalar Perturbations:** Such perturbations are caused by density fluctuations in the energy distribution of the plasma that permeated space when the CMB radiation was emitted. Such fluctuations gave rise to regions with different effective temperatures, with the fluid flowing from areas of hot effective temperature to areas of cold effective temperature. These regions are set up due to the fact that overly dense regions require that the photons must climb out of a gravitational potential well, losing energy in the process, and resulting in them being viewed as cold regions. The temperature pattern that is observed contains a quadrupole component, [13]. These perturbations give rise to curl-free (E-mode or electric mode) polarisation states following Thomson scattering of the quadrupole radiation.

2. **Vector Perturbations:** Vector perturbations are caused by vorticity in the plasma, i.e. if the fluid velocity is represented by $\mathbf{v}$, then $\nabla \cdot \mathbf{v} = 0$ and $\nabla \times \mathbf{v} \neq 0$. This vorticity is damped by inflation, the expansion of the Universe, and it is expected that the contribution to the quadrupole from this mechanism is negligible. The quadrupole is generated in this case by the Doppler effect acting upon the bulk motion of the plasma.

3. **Tensor Perturbations:** Gravitational waves, produced by inflation, introduce perturbations that stretch and compress space in orthogonal dimensions, [14]. Tensor fluctuations can be mathematically interpreted as perturbations transverse to the metric. This stretching and compression results in the wavelength of the photons changing, which

causes a change in the energy of the radiation and therefore the effective temperature of the region. Since different regions experience different degrees of stretching and compression, regions of different temperature are created and the temperature variation assumes a quadrupole form. The resulting polarisation following Thomson scattering is found to be divergence-free (B-mode or magnetic mode) or curl-free, E-mode.

The polarisation pattern that is measured on the sky can be separated into two components, the E and B-modes that were described above. The E-mode components has zero curl and is also referred to as the gradient mode. It displays no handedness and can have its origin in either scalar or tensor perturbations. The B-mode is caused only by vector or tensor perturbations, and is gradient free. It is also known as the curl mode and has handedness. If the normal vector of each polarisation is $\hat{\mathbf{n}}$, then under the transformation $\hat{\mathbf{n}} \rightarrow -\hat{\mathbf{n}}$ E-modes remain unchanged as long as the harmonic l is even, since they have $(-1)^l$ parity. B-modes will undergo a sign change however, as they have $(-1)^{l+1}$ parity. The E and B-modes therefore form the two components which the measured polarisation pattern consists of. These two components are shown on the right hand side of figure 1.1. The components of the polarisation which are parallel or perpendicular to the direction of travel of the wave are known as the E-modes of the polarisation. Any component which crosses the direction at forty five degrees is known as a B-mode of the polarisation. The overall polarisation pattern consists of a random superposition of these components. By examining the measured polarisation patterns, it is possible to decompose the polarisation map into each of these components. These components can then be analysed by means of decomposing them into spherical harmonics.

1.3.4 Analysis of the Anisotropies

To understand the analysis of the CMB from measurable data, a simple situation is considered involving the most intuitive observable that is associated with the CMB, its intensity as a function of frequency and direction on the sky, $\hat{\mathbf{n}}$. The CMB spectrum represents that of an excellent blackbody with an almost constant temperate T across the sky, [15], and so the intensity is usually described as a fluctuation in this temperature as a function of the direction on the sky, $\Theta(\hat{\mathbf{n}}) = \frac{\Delta T}{T}$.

A multipole expansion is the name given to a mathematical series that is used to represent a function that varies with angles, typically two angles projected onto a sphere. Often these series can be truncated, as the higher order multipole moments (terms in the expansion) contribute little and so by ignoring them one can still approximate the original function to a high fidelity. Most often the series is written as a sum of spherical harmonics, and so a function of two angles, say $f(\theta, \phi)$, can be written as

$$f(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} C_l^m Y_l^m(\theta, \phi), \tag{1.11}$$

where $Y_l^m(\theta,\phi)$ are the standard set of spherical harmonics and C_l^m are constant coefficients which are found from the function being expanded. Applying this to the temperature fluctuation field of the CMB, and assuming that the fluctuations are Gaussian, then the multipole moments of the temperature field

$$\Theta_{lm} = \int d\hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}})\,\Theta(\hat{\mathbf{n}}) \tag{1.12}$$

can be fully characterised by the power spectrum according to

$$< \Theta_{lm}^* \Theta_{l'm'} >= \delta_{ll'}\delta_{mm'} C_l. \tag{1.13}$$

The values of these spectra as a function of l are independent in a given realisation, and so analysis is typically performed in harmonic space which reduces to ordinary two dimensional Fourier analysis when small sections of the sky are considered, in which the curvature of the sky is negligible. l is known as the Fourier wavenumber in this case. In Fourier space the angular wavelength is given by $\theta = 2\pi/l$, and so large multipole moments (large l) correspond to small angular scales and vice versa. $l = 10$ is roughly 10 degrees on the sky and $l = 100$ is roughly 1 degree on the sky. Thus, to examine the fine structure of the CMB, the measured spectra must be sampled sufficiently to recover the higher order multipole moments. In the Fourier approximation, the variance of the field is given by $\int d^2 l C_l / (2\pi)^2$, and so the power spectrum is typically represented by

$$\Delta T^2 = \frac{l(l+1)}{2\pi} C_l T^2, \tag{1.14}$$

which corresponds to the power per logarithmic interval per wavenumber for $l >> 1$.

A typical CMB power spectrum is shown in figure 1.2, which displays the temperature anisotropy in terms of multipole moment, as measured by several instruments, including a theoretical model which is represented by the solid line. A careful study of the structure of this spectrum can reveal a significant amount of information. For example, the angular scale of the first peak in the spectrum can be used to determine the curvature of the universe. The second peak can be used to determine the reduced baryon density and the third peak can be used to investigate dark matter density. Information regarding the primordial density perturbations can also be found by examining the precise location of the peaks with respect to the multipole moment l. By examining the angular location of the peaks and taking the ratio of the locations of subsequent peaks, the precise types of density perturbations that were present can be investigated. Other information can be found at higher values of l. A similar analysis can be carried out on the E and B-mode maps to yield additional information which is necessary if inflation is to be understood, for example. The power spectra used in these cases

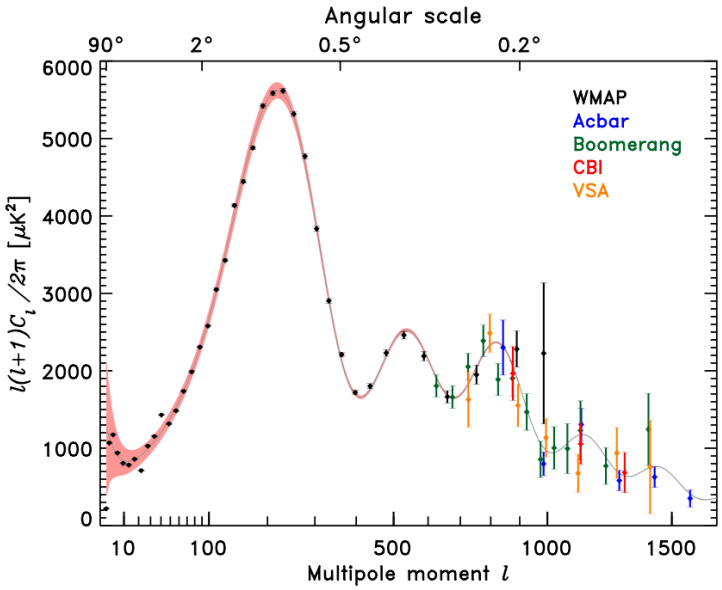

Figure 1.2: CMB power spectrum as a function of multipole moment, l, taken from a collection of WMAP Data Product images made available by NASA, [1].

are defined according to

$$< E^*_{lm} E_{l'm'} >= \delta_{ll'}\delta_{mm'} C_l^{EE}, \tag{1.15}$$

$$< B^*_{lm} B_{l'm'} >= \delta_{ll'}\delta_{mm'} C_l^{BB}, \tag{1.16}$$

$$< \Theta^*_{lm} E_{l'm'} >= \delta_{ll'}\delta_{mm'} C_l^{\Theta E}, \tag{1.17}$$

where C_l^{EE} is the set of constant coefficient used in the expansion of the function concerning the E-modes for example. The spectrum containing both E and Θ represents the correlation between the temperature fluctuations and the E-mode polarisation. The correlation of B with either of these two functions is not considered, as due to the handedness of the B-modes, such a correlation is not expected to exist. Of note is the fact that the polarisation signals (EE and BB) are smaller than the temperature fluctuation signal by approximately 1-2 orders of magnitude, as the polarised radiation was only produced towards the end of recombination. Since the relevant photons could not diffuse far before the end of recombination, their spectra have low values of l, as these correspond to large angular scales on the sky. Thus, to measure the polarisation signals, one must go to large values of l which is challenging as this requires sensitive, highly resolved measurements.

It is clear that in order to obtain information about the CMB and the origins of the universe, that it is necessary to increase the sensitivity and resolution of the measurements so as larger values of l can be considered, in addition to the structure of the spectrum being more clearly resolved at the values of l that measurements have previously retrieved. It is also necessary to

measure the polarisation of the CMB in order to determine information that cannot be found by considering the temperature alone. Future missions will focus on these goals.

1.4 CMB Missions

A brief description will now be given of several CMB missions that have taken place to date, and of a future mission that is planned. A future far infrared mission, SPICA, will also be outlined as it shares similar goals to those proposed for future microwave/sub-millimetre CMB missions. In order to accurately probe the conditions of the early Universe, it is necessary to examine the fluctuations, which manifest themselves as the multipole moments of the CMB radiation, on many angular scales. This, as outlined earlier, is known as the angular power spectrum of the CMB, and physically allows the amplitude of the structures to be examined on different mass scales. Taking the COBE mission as an example, [16], although it was the first mission to successfully detect the anisotropies, it was only able to do so for broad fluctuations as it had an angular resolution of 7°. This corresponds to large scale masses. Later missions were/are being designed to have a smaller angular resolution, allowing the fluctuations to be examined on a smaller scale. Many experiments have been designed to measure the CMB since its original detection by Penzias and Wilson in 1964, [17], and some of these missions will now be briefly examined.

WMAP

The **W**ilkinson **M**icrowave **A**nisotropy **P**robe (WMAP) satellite was launched by NASA in June 2001 and placed into a Lissajous orbit around the second Lagrangian point, [1]. This point was chosen so as to minimise the background noise from the Sun, Moon and the Earth, as well as to increase thermal stability. The mission lasted for 12 years, 11 months and 12 days.

The satellite spanned 5 bands, covering the 23-94 GHz range. In the highest frequency band, the W band, WMAP had an angular resolution of $< 0.25°$, which is a significant improvement on COBE. It also offered enhanced sensitivity and accuracy relative to COBE. This was achieved due to the fact that WMAP observed in five different bands, as doing this allowed signals from synchrotron radiation, free-free and dust emissions to be subtracted. This allows the true CMB spectrum to be determined.

Many significant results were obtained from the data taken by WMAP, [18]. The data indicated that the first stars were born much earlier than previously thought, approximately 200 million years after the Big Bang. It also showed the age of the Universe to be 13.7 billion years old. These results are consistent with those predicted by the theories of the Big Bang and inflation, and so serve as powerful evidence that these theories are correct. This agreement

implies that the Universe has a density that is very close to the critical density. The critical density is the density that the Universe requires in order to be flat, assuming a zero cosmological constant (Λ). An expression for the critical density is found by solving the first Friedmann equation for density, assuming $\Lambda = 0$ and also $k = 0$ (noting that $\Lambda + \Omega = 1$), where k is the curvature of the Universe which when set to zero results in a flat Universe (no curvature). The critical density is therefore given by $\rho_c = \frac{3H^2}{8\pi G}$, where H is the Hubble parameter. Comparing the actual density of the Universe to the critical density gives rise to a density parameter of $\frac{\Omega}{\Omega_c} \approx 1$, giving an approximately flat Universe which will therefore expand forever but at a rate which is constantly reducing, asymptotically approaching zero.

The Planck Surveyor

As discussed in section 1.1, the Planck Surveyor satellite was launched by ESA in 2009 in order to make full sky maps of the CMB anisotropies, [19]. With an angular resolution as low as 5 arcminutes, it offered higher sensitivity and resolution to that of WMAP, whilst also being sensitive to fluctuations in the polarisation of the CMB, which is critical. The polarisation fluctuations occur at a level of one to two orders of magnitude below those of the temperature fluctuations, and so the increased sensitivity and resolution of the instrument are vital.

The data obtained from Planck is still being processed, however results are regularly being published. Of note are the results indicating the composition of the Universe. The composition of the Universe was measured by Planck to be 4.9% normal baryonic matter, 26.8% dark matter, and 68.3% dark energy. Planck also derived a value for the Hubble parameter from its measurements, which was found to be $H_0 = 67.3 \pm 1.2$ kms^{-1}Mpc^{-1} [20].

QUaD-originally known as QUEST

The determination of the polarisation properties of the CMB is a problem that is receiving considerable attention from all over the world. The typical work order that is followed when carrying out such measurements is to make initial measurements from the ground and to then use satellite missions to follow on from these initial measurements.

QUaD is an example of such a ground based experiment, [21]. QUaD stands for QUEST (**Q** and **U** **E**xtragalactic **S**urvey **T**elescope) **a**t **D**ASI (**D**egree **A**ngular **S**cale **I**nterferometer). QUaD carried out observations for 3 years, making its first science observations in May 2005. QUad is stationed approximately 1 km from the geographic South Pole. It was located here due to the fact that it is one of the best ground based sites for millimeter observations, owing to the low water content of the atmosphere, meaning that the millimetre radiation does not get attenuated strongly by the atmosphere. QUaD is famous for the first detection of the CMB polarisation, and has made the most accurate measurements of them yet. NUI Maynooth was involved with this project through a collaboration with Cardiff and Stanford Universities.

QUaD mapped a region of the sky with a low foreground signal at frequencies of 100 and 150 GHz. The benefit of observing at such frequencies is that they cover the band in which astronomical signals from sources other than the CMB are at a minimum, whereas the signal associated with the CMB is at its strongest. QUaD was designed to examine the polarisation of the CMB. This signal is one to two orders of magnitude lower than the temperature anisotropy, so it is necessary to observe at the frequencies at which the signal is the strongest. Using the data measured by QUaD, it was possible to calculate angular power spectra for the E-modes and to place an upper limit on the B-modes of the CMB polarisation. This was achieved by measuring the 2×2 tensor that contains the temperature and polarisation anisotropies of the CMB. Decomposing this tensor results in a scalar field which contains information on the E modes. Information about the B modes is contained in a pseudo-scalar field.

QUaD made use of an on-axis Cassegrain telescope design. The angular resolution of QUaD is determined by the diameter of the primary mirror in the Cassegrain configuration. The 2.6 m primary mirror gives an angular resolution of 4.6 arcmin at a frequency of 150 GHz. This resolution is extremely important as it determines the maximum value of l that can be realised in the multipole expansion when the angular power spectrum is being constructed, as outlined earlier.

The detection system consisted of two cryogenic re-imaging lenses with anti-reflective coating, which fed a focal plane array of 31 pixels. Each pixel consisted of a corrugated horn which fed two Polarisation Sensitive Bolometers (PSBs) (kept at less than 1 K) that detected the power incident on them by measuring changes in resistance that are due to the heating effect of the incoming radiation. The two PSBs were arranged so as to be sensitive to both orthogonal linear polarisations as the sky is scanned. The total intensity, which corresponds to the temperature, was obtained by summing the output of the PSBs, whilst the difference of the two outputs allowed the removal of the equally split unpolarised light, leaving only the polarised component. Each of the 31 pixels was monochromatic, with 12 dedicated to the 100 GHz observing band, with the remaining 19 observing at 150 GHz.

QUaD successfully measured the E-mode polarisation, achieving a very strong detection. An upper limit was also placed on the B-mode polarisation, with the detection being at a much lower level than that of the detection of the E-mode polarisation, [22–24]. This is consistent with the standard cosmological model which predicts that the polarisation should be dominated by the E-mode contribution, however the map of the B-mode component is also consistent with that of random chance which could have been caused by noise in the instrument. In either case, QUaD provided encouraging results in terms of successfully separating the polarisation contributions by achieving such a strong E-mode detection.

BICEP2

The process of inflation must have produced a Cosmic Gravitational-Wave Background (CGB) at very early times. These gravitational waves would have made a unique and measurable contribution to the polarisation of the CMB, and so if this signature could be measured then it would provide indirect evidence in favour of inflation. The gravitational waves are tensor fluctuations, as described earlier, and so their contribution to the CMB will have a non-zero curl component which will result in a contribution to the B-mode polarisation. BICEP2 (Background Imaging of Cosmic Extragalactic Polarisation) is located at the south pole, where the water content of the atmosphere is sufficiently low to avoid significant attenuation of the CMB radiation, and operated from 2010 to 2012 with the purpose of measuring (or placing an upper limit on) the B-mode polarisation of the CMB in order to detect the signature from the gravitational waves that would provide evidence of inflation, [25].

BICEP1 was a prototype for future experiments, with each subsequent experiment being more sensitive to the B-mode polarisation, [26]. Whereas BICEP1 had 98 detectors, BICEP2 uses 512 antenna coupled detectors on a focal plane that is the same size as that of BICEP1 (achieved by an improved detector design), leading to an increase in sensitivity of approximately one order of magnitude. The detectors are polarisation sensitive TES bolometers with the 512 detectors arranged in pairs that are sensitive to orthogonal polarisations giving rise to 256 pixels, detecting radiation at 150 GHz. The focal plane is kept at 250 mK by means of a sub-Kelvin generator, allowing background limited sensitivity. The telescope used by BICEP2 is of the small aperture (26 cm), refracting type, which allows for excellent control of systematic effects. The optics pertaining to the telescope are cooled to 4 K.

The data from BICEP2 is currently being analysed. In March 2014 it was announced that B-modes supporting inflation (showing a tensor to scalar ratio of $r = 0.2$) had been detected, however it has since been claimed that the polarisation referred to in these findings may have been the result of light that had scattered off dust located between stars in the Milky Way due to errors in foreground subtraction in the analysis. The data from BICEP2 is being analysed on an ongoing basis, [27].

COrE-Cosmic Origins Explorer

COrE was a proposed class M space based satellite mission aiming to carry out a high precision full sky survey of the polarisation at microwave and sub-millimetre frequencies. In doing so it would have allowed the probing of the cosmic origins, as well as investigations into neutrino masses and the origin of the Universe, [28–30]. A similar mission was also proposed, called PRISM (Polarised Radiation Imaging and Spectroscopy Mission), [31]. Both PRISM and COrE were rejected, and so now a new mission is proposed, COrE+, which combines the capabilities of both COrE and PRISM. A number of different designs and science goals are currently being

considered. The COrE+ proposal will now be considered, as it is extremely relevant for future missions.

COrE's predicted performance (relevant, as the COrE+ proposal includes COrE capabilities) is such as to allow unprecedented polarisation sensitivity and foreground separation capabilities, allowing the sensitivity of B mode measurements to be increased by two orders of magnitude compared to current missions. This will allow the angular power spectrum of these features to be expanded to much higher orders than are currently available, meaning the initial conditions of the Universe can be examined at a whole new level of detail. Non-Gaussianity will also be investigated using this data, allowing the existence of non-linear physics during inflation to be proven or discounted. It will also allow the sum of the neutrino masses to be measured to within 0.03 eV.

The higher frequency channels (where the total band is 45-795 GHz) will allow the investigation of star formation and the connection between this phenomenon and magnetic fields. The galactic magnetic field will also be made examinable, as well as the properties of interstellar dust. This latter point means that it will be possible to determine the initial conditions that existed for stars to form in the more diffuse portions of the interstellar medium. It will also enable the analysis of a large number of new galactic and extragalactic polarised point sources.

Future missions with COrE capability, such as COrE+ will house the system as a single instrument on one spacecraft. Like Herschel, it will orbit around the second Lagrangian point of the Sun/Earth system. The instrument will be capable of measuring a full sky survey every 6 months.

The COrE+ proposal, and so future missions with COrE capabilities, builds on the knowledge acquired during the design and lifespan of both Planck and Herschel, and was proposed in response to the Cosmic Vision 2015-2025 call, placed by the European Space Agency. The science goals of this instrument will be achieved by using 6384 detectors to observe at a wide range of bands with central frequencies ranging from 45-795 GHz. A typical bandwidth for each observing channel is 15 GHz. This will allow better quality maps to be developed than those returned by Planck, despite the fact that the signal being measured by COrE, the polarisation signal, is considerably weaker. These results will be achieved by having twice the frequency resolution and an increase in sensitivity of approximately a factor of 30 when compared to Planck, even though the angular resolution and frequency coverage are similar. This improvement is made possible by the increase in the number of detectors, as outlined above, but also by the fact that each detector will operate in the region of the photon shot noise limit.

Another reason for the increased performance is that systems with COrE capability will be designed with an emphasis on rejecting systematic error within the detector design. Planck used two different detectors to measure the intensity of the orthogonal components of a linearly polarised signal. By subtracting these intensities, the polarisation of the signal could be

measured. This introduced errors to the measurement. Proposed future missions will measure the polarisation by modulating the polarisation signal by means of a rotating half-wave plate, situated between the sky and the detector horns. The result of this is that the polarisation will be measured directly, reducing the inherent errors that can occur when an indirect process is used that involves processing using an optical system. The direct measurement also brings with it time saving, prolonging the length of the mission.

SPICA/SAFARI

The SPace Infrared telescope for Cosmology and Astrophysics (SPICA) is a proposed infrared space telescope that is aiming for launch in 2025, with a mission duration of 3 years. This type of future mission proposal will be the next generation of future far IR telescope after the Herschel space telescope. It is intended to also orbit the second Lagrangian point of the Sun/Earth system. The main science objectives of such a mission will involve carrying out astronomy at far infrared wavelengths to enable progress in the fields of cosmology and astrophysics, building on Herschel. Science goals will involve the birth and conditions necessary for the evolution of galaxies, stars, and planetary systems, and the evolution of matter, [32–34]. It will also allow research into the life cycle of interstellar dust. SPICA is proposed to have a specific focus on planetary system formation and galactic evolution.

It is difficult to observe the formation of planetary systems, as the regions in which this occurs are obscured by dust. An analysis of the gases in this region, such as O, OH and H_2O lines will allow the boundary conditions for planetary formation to be determined. In order to achieve this, it is necessary for future instruments to be sensitive to these spectral lines, and SPICA is designed for sensitivity in these regions. Young star systems contain dusty disks which could give rise to terrestrial planets. Analysis of this dust via far infrared (FIR) spectroscopy (using SAFARI) would then possibly reveal information about the processes that give rise to such planets. The sensitivity of the instruments on the proposed SPICA mission is sufficient to allow a detailed mapping of the mineral content of nearby star systems, in addition to the distribution of grain sizes in these regions. SPICA would also be able to determine the role of stellar winds in the evolution of massive stars in addition to determining the life cycle of dust, i.e. the formation of dust around evolved stars to its distribution throughout the ISM, all the way to its destruction by supernova generated shock waves. The coronagraph proposed for SPICA will also allow imaging in the 3-27 μm range of young gas giants, allowing stellar evolution models to be assessed.

A significant quantity of the cosmic infrared background radiation is believed to originate from the integrated luminosity of high red-shift galaxies, [35]. These galaxies present issues for current galactic evolution models, and so must be studied in order to resolve these issues. As the instruments proposed for SPICA allow for observations in the mid and far infrared, it would be possible for SPICA to analyse star formation and black hole accretion for higher and

higher values of redshift, z. This equates to looking back in time at the oldest galaxies, and so it would allow a study of galactic evolution. These observations would allow the examination of the connection between starbursts and active galactic nuclei. SAFARI would be used to carry out FIR spectroscopy in order to characterise the gas from these regions in terms of temperature, density and composition, as well as allowing the characterisation of the radiation emanating from AGN and starbursts.

The necessity for previous far infrared missions to use liquid cryogens to cool the telescope and instrument to less than 5 K has meant that these missions were limited to aperture sizes of less than one metre. This meant that the spatial resolution of the systems was poor, with limited sensitivity. The Herschel mission provided a compromise by employing a 3 m mirror. This resulted in an increase in spatial resolution, but with limited increase in sensitivity relative to previous missions in the 55-210 μm range, owing to the fact that the larger aperture resulted in a cooled temperature of 80 K. By using mechanical cooling, SPICA might employ a 3 m class telescope that can be cooled to less than 6 K, resulting in the elimination of non-astronomical photon noise. This, combined with highly sensitive detectors, will allow SPICA to achieve a level of sensitivity that is limited only by the sky in the 5-210 μm range. The 3 m primary mirror is monolithic and the tolerancing on the telescope and mirrors is such that the system can realise diffraction limited performance at the shortest wavelength of 5 μm. The revolutionary cooling systems used to achieve this are made up of a combination of passive and active cooling systems. The passive cooling system uses radiators in combination with thermal and solar shields, whilst the active cooling systems use several mechanical coolers.

The telescope itself is proposed to be a Ritchey-Chretien design. As previously stated, it is of diameter 3 m and cryogenically cooled to 5 K. The field of view is of radius 5 arcmin. The proposed telescope will contain seven instruments.

1. MIRACLE: A camera and spectrophotometer for operation in the mid infrared, $\lambda = 5 - 38$ μm.

2. SAFARI: An imaging Fourier Transform Spectrometer for operation in the far infrared, $\lambda = 34 - 210$ μm.

3. MIRHES: A high resolution spectrometer for use with $\lambda = 4 - 18$ μm.

4. MIRMES: A high resolution spectrometer for operation in the mid infrared, $\lambda = 10.32 - 36.04$ μm.

5. SCI: A coronograph designed for operation in the mid infrared, $\lambda = 5 - 27$ μm.

6. BLISS: A spectrometer for operation in the far infrared/sub-millimetre, $\lambda = 38 - 430$ μm.

7. FPC: dedicated focal plane cameras

These instruments will generate a total amount of science data of 350 Gbit/day. Of interest in this thesis is the SAFARI instrument, which will now be considered.

SAFARI (Spica FAR infrared Instrument) is a proposed Mach-Zehnder imaging Fourier Transform spectrometer that covers the 34-210 μm band. It gives a spectral resolution of R 10 to 10000, offering a field of view of 2 arcmin by 2 arcmin. Such spectrometers are capable of spatial multiplexing, allowing a high mapping speed. It is also not difficult to implement imaging capabilities with such spectrometers. SAFARI may be broadly considered to consist of the following elements.

- A front end optical telescope, cooled to 4.5 K.
- The Mach-Zehnder interferometer, cooled to 4.5 K. An optical path difference of up to 124 mm can be realised in the spectrometer by means of a magnetic bearing mechanism.
- A back end optical camera telescope, cooled to 4.5 K.
- A focal plane array with accompanying detector array that allows for observation in three bands. The three bands are 35-60 μm with 64×64 pixels, 60-110 μm with 38×38 pixels and 110-210 μm with 20×20 pixels. The pixel design that will be considered in this thesis is used in the 35-60 μm array.
- Separate electronic systems to control the readout of the TES detectors, the control of the spectrometer, the control of the 50 mK cooling system which is used for the detection focal plane, and a final electrical system to control the management and interfacing of the instrument with the satellite.

The focal plane array will be of interest in this thesis. The camera focusses the incident image onto an array of (pyramidal) feed horns which feed an array of detector cavities. Two types of detectors are currently under consideration for use in the SAFARI system. Sensitivity is critical in the context of the scientific goals of the instrument. The two systems being considered are Transition Edge Sensor (TES) bolometers and photoconductors. The TES bolometer case is assumed in this thesis. They have several advantages compared to photoconductors. Their spectral response is such as to allow good response at short wavelengths (40 μm) where the performance of the photoconductors is poor. The absorbers used with TES bolometers are also of low cross-section area compared to that of the photoconductors. This means that the TES bolometers are less likely to be effected by unwanted cosmic rays due to the lower hit rate.

The proposed SAFARI instrument will use the mid/far infrared portion of the electromagnetic spectrum to study the processes that drive galactic evolution. These processes are 'hidden' behind dust that is opaque at other frequencies, which is why the mid/far infrared is so powerful when these processes are required to be studied. Currently this proposal will be aimed at a future M class mission.

Future CMB/Far-Infrared Technology

As shown above, the emphasis for future CMB and far IR missions is to increase the sensitivity and resolution of the measurements being used to generate the sky maps. Regardless of what signal is being measured, be it the temperature anisotropy signal or the polarisation signal, it is vital to take more sensitive measurements. The angular power spectrum that is generated from this will have higher values of l contained in it then ever previously measured, and so this will reveal entirely new levels of detail about the Universe and its history.

The CMB polarisation signal is currently a major driving force, and to realise these improvements in performance, it is necessary for the technology used in the detector systems to improve. The angular resolution and sensitivity of the systems must be improved, whilst maintaining polarisation sensitivity in the detected signal. One element that must be examined in detail is that of the feedhorn antennas used. To measure the signal correctly, these horns must realise high performance in terms of several performance metrics. They must be sensitive to low levels of radiation, and so their return loss must be low. They must be polarisation sensitive, so the distinction between the cross and co-polar fields must be high. The horns must also be robust so as to be independent of azimuthal rotations, so the co-polar beam must be symmetric. It is also necessary to use thousands of such horns, arranged at the focal plane of the instrument. This increases the sensitivity of the instrument, but also allows more measurements to be taken quickly, resulting in extended mission durations. Naturally, in terms of structural integrity, launch limitations and mission duration requirements, the weight and thermal penalties of the horns must be low. The corrugated horns currently utilised for this purpose are not therefore the ideal choice for such missions. They are also costly, and when many thousand are required, this cost can become prohibitive. One aim of this thesis is to develop an optimisation routine that will allow a smooth-walled horn to be designed (using readily available software and a desktop pc) within hours, that satisfies these requirements at the expense of the performance requirements that are not so critical. Although for the purposes of this thesis feed horns are considered to couple to the detectors, other pixel designs may satisfy the performance requirements. Examples of such alternative systems are bare arrays that are phased together and lens antennas or other refractive elements instead of feed horn antennas.

In order to maximise the amount of the desired signal that can be detected, the throughput of the horn antennas should also be maximised. This involves using few or many-moded horns. In order to use these in space missions, it is necessary to characterise their performance carefully and to use this information to further develop and refine simulation techniques for their design. Measuring multi-moded horns properly currently presents a challenge, and so this thesis will also aim to examine this issue.

Many future detector systems will be based on integrating cavities which house detectors

such as bolometers. Each pixel of such a system will consist of a horn antenna feeding such a cavity. These structures need to be analysed and understood in detail. Currently, a prototype of such a pixel for the SAFARI instrument is being experimentally tested by the Space Research Organisation of the Netherlands (SRON). These pixels include a gap between the horn and the cavity, which is capable of acting as a mechanism to allow power to escape from the system. Such systems will be examined in this thesis, with the aim of including such a gap to model its effects on the system. This is necessary if the performance of the pixel is to be truly understood. The Planck bolometers were also housed in a similar cavity and so the analysis presented here is also relevant to that mission.

1.5 Overview of Thesis and Author's Contribution

Here, a brief overview of each chapter is given, in addition to the author's specific contribution to each chapter, where appropriate.

Chapter 2

Chapter two gives an introduction to the mode-matching technique for cylindrical waveguides, in addition to details on how this technique is implemented in the mode-matching software known as SCATTER, developed at NUI Maynooth in the Terahertz Space Optics Group. This software forms a basis for the majority of this thesis, and so it is examined in detail in this chapter. Details are also given for another software package that is currently being developed within the group; the Maynooth Optical Design and Analysis Laboratory (MODAL). This is a software package allowing the design, analysis and testing of quasi-optical systems using physical optics, mode-matching and Gaussian optics.

Chapter 3

Two significant modifications were made to SCATTER by the author, in order to increase the efficiency of the code. One of these modifications was the implementation of a cylindrical geometry mode-matching calculation in Matlab that utilises a dynamic number of waveguide modes based on the radius of each section of waveguide that forms the overall structure and the frequency of operation. This approach is used to significantly reduce the computational time required for a given waveguide structure when compared to the time required using the traditional approach of a fixed number of modes. The other modification was the implementation of a singular value decomposition (SVD) based method of field pattern calculation. The SVD technique is applied rigorously to further minimise the computational required when compared with the traditional field calculation. In this chapter, the underlying theory for each modification is developed, and details are given on the implementation of each modification, along with

the verification of each method. The Matlab version of SCATTER, with these modifications made, is shown to run in a fraction of the time that is required using traditional analysis.

Chapter 4

One goal of this thesis is to develop an efficient method to obtain a horn antenna design subject to the particular design requirements of a CMB/Far-IR mission. Following an investigation of several optimisation techniques, the global optimisation techniques of the genetic algorithm and simulated annealing were selected by the author. They were applied to several horn designs that form the basis of the optimisation process and these designs were subjected to the optimisation algorithms using a variety of optimisation metrics that are of interest to CMB/Far-IR missions in order to establish which optimisation algorithm/horn design/performance metric gives the most satisfactory result in terms of the requirements of such future missions. It was found that using a piecewise conical profile horn with the genetic algorithm gave rise to the designs with the most potential, so in this chapter, this horn design is allowed to be modified by the genetic algorithm subject to various performance requirements in order to see which yields the best results and to therefore establish the limitations of this design approach.

Chapter 5

Using the findings of chapter 4 along with a particular set of requirements for a future CMB/Far-IR mission, a horn was designed to meet these requirements in addition to meeting the requirements of a short design time on a standard desktop PC using only standard Matlab algorithms. The resulting horn was simulated extensively to determine its performance and the design was manufactured by Rutherford Appleton Laboratories in the UK. The results of the measurement campaign carried out on the horn are presented and compared to the simulated results and design requirements. The alignment process for the measurement testbed is also described.

Chapter 6

In this chapter the inclusion of a free space gap within the framework of the mode-matching technique is considered, along with the other elements that are required in order to simulate a single SAFARI like pixel. The theory for implementing each element was developed by the author and is presented. The associated code was developed and verified by the author by simulating a simple cavity structure and comparing it to results obtained from a full finite difference time domain electromagnetic solver. The verified code is then applied to a model of the experimental setup in use at the Space Research Organisation of the Netherlands (SRON), where a prototype SAFARI pixel is being analysed, to characterise the measurements, including the effect of the free space gap.

Chapter 7

In this chapter, the theory behind replacing the standard source of illumination used by SCATTER with a plane wave that can be used to provide off-axis illumination is developed by the author and presented. The code was implemented and applied to a back to back corrugated horn in both single and few-moded configurations for various angles of illumination. The results of these simulations are analysed in detail and the correct functioning of the code is verified by this.

Chapter 8

In this chapter, the development and testing of a 100 GHz Golay cell based measurement testbed is described, in which the author was responsible for the development of the code for controlling the scanning mechanism. The system was tested by measuring the beam patterns of several horns, with the results obtained from measuring a corrugated horn presented in order to verify the correct functionality of the scanning and processing code. The author designed a smooth-walled conical horn antenna for few-moded use at 100 GHz. This was manufactured at the University of Manchester and was measured by the author in a back to back configuration for various angles of illumination using the Golay cell testbed and a corrugated horn as the source. Equivalent simulations were carried out on this horn arrangement using the code developed in chapter 7 and these results are compared with the measured results, further verifying the code.

Chapter 2

Mode Matching Technique: SCATTER

In this chapter, a brief introduction will be given to the modal matching technique, as well as the implementation of this technique in SCATTER; software developed in-house [36,37] that allows the performance of waveguide structures to be evaluated. At the wavelengths encountered in the microwave and sub-millimetre portions of the spectrum, these waveguide structures primarily appear in two forms, cylindrical or rectangular, with the structure being hollow and conducting in both cases. While it is possible to apply the modal matching technique to either geometries, cylindrical geometries are considered primarily in this thesis. This approach forms part of an integrated quasi-optical system design and analysis package that is also being developed in-house, **M**aynooth **O**ptical **D**esign and **A**nalysis **L**aboratory (MODAL), [38–40]. This software package will be used for portions of the work presented in this thesis, and so in this chapter it will also be briefly presented.

Cylindrical waveguides are excellent candidates for use in sub-millimetre systems, as the circular symmetry means that the waveguide has no preferred orientation. This in turn means that it can receive both polarisations of the orthogonal waveguide modes.

Horn antennas can be analysed using the mode matching technique when the feed horn is thought of as a series of discrete waveguide sections of circular cross-section. These sections, when cascaded together, recreate the complete horn. By analysing each section, and by allowing the output of one section to be the input for the next, it is possible to analyse the horn antenna as a whole by using the modal matching technique. In this way, both single-moded (TE_{11} only) and multi-moded (TE_{11} and higher) horns can be modelled. (TE and TM refer to the cylindrical forms of the transverse electric and transverse magnetic waveguide mode sets respectively. These will be defined and discussed in detail in sections 2.1.1 and 2.1.2). This effective discrete sampling of the horn means that the representation of the structure is only approximate. Finer sampling (sections of smaller length) allows more accurate representation of the horn, however this comes at the expense of computational time and power. We have found that it is sufficient for most applications to use a number of sections such as there are ten sections per wavelength [36], ensuring an accurate representation of the horn. The modal

matching technique shall now be examined.

2.1 The Principles of Mode-Matching: Scattering in Cylindrical Waveguides

Consider initially the junction between two cylindrically symmetric waveguides, in which the second waveguide has a larger radius than the first. To model this junction, the reflection and transmission coefficients at the junction must be found for all modes incident from either side. To achieve this, the electric and magnetic fields in the plane of the junction are expanded, using an infinite series of the normal propagating and evanescent modes appropriate to each waveguide as a basis set, which will be discussed in more detail later. The number of modes that physically propagate through the section will be determined by the cross-section diameter of the waveguide and the frequency of the radiation.

In order to satisfy Maxwell's equations, a clear requirement, we must have that the transverse components of both the electric and magnetic fields are continuous across the junction between the two waveguides. Using this, the transmission and reflection coefficients for each junction can be found. In this case we would assume some set of input coefficients for each mode and then use these and the required boundary conditions in order to determine the output coefficients of the section for each mode. If subsequent sections existed, these output coefficients would be used as the input set for the next section. For the analysis of a complete horn, this process would be repeated all the way through the horn, starting with some initial set of input modal coefficients that represent the initial field. The relationship between input and output mode coefficients for a given section is given by the scattering matrices for that junction. Each junction can have its scattering matrices calculated, and then at the end these can be cascaded to determine the overall scattering matrices for the horn, relating the input field to the output field (electric or magnetic).

Clearly, the number of modes to be considered in the calculation must be finite, and therefore determined. The number of propagating and evanescent modes can be found by carrying out a stability analysis of the results (increasing the number of modes until the predicted performance converges to a stable solution). This will yield the minimum number of modes required to describe the overall horn adequately (i.e. the number of modes required to describe the *widest* section). The simplest solution is to use this number of modes (plus several, perhaps two or four, additional modes in order to allow for evanescent modes) for each junction, as this keeps the matrices square, and simplifies the arithmetic. For a more computationally efficient analysis, the minimum number of modes required for each section (with several additional modes to account for evanescence again considered) can be calculated and used based on some predetermined accuracy level, however this would result in rectangular matrices. Although

this would speed up the computation, the mathematics would be more complex, however this will be discussed later. The number of modes required to accurately represent a waveguide section can be taken to be approximately ten times the ratio of the waveguide diameter to the wavelength, [41]. Alternatively, the number of modes can be varied until there is negligible difference between subsequent coefficient values, however this is naturally more computationally intensive and time consuming. This must all be carried out using at least ten sections per wavelength. With these measures in place, the calculation process can be assumed to be convergent and the results trusted. This thesis will at various points consider cavity structures, which are realised within the mode-matching technique as a sections or sections of waveguide, with a small section at the end which cannot support any modes. This effectively closes the cavity. Since resonances can readily occur within cavities, resulting in power being scattered into modes which may not normally be excited, additional modes should be considered when cavities are being examined in order to account for this possibility.

Cylindrical waveguide modes are in use due to the circular symmetry of the waveguide sections in question, and so pure TE and TM modes can be defined which form a complete set of solutions to the wave equations for waveguides of circular cross-section. For a waveguide of finite length, the orthogonality relationship between the magnetic and electric modes can therefore be expressed as

$$\iint_S \mathbf{e}_n \times \mathbf{h}_m^* d\mathbf{S} = P_{nm}\delta_{nm}, \tag{2.1}$$

where P_{nm} represents the total power carried by a mode of corresponding subscript. $\mathbf{e}_n$ and $\mathbf{h}_m^*$ represent the transverse electric and magnetic fields for two linearly independent solutions of Maxwell's equations, and the asterisk represents the complex conjugate, [36]. This shows that for a lossless guide, the power flow within it is represented by the sum of the power carried by each individual mode, as expected. These figures should remain constant in a section of uniform radius, as there will be no scattering between modes and no resistive losses. Modal matching is a more universal method of simulating waveguide structures than other methods (such as the uniform surface impedance approximation) due to this, as it does not require regular fields at all points, as the fields along the horn are well defined and higher order modes are automatically considered, subject to the correct number of modes being chosen.

2.1.1 Modal Content of Cylindrical Waveguides

It is assumed [37, 42] that the cylindrical coordinate system is defined so that the cross-section of the waveguide lies in the x-y plane, leaving the z direction as the propagation axis, as shown in figure 2.1. This cylindrical system is easier to operate within than a spherical system, due to the resulting form of the field equations, and so the approximation of the horn to a series of cascaded waveguide sections is used to allow this form to be used, as the cylindrical system is the most natural representation of such a system. A waveguide that is homogeneous and uniform

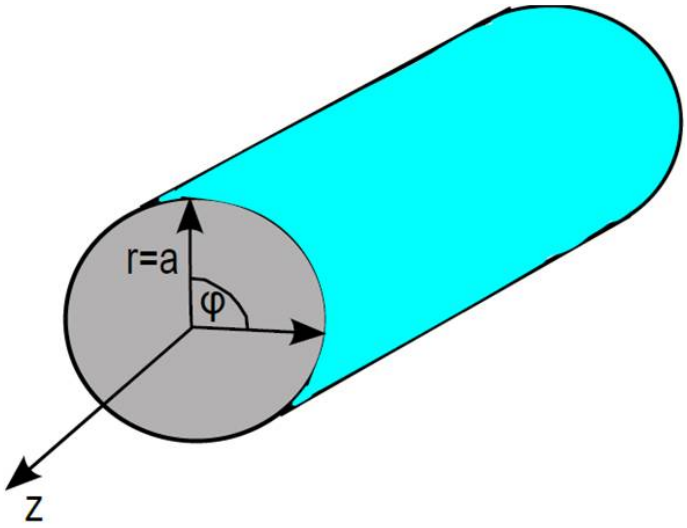

Figure 2.1: Cylindrical coordinate system used in SCATTER for circular waveguides

in cross-section supports two orthogonal mode sets. These are known as TE and TM modes, transverse electric and magnetic modes, where respectively, E_z=*0* and H_z=*0*. This means that the electric and magnetic fields present in such waveguides can be decomposed into axial components, E_z and H_z, and transverse components E_t and H_t. The transverse components are a function of r, ϕ and z, the axial distance, azimuthal and propagation distance respectively. As such the field can be fully described at any point in a particular plane. Each waveguide can have an infinite number of evanescent modes present, but the number of (real) propagating modes is dependant on the diameter of the waveguide. Each mode has a cut-off frequency below which it does not propagate in a waveguide of a certain radius. This relationship will be illustrated when the form of these modes has been defined. Higher order modes have higher cut-off frequencies, and larger waveguides support higher frequencies, so the larger the waveguide, the more modes are supported. In circular cylindrical coordinate systems, the equations for TE and TM modes that relate the transverse and longitudinal components of the electric and magnetic fields are given below, [43]

$$E_r = \frac{-j}{k_c^2}\left[\beta\frac{\partial E_z}{\partial r} + \frac{\omega\mu}{r}\frac{\partial H_z}{\partial \phi}\right] \tag{2.2a}$$

$$E_\phi = \frac{j}{k_c^2}\left[\frac{-\beta}{r}\frac{\partial E_z}{\partial \phi} + \omega\mu\frac{\partial H_z}{\partial r}\right] \tag{2.2b}$$

$$H_r = \frac{j}{k_c^2}\left[\frac{\mu\epsilon}{r}\frac{\partial E_z}{\partial \phi} - \beta\frac{\partial H_z}{\partial r}\right] \tag{2.2c}$$

$$H_\phi = \frac{-j}{k_c^2}\left[\mu\epsilon\frac{\partial E_z}{\partial r} + \frac{\beta}{r}\frac{\partial H_z}{\partial \phi}\right], \tag{2.2d}$$

where

$$k_c^2 = \gamma^2 + k^2 = k^2 - \beta^2, \tag{2.3}$$

and k_c is the cut-off wavenumber of the mode. γ is known as the propagation constant, which is given by $\gamma = j\beta$ when the wave is propagating. β is real if no attenuation occurs, and for evanescent modes is pure imaginary.

The situation is highly symmetric, so $\mathbf{E}_z$ and $\mathbf{H}_z$ can be separated and expressed as the product of a function of the coordinates r and ϕ, with a phase term that is a function of propagation distance and time, $\exp\left[j\left(\omega t - \beta z\right)\right]$. When the wave equation (in the cylindrical form) is solved for modal solutions, the solutions are known as modes and are products of the Bessel functions and trigonometric functions, [37]. This will now be examined.

2.1.2 TM and TE Mode Analysis

We begin by examining the TM modes. As shown in [43], the Helmholtz equation can be reduced to

$$\nabla_t^2 E_z = \frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial E_z}{\partial r}\right) + \frac{1}{r^2}\frac{\partial^2 E_z}{\partial \phi^2} = -k_c^2 E_z. \tag{2.4}$$

This equation is solved by means of the separation of variables, resulting in the following expression for E_z:

$$E_z\left(r,\phi\right) = AJ_n\left(k_c r\right)\begin{pmatrix}\cos n\phi \\ \sin n\phi\end{pmatrix}\exp\left(j\left(\omega t - \beta z\right)\right), \tag{2.5}$$

where $J_n\left(k_c r\right)$ is a decaying oscillatory function known as a Bessel function of the first kind, of order n. Both of the orthogonal modes have the same radial variation, and the sin and cos terms account for the azimuthal variation in E_z. The sections are assumed to be perfect conductors, and so the field must satisfy the boundary conditions that $E_z = E_\phi = 0$ on the walls of the waveguide, at which r is equal to the radius of the circular section of waveguide, a. Using this condition and equations 2.2a-2.2d, it is found, [36], that for TM modes of order n and l,

$$E_r = Z_{TM}H_\phi = \frac{-j\beta}{k_c}AJ_n'\left(k_c r\right)\begin{pmatrix}\cos n\phi \\ \sin n\phi\end{pmatrix}\exp\left(j\left(\omega t - \beta z\right)\right) \tag{2.6a}$$

$$E_\phi = -Z_{TM}H_r = \frac{-j\beta n}{k_c^2 r}AJ_n\left(k_c r\right)\begin{pmatrix}-\sin n\phi \\ \cos n\phi\end{pmatrix}\exp\left(j\left(\omega t - \beta z\right)\right), \tag{2.6b}$$

where $J_n'\left(z\right) = dJ_n\left(z\right)/dz$ is the derivative of the Bessel function $J_n\left(z\right)$. Z_{TM} represents the ratio of the transverse electric field to the transverse magnetic field for a given mode, also known as the characteristic waveguide impedance, given by equation 2.7

$$Z_{TM} = \frac{\beta}{\omega\epsilon} = \frac{1}{\omega\epsilon}\sqrt{k^2 - k_c^2} = Z_0\sqrt{1 - \left(\frac{p_{nl}\lambda}{2\pi a}\right)^2}, \tag{2.7}$$

and Z_0 is the characteristic impedance of the medium, given by $\sqrt{\frac{\mu}{\epsilon}}$. If the boundary condition is to be satisfied, then the argument of the Bessel function at $r = a$ must be equal to one of its roots, in other words

$$k_c a = \frac{2\pi a}{\lambda_c} = \omega_c \sqrt{\mu\epsilon} a = p_{nl}, \tag{2.8}$$

where $J_n(p_{nl}) = 0$, so p_{nl} is the l^{th} root of $J_n(z)$.

Using this form of the TM modes, the cut-off wavelength can be expressed as

$$\lambda_{TM_{mn}} = \frac{2\pi r}{p_{nm}}. \tag{2.9}$$

If the wavelength of the radiation in the waveguide is less than the cut-off wavelength for a particular mode, then that mode will be supported within the waveguide section and it will propagate. Equation 2.9 can be easily rewritten in terms of frequency or wavenumber, which may be more intuitive to work with at times.

There are an infinite number of orders possible for Bessel functions, and for each order there are an infinite number of roots of the functions, so it follows that there are an infinite number of possible modes. We can re-express the transverse components of the TM modes (when C_{nl} is set to $(-j\beta A_{nl})/k_c$) as

$$e_{nl,r} = \frac{-j\beta}{k_c} A_{nl} J_n'(k_c r) \begin{pmatrix} \cos n\phi \\ \sin n\phi \end{pmatrix} = C_{nl} J_n'(k_c r) \begin{pmatrix} \cos n\phi \\ \sin n\phi \end{pmatrix} \tag{2.10a}$$

$$e_{nl,\phi} = \frac{-j\beta}{k_c^2 r} A_{nl} J_n(k_c r) \begin{pmatrix} -n \sin n\phi \\ n \cos n\phi \end{pmatrix} = C_{nl} \left[\frac{n J_n(k_c r)}{k_c r}\right] \begin{pmatrix} -\sin n\phi \\ \cos n\phi \end{pmatrix} \tag{2.10b}$$

$$h_{nl,r} = -\frac{e_{nl,\phi}}{Z_{TM}} \tag{2.10c}$$

$$h_{nl,\phi} = \frac{e_{nl,r}}{Z_{TM}}, \tag{2.10d}$$

and each mode is normalised according to

$$\int_0^{2\pi} \int_0^a \left(|e_{nl,r}|^2 + |e_{nl,\phi}|^2\right) r dr d\phi = 1. \tag{2.11}$$

It should be noted that this does not give rise to normalisation of the power flow, $W = \int \mathbf{e}_t \times \mathbf{h}_t^* . d\mathbf{S}$, where $d\mathbf{S} = \hat{\mathbf{k}} r dr d\phi$, as the impedance of the modes, Z_{TM}, is itself a function of the radius of each waveguide section, and W is complex for evanescent modes which will likely be present in any case. This therefore determines the values of C_{nl} according to

$$C_{nl} = \left[\int_0^a \left\{ J_n'^2(k_c r) + \left[\frac{n J_n(k_c r)}{k_c r}\right]^2 \right\} (1 + \delta_{n0}) \pi r dr \right]^{-\frac{1}{2}}. \tag{2.12}$$

Using the following recurrence relations, taken from [44],

$$J_n'(x) = \frac{1}{2}\left\{J_{n-1}(x) - J_{n+1}(x)\right\} \tag{2.13a}$$

$$\frac{nJ_n(x)}{x} = \frac{1}{2}\left\{J_{n+1}(x) + J_{n-1}(x)\right\}, \tag{2.13b}$$

and the standard integrals and associated recurrence relationships detailed in [36], equation 2.12 can be reduced to

$$C_{nl} = \left[\left(\frac{1+\delta_{n0}}{2}\right)\pi a^2 J_{n+1}^2(p_{nl})\right]^{-\frac{1}{2}}, \tag{2.14}$$

when $J_n(k_c a) = 0$ is substituted for $k_c a = p_{nl}$, as this implies that $J_{n+1}(x) = -J_{n-1}(x)$. Using equation 2.14, the transverse component of the electric and magnetic fields for the TM modes can be re-expressed as

$$\begin{aligned}\mathbf{e}_{nl} = &\sqrt{\frac{2}{(1+\delta_{n0})\pi a^2 J_{n+1}^2(p_{nl})}} \times \exp\left[j(\omega t - \beta z)\right] \times \\ &\left[J_n'\left(p_{nl}\frac{r}{a}\right)\begin{pmatrix}\cos n\phi \\ \sin n\phi\end{pmatrix}\hat{\mathbf{r}} + \frac{nJ_n\left(p_{nl}\frac{r}{a}\right)}{p_{nl}\frac{r}{a}}\begin{pmatrix}-\sin n\phi \\ \cos n\phi\end{pmatrix}\hat{\phi}\right]\end{aligned} \tag{2.15a}$$

$$\begin{aligned}\mathbf{h}_{nl} = &\sqrt{\frac{2/Z_{TM}}{(1+\delta_{n0})\pi a^2 J_{n+1}^2(p_{nl})}} \times \exp\left[j(\omega t - \beta z)\right] \times \\ &\left[J_n'\left(p_{nl}\frac{r}{a}\right)\begin{pmatrix}\cos n\phi \\ \sin n\phi\end{pmatrix}\hat{\mathbf{r}} + \frac{nJ_n\left(p_{nl}\frac{r}{a}\right)}{p_{nl}\frac{r}{a}}\begin{pmatrix}-\sin n\phi \\ \cos n\phi\end{pmatrix}\hat{\phi}\right].\end{aligned} \tag{2.15b}$$

These are the equations that are required to describe the transverse field components of the TM modes.

To retrieve the corresponding equations for the TE modes, where $E_z = 0$, it is necessary to examine the Helmholtz equation for the $\mathbf{H}$ field and re-arrange it to obtain the Laplacian for H_z, which is shown in equation 2.16.

$$\nabla_t^2 H_z = \frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial H_z}{\partial r}\right) + \frac{1}{r^2}\left(\frac{\partial^2 H_z}{\partial \phi^2}\right) = -k_c^2 H_z. \tag{2.16}$$

The solutions to this equation are of the form

$$H_z(r,\phi) = BJ_n(k_c r)\begin{pmatrix}\sin n\phi \\ \cos n\phi\end{pmatrix}\exp\left[j(\omega t - \beta z)\right], \tag{2.17}$$

where B is a function of z. Considering the impedance of the TE modes, given by $Z_{TE} = Z_0/\sqrt{(p_{nl}'\lambda/2\pi a)^2 - 1}$, the field components can be calculated according to

$$E_r = Z_{TE} H_\phi \tag{2.18a}$$

$$E_\phi = -Z_{TE} H_r. \tag{2.18b}$$

From here, the expressions for the transverse components of the electric and magnetic fields for the TE modes (at the aperture of the horn antenna) are found using the same approach and steps as for the TM mode calculation detailed above. The individual steps can be found in [36]. The components are given by

$$\mathbf{e}_{nl} = D_{nl} \left[\frac{n J_n\left(p'_{nl}\frac{r}{a}\right)}{p'_{nl}\frac{r}{a}} \begin{pmatrix} \cos n\phi \\ -\sin n\phi \end{pmatrix} \hat{\mathbf{r}} - J'_n\left(p'_{nl}\frac{r}{a}\right) \begin{pmatrix} \sin n\phi \\ \cos n\phi \end{pmatrix} \hat{\phi} \right] \times \exp\left[j\left(\omega t - \beta z\right)\right] \tag{2.19a}$$

$$\mathbf{h}_{nl} = D_{nl} Z_{TE} \left[J'_n\left(p'_{nl}\frac{r}{a}\right) \begin{pmatrix} \cos n\phi \\ -\sin n\phi \end{pmatrix} \hat{\mathbf{r}} - \frac{n J_n\left(p'_{nl}\frac{r}{a}\right)}{p'_{nl}\frac{r}{a}} \begin{pmatrix} \sin n\phi \\ \cos n\phi \end{pmatrix} \hat{\phi} \right] \times \exp\left[j\left(\omega t - \beta z\right)\right], \tag{2.19b}$$

where D_{nl} is given by

$$D_{nl} = \left[(1+\delta) \frac{\pi a^2}{2} \left[1 - \left(\frac{n}{p'_{nl}}\right)^2 \right] J_n^2\left(p'_{nl}\right) \right]^{-\frac{1}{2}}. \tag{2.20}$$

The cutoff wavelengths for the TE modes are found by replacing p in equation 2.9 with p', such that

$$\lambda_{TE_{mn}} = \frac{2\pi r}{p'_{nm}}. \tag{2.21}$$

This gives a complete description of the transverse components of the fields for TE and TM modes, and so the cut-off wavenumbers for the modes can now be examined in order to determine the order in which the modes turn on in a waveguide section. The highest wavenumber supported by a waveguide at its central frequency is given by $k_c a$, and so it will only support a finite number of real, propagating modes. Modes with cut-off wavenumbers above the maximum supported by the waveguide will not propagate in the waveguide. Waveguides with larger radii will support more modes at a given wavelength. The number of modes supported also varies with the wavelength of the radiation passing through it, for a fixed waveguide radius. As wavelength decreases, the number of modes supported will increase, so to excite more modes in a certain waveguide, a larger frequency is required. The lowest order fundamental mode which can propagate in a **circular** waveguide, is the TE_{11} mode, where $n = l = 1$. Table 2.1 shows the first ten modes supported in a circular waveguide in the order in which they 'turn on', i.e. sorted in order of increasing cut-off wavenumbers.

Modes which have the same cut-off wavenumbers are said to be degenerate, so long as they have different field distributions.

Hybrid Modes

When the modal content of a corrugated horn is analysed, for example using the singular value decomposition technique which will be described in section 3.2, it is found that the modes

Mode	$k_c a$
TE_{11}	1.8412
TM_{01}	2.4048
TE_{21}	3.0542
TM_{11}	3.8317
TE_{01}	3.8317
TE_{31}	4.2012
TM_{21}	5.1356
TE_{41}	5.3176
TE_{12}	5.3314
TM_{02}	5.5201

Table 2.1: Cut-off wavenumbers for the first ten propagating modes in a circular waveguide

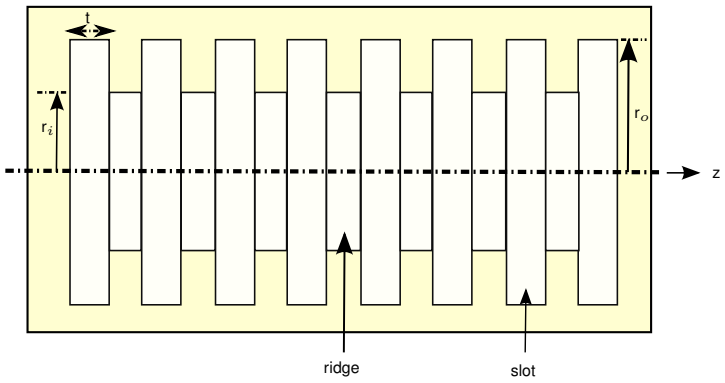

Figure 2.2: Schematic of a corrugated surface waveguide

supported tend to not be TE or TM modes, but rather some linear combination of TE and TM modes. Such modes are known as hybrid modes. An example of a hybrid mode is the HE_{11} mode, which is the mode present in a single-moded corrugated horn. This mode consists of contributions from several TE and TM modes, however it can be represented to a high degree of accuracy by a linear combination of the TE_{11} and TM_{11} modes, in the ratio 85/15.

Figure 2.2 shows a schematic of a waveguide with a corrugated surface. The corrugations are sufficiently dense in terms of wavelength so that the waveguide presents a continuous surface to the waveguide modes. The corrugations act so as to render the surface reactance components of the structure anisotropic such that if the structured is viewed from the azimuthal direction then the average impedance appears to be zero, however if viewed along the axial direction then the uniformity of the impedance is lost due to the presence of the slots (corrugations), with the impedance at any point along the waveguide depending on the depth of the slot.

The spatial dependence of the reactance means that in order to satisfy the boundary conditions required by Maxwell's equations, the magnetic and electric fields cannot be zero in the propagation direction, and so the propagating mode set is not formed of TE and TM modes. The actual mode set is a mixture of these two mode sets, the hybrid modes that were intro-

duced above. A rigorous mathematical treatment of the hybrid modes and their implications for waveguide structures that utilise them can be found in [37] and [45].

The use of TE and TM in modelling waveguide structures will now be examined.

2.2 Matrix Description of Mode Scattering

The impact of changes in the dimensions of a waveguide on a propagating field expressed as a set of TE/TM modes is now considered. A horn of circular cross-section and of arbitrary shape (corrugated, smooth walled, profiled etc.) can be considered as a cascaded series of cylindrical sections of varying radii, and so it can be modelled as a series of discontinuities. This can be modelled efficiently using the scattering matrix approach, [46].

To visualise this approach, consider figure 2.3. The block, **S**, is some arbitrary element which has inputs **A** and **C** and outputs **B** and **D**. Each side of the block has one input and one output. If each input and output is considered to be a set of modes, then each of these inputs/outputs

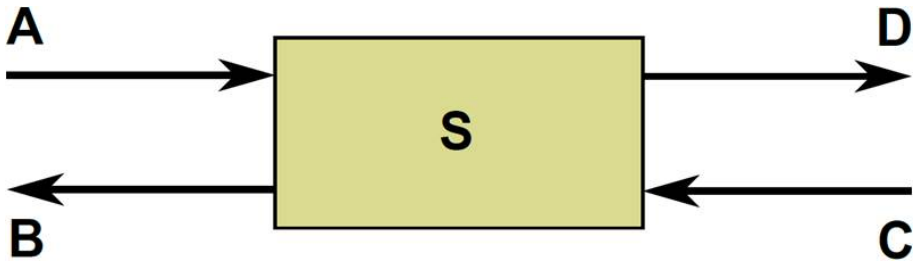

Figure 2.3: Schematic of scattering matrix representation for a block element

can be considered to be a column matrix, whose entries are mode coefficients that represent how much of each mode is in each of the inputs/outputs. These column matrices are represented as $[A]$,$[B]$,$[C]$ and $[D]$, with the labelling convention established in figure 2.3 maintained. $[S]$ now represents the overall scattering matrix that governs the scattering between each of the four input/output matrices, such that

$$\begin{bmatrix} [B] \\ [D] \end{bmatrix} = [S] \begin{bmatrix} [A] \\ [C] \end{bmatrix}. \tag{2.22}$$

The function of the scattering matrix is clear from this equation. It relates the two input fields, **A** and **C** to the two output fields, **B** and **D**. For a given input field, it determines how much of each mode is scattered into each of the two possible output fields, i.e. it determines how much of that input field is reflected (outputs on the same side as the input) and how much of it is transmitted (outputted on the opposite side to the input). The scattering matrix can be considered to consist of four sub-matrices that represent this overall process, as shown in

equation 2.23.

$$[S] = \begin{bmatrix} [S_{11}] & [S_{12}] \\ [S_{21}] & [S_{22}] \end{bmatrix}. \tag{2.23}$$

These matrices are of infinite dimension, as an infinite number of modes are possible, however in calculations the number of modes is specified, and so the matrices have finite dimensions appropriate to this. If n is the number of modes included, then the scattering matrices consists of $4n \times 4n$ elements, using n modes for each of the TE and TM modes. Each of the sub-matrices in the scattering matrix is of dimension $2n \times 2n$. There is assumed to be zero reflection at the waveguide aperture, and so $[C]$ is set to zero, however for individual waveguide sections this is clearly not the case as there will be a field incident on this side of each junction within the horn due to impedance mismatches occurring at the various junctions. The amplitude coefficients for transmission and reflection on the source and output sides of the horn section are calculated according to

$$\begin{bmatrix} [B] \\ [D] \end{bmatrix} = \begin{bmatrix} [S_{11}] & [S_{12}] \\ [S_{21}] & [S_{22}] \end{bmatrix} \begin{bmatrix} [A] \\ [C] \end{bmatrix} = \begin{bmatrix} [S_{11}]\,[A] + [S_{12}]\,[C] \\ [S_{21}]\,[A] + [S_{22}]\,[C] \end{bmatrix}, \tag{2.24}$$

which gives that

$$[B] = [S_{11}]\,[A] + [S_{12}]\,[C]\,, \quad [D] = [S_{21}]\,[A] + [S_{22}]\,[C]\,. \tag{2.25}$$

If a horn is single-moded, then usually only the TE_{11} mode exists in the (fundamentally moded) waveguide that feeds the horn, and so for this section, $[A]$ will have only one non-zero element; the first one. All of the remaining entries $(2n-1)$ in the column will be zero and are present to account for the additional modes included at the aperture. In this case, the matrix appears as

$$[A] = \begin{bmatrix} 1 \\ 0 \\ 0 \\ \cdots \\ \cdots \\ \cdots \end{bmatrix}. \tag{2.26}$$

The non-zero entry corresponds to the TE_{11} mode.

In the case of over-moded horns that are fed by oversized waveguides (capable of supporting multiple modes at the frequency in question), the matrix will look slightly different. In this case, it will contain a number of non-zero elements equal to the number of modes that can propagate in the waveguide. These entries will be located appropriately within the matrix. For such a horn feeding an integrating cavity, and assuming no mode scattering at the cavity end (output) of the horn, each mode is considered as being independent, and the matrix has the

following form,

$$[A] = \begin{bmatrix} 0 \\ 0 \\ a_n \\ \cdots \\ \cdots \\ \cdots \end{bmatrix}. \tag{2.27}$$

Here, a_n is the coefficient for the corresponding mode, with its value chosen in such a way as to ensure that the same normalised power is carried by all of the modes. The modes are chosen so that normalisation with respect to $\int \left(|E_{nl,r}|^2 + |E_{nl,\phi}|^2\right) r dr d\phi = 1$ is obtained, as opposed to with respect to power, so for different modes the a_n will be different. This results in the production of partially coherent fields, and so the overall intensity can be found by adding the resulting scattered modal fields in quadrature at the aperture, following propagation through the horn.

Consider the scattering of modes at a step discontinuity in the waveguide, using the approach outlined in [46]. The waveguides in question here are circularly symmetric, and due to this scattering cannot occur between modes which are of different azimuthal order. This simplifies the form of the scattering matrix, $[S]$, as each azimuthal order can be represented in its own independent matrix. As previously mentioned, the number of modes that can propagate in a section of waveguide depends on the diameter of the section for a fixed wavelength, and so from section to section, as the diameter changes, the number of modes supported changes. If however, these discontinuities are closely spaced, then any evanescent modes transmitted by the first discontinuity will contribute to the incident field at the next discontinuity by a non-negligible amount. They must therefore be accounted for.

Figure 2.4 illustrates a step discontinuity in a waveguide system, where the waveguide radius increases from a to b. This junction is the 'black box' previously considered. A different number

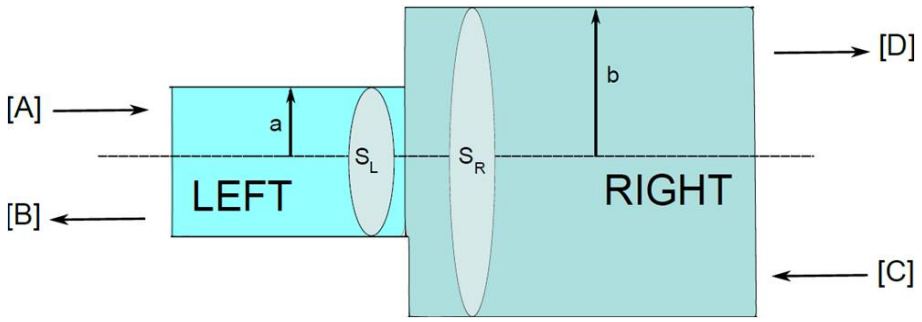

Figure 2.4: Schematic of a waveguide step discontinuity, showing both left and right sides

of modes can be used to represent each side of the junction. This has computational benefits, as unnecessary modes are not considered, however it results in a rectangular scattering matrix

for that junction, and so complicates the field matching step. If equal numbers of modes are used in all sections, the scattering matrices are square, and the overall process is simplified. In the discontinuity (junction) plane, the electric and magnetic fields are given by

$$\mathbf{E}_L = \sum_{n=1}^{N} \left\{ A_n e^{-j\beta_n z} + B_n e^{j\beta_n z} \right\} \mathbf{e}_{nL} \tag{2.28a}$$

$$\mathbf{H}_L = \sum_{n=1}^{N} \left\{ A_n e^{-j\beta_n z} - B_n e^{j\beta_n z} \right\} \mathbf{h}_{nL} \tag{2.28b}$$

$$\mathbf{E}_R = \sum_{n=1}^{N} \left\{ D_n e^{-j\beta_n z} + C_n e^{j\beta_n z} \right\} \mathbf{e}_{nR} \tag{2.28c}$$

$$\mathbf{H}_R = \sum_{n=1}^{N} \left\{ D_n e^{-j\beta_n z} - C_n e^{j\beta_n z} \right\} \mathbf{h}_{nR}, \tag{2.28d}$$

where $\pm\beta_n$ is the propagation constant encountered previously, whose sign is determined by the direction of travel. The terms $\mathbf{e}_n e^{\pm j\beta z}$ and $\mathbf{h}_n e^{\pm j\beta z}$ represent the transverse fields for the n^{th} mode, and the subscripts L and R represent the left and right sides of the junction, as indicated in figure 2.4. A_n and B_n are respectively the forward and reflected amplitude coefficients for the n^{th} mode on the left side of the junction, with D_n and C_n fulfilling the equivalent purpose for the right side of the junction. It should be noted that the junction is located at $z = 0$. We now consider three relationships conditions;

1. At the interface of the junction, the transverse electric and magnetic fields must be continuous.

2. In the larger waveguide, the transverse electric field must be zero outside the common free space region between the two sections, due to the presence of the conducting waveguide wall.

3. The modes have an orthogonality relationship, namely $\int_S \mathbf{e}_n \times \mathbf{h}_m^* . d\mathbf{S} = 0$ for $n \neq m$.

These conditions allow the calculation of the mode coefficients, according to

$$\sum_{n=1}^{N} (A_n + B_n) \mathbf{e}_{nL} = \sum_{n=1}^{N} (D_n + C_n) \mathbf{e}_{nR} \tag{2.29a}$$

$$\sum_{n=1}^{N} (A_n - B_n) \mathbf{h}_{nL} = \sum_{n=1}^{N} (D_n - C_n) \mathbf{h}_{nR}. \tag{2.29b}$$

The transverse fields are continuous across the junction, so combining this with the orthogonality requirement (with $d\mathbf{S}$ pointing in the z direction), results in the following set of simultaneous

equations being formed,

$$\sum_{n=1}^{N}(A_n+B_n)\int_{S_L}\mathbf{e}_{nL}\times\mathbf{h}^*_{n'R}.d\mathbf{S}=\sum_{n=1}^{N}(D_n+C_n)\int_{S_L}\mathbf{e}_{nR}\times\mathbf{h}^*_{n'R}.d\mathbf{S}. \tag{2.30}$$

Since at the conducting wall (to satisfy the boundary condition that the transverse electric field is zero there), over S_R-S_L, $\sum_{n=1}^{N}(D_n+C_n)\,\mathbf{e}_{nR}=0$, the integral over S_L can be transformed into an integral over S_R. As a result of this, we have that

$$\sum_{n=1}^{N}(A_n+B_n)\int_{S_L}\mathbf{e}_{nL}\times\mathbf{h}^*_{n'R}.d\mathbf{S}=\sum_{n=1}^{N}(D_n+C_n)\int_{S_L}\mathbf{e}_{nR}\times\mathbf{h}^*_{n'R}.d\mathbf{S}. \tag{2.31}$$

This allows equation 2.30 to be re-expressed as

$$\sum_{n}(A_n+B_n)\,P_{n'n}=\sum_{n}(D_n+C_n)\,Q_{n'n} \tag{2.32a}$$

$$[P]\,[[A]+[B]]=[Q]\,[[D]+[C]]\,, \tag{2.32b}$$

where equation 2.32b is equation 2.32a expressed in matrix form, and the power coupling terms are given by

$$P_{n'n}=\int_{S_L}\mathbf{e}_{nL}\times\mathbf{h}^*_{n'R}.d\mathbf{S}, \tag{2.33a}$$

$$Q_{n'n}=\left(\int_{S_R}\mathbf{e}_{nR}\times\mathbf{h}^*_{n'R}.d\mathbf{S}\right)\delta_{n'n}. \tag{2.33b}$$

$[P]$ is a square matrix that contains the mutual power that is coupled between the modes on the left side of the junction to those on the right. For modes on the right of the junction, the self coupled power between themselves is given by $[Q]$, which is a diagonal matrix. For the continuity of the magnetic fields, a similar calculation results in the following corresponding expression,

$$[R]^*\,[[A]-[B]]=[P]^\dagger\,[[D]-[C]]\,, \tag{2.34}$$

where $R_{nn'}$ is the self coupled power on the right of the junction, given by

$$R_{nn'}=\left(\int_{S_L}\mathbf{e}_{n'L}\times\mathbf{h}^*_{nL}.d\mathbf{S}\right)\delta_{n'n}. \tag{2.35}$$

The power coupling integrals, $[P]$, $[Q]$ and $[R]$, are evaluated by involving the surface integral

relation,

$$\int_S \mathbf{e} \times \mathbf{h}^* . d\mathbf{S} = \int_0^{2\pi} \int_0^a \left(E_r H_\phi^* - E_\phi H_r^* \right) r dr d\phi, \tag{2.36}$$

where $\mathbf{e} \times \mathbf{h}^*$ is recognised as the Poynting vector. Integrating this over a surface gives the power flowing through that surface, and in its general form it can account for resistive losses due to any walls present in the structure under consideration, for example waveguide walls, [43]. The Poynting vector is real for propagating (real) modes, so when a lossless medium is used, as is assumed in SCATTER for the waveguide walls, it is found that there is no energy dissipation, and the average power that enters one side of the junction equals the power leaving the other side, [36].

For evanescent (non-propagating) modes, the Poynting vector is *purely* imaginary. In the case of TE modes, it is found that for a given length of waveguide, the amount of time averaged magnetic energy stored in that length is greater than the equivalent electric energy stored in the same length. The opposite is true for evanescent TM modes. The net energies can be viewed as being reactive, occurring due to the inductive and capacitive nature of the impedances of the evanescent modes themselves. Therefore, for evanescent modes, $\int_S \mathbf{E} \times \mathbf{H} . d\mathbf{S}$ physically represents energy flowing backwards and forwards within the length of waveguide in question, to supply the instantaneous changes in the net stored energy in the volume of waveguide.

To evaluate the power scattered between the modes, the modal power integral, P, must be evaluated. This is a computationally intensive calculation, however by taking an analytical approach, closed form expressions for the integrals can be found that significantly reduces the computational effort required. Expressions must therefore be derived, using these closed form of the integrals, that allow the calculation of all possible modal scattering transitions (self and mutual power coupling), i.e. TE to TE, TE to TM, TM to TE and TM to TM.

For a waveguide of radius a feeding into a waveguide of radius b, with $b > a$, the mutual power coupling of the modes is given by

$$P_{nl',l} = \int_0^{2\pi} \int_0^r \left(E_r H_\phi^* - E_\phi H_r^* \right) r dr d\phi = \int_0^{2\pi} \int_0^a \left(e^a_{nl',r} \frac{e^b_{nl',r}}{Z^b_{nl'}} + e^a_{nl',\phi} \frac{e^b_{nl',\phi}}{Z^b_{nl'}} \right) r dr d\phi. \tag{2.37}$$

Using the transverse field equations that govern TE to TE and TM to TM mode scattering, this integral can be written in the general form

$$P_{nl',l} = (1 + \delta_{n0}) \pi \frac{A_L A_R^*}{Z^b} \int_0^a \left[\frac{n J_n (k_L r)}{(k_L r)} \frac{n J_n (k_R r)}{(k_R r)} + J_n' (k_L r) J_n' (k_R r) \right] r dr, \tag{2.38}$$

where A is the modal normalisation factor, that in the case of TM modes refers to C_n, and

in the case of TE modes refers to D_n, both of which were previously derived. $k_L = k_c$ for a waveguide of radius a, and $k_R = k_c$ for a waveguide of radius b. Rewriting the two integral components of equation 2.38 using the previously encountered recurrence relationships, allows the standard Bessel integral and expansion forms to be used, yielding equation 2.39, which represents the mutual power coupling between modes, [36, 37].

$$P_{nl',l} = \frac{(1+\delta_{n0})\,\pi A_L A_R^* a}{Z^b\left[k_R^2 - k_L^2\right]}\left[k_R J_n'\left(k_L a\right) J_n\left(k_R a\right) - k_L J_n'\left(k_R a\right) J_n\left(k_L a\right)\right]. \tag{2.39}$$

If we use that for TM modes, $J_n\left(k_L a\right) = 0$, and for TE modes that $J_n'\left(k_L a\right) = 0$, and if A_L and A_R^* are replaced with the appropriate amplitude coefficients, then equation 2.39 simplifies to

$$P_{TM-TM} = \frac{(1+\delta_{n0})\,\pi C_{nl} C_{nl'} a}{Z_{TM}^{b*}\left[\left(\frac{p_{nl'}}{b}\right)^2 - \left(\frac{p_{nl}}{a}\right)^2\right]}\left(\frac{p_{nl'}}{b}\right) J_n'\left(p_{nl}\right) J_n\left(\frac{p_{nl'} a}{b}\right), \tag{2.40a}$$

$$P_{TE-TE} = \frac{-(1+\delta_{n0})\,\pi D_{nl} D_{nl'}}{Z_{TE}^{b*}\left[\left(\frac{p'_{nl'}}{b}\right)^2 - \left(\frac{p'_{nl}}{a}\right)^2\right]} p'_{nl} J_n\left(p'_{nl}\right) J_n'\left(\frac{p'_{nl'} a}{b}\right). \tag{2.40b}$$

To find the coupling between the remaining combinations of modes, a similar procedure is used, [37], resulting in the equations

$$P_{TE-TM} = \frac{\pi D_{nl} C_{nl'}}{2 Z_{TM}^{b*}} J_n\left(p_{nl}\right) J_n\left(\frac{p_{nl'} a}{b}\right), \tag{2.41a}$$

$$P_{TM-TE} = 0. \tag{2.41b}$$

Equation 2.41b shows that the TM modes to the left of the junction do not scatter any of their power into the TE modes on the right of the junction. The elements of the overall power scattering matrix, $[P]$, are therefore arranged as follows, in ascending order with respect to n and l,

$$P = \begin{bmatrix} P_{TE-TE} & P_{TE-TM} \\ P_{TM-TE} & P_{TM-TM} \end{bmatrix} \tag{2.42}$$

$[Q]$ and $[R]$ are diagonal matrices that correspond to the self coupling on the right and left hand sides of the junction respectively. Their elements are calculated using the transverse field components, as was done in the mutual coupling case. The fact that an orthogonal set of modes is being used is confirmed by the fact that these matrices are diagonal. For TM to TM coupling on the right hand side of the junction, the power coupling integral is written as

$$R_{TM-TM} = \frac{C_{nl}^2}{Z_{TM}^{a*}} \int_0^a \left[J_{n-1}^2\left(k_L r\right) + J_{n+1}^2\left(k_L r\right)\right] r dr. \tag{2.43}$$

Using the recurrence relations previously encountered, this can be re-expressed as

$$R_{TM-TM} = \frac{C_{nl}^2}{Z_{TM}^{a^*}} a^2 J_{n+1}^2 \left(p_{nl}\right) = \frac{1}{Z_{TM}^{a^*}}. \tag{2.44}$$

Similarly, Q_{TM-TM} is found to be given by

$$Q_{TM-TM} = \frac{1}{Z_{TM}^{b^*}}. \tag{2.45}$$

Using the fact that $J_n\left(p'_{nl}\right) = 0$, it is also found that

$$R_{TE-TE} = \frac{1}{Z_{TE}^{a^*}} \tag{2.46}$$

$$Q_{TE-TE} = \frac{1}{Z_{TM}^{b^*}}. \tag{2.47}$$

The form of the power integrals is now known, and so combining this with equations 2.32b and 2.34 allows the scattering matrix components to be calculated. $[A]$ and $[C]$ are arbitrary, so if $[C]$ is set equal to zero, then using this and setting $[D] = [S_{21}]\,[A]$ and $[B] = [S_{11}]\,[A]$ in equations 2.32b and 2.34, we retrieve the following simultaneous equations

$$[P]\,[S_{11}] - [Q]\,[S_{21}] = -\,[P] \tag{2.48a}$$

$$[R]^*\,[S_{11}] - [P]^\dagger\,[S_{21}] = -\,[R]^*\,. \tag{2.48b}$$

Solving these simultaneous equations for the scattering matrix components yields

$$[S_{21}] = 2\left[[Q] + [P]\,[R^*]^{-1}\,[P]^\dagger\right]^{-1}[P] \tag{2.49a}$$

$$[S_{11}] = \left[[R^*] + [P]^\dagger\,[Q]^{-1}\,[P]\right]^{-1}\left[[R^*] - [P]^\dagger\,[Q]^{-1}\,[P]\right] \tag{2.49b}$$

As an alternative, $[A]$ can be set equal to zero. This then gives that

$$[B] = [S_{12}]\,[C] \tag{2.50a}$$

$$[D] = [S_{22}]\,[C]\,, \tag{2.50b}$$

which, when solved simultaneously, yields the following scattering matrix elements

$$[S_{12}] = 2\left[[R^*] + [P]^\dagger\,[Q]^{-1}\,[P]\right]^{-1}[P]^\dagger \tag{2.51a}$$

$$[S_{22}] = -\left[[Q] + [P]\,[R^*]^{-1}\,[P]^\dagger\right]^{-1}\left[[Q] - [P]\,[R^*]^{-1}\,[P]^\dagger\right]. \tag{2.51b}$$

The overall scattering matrix, $[S]$, therefore consists of sub-matrices whose elements are given

by equations 2.49a, 2.49b, 2.51a and 2.51b.

The above analysis assumes that the radiation is passing into a waveguide with a larger radius than the waveguide that the radiation is passing into it from, and that the modal scattering occurs at the interface between these two waveguides. If the modes are travelling from a larger waveguide into a smaller waveguide, as will happen when corrugations are under consideration, then a re-calculation of the scattering matrix is required. In this case, we have that $a > b$. This case is modelled by considering the situation to be totally equivalent to the situation where $a < b$, but the direction of propagation of the modes has been reversed. We therefore have that

$$[A'] = [C] \tag{2.52a}$$

$$[B'] = [D] \tag{2.52b}$$

$$[C'] = [A] \tag{2.52c}$$

$$[D'] = [B]\,, \tag{2.52d}$$

where the prime indicates the case where $a > b$, and the unprimed case indicates that $a < b$. Using these equations gives

$$\begin{bmatrix} [B'] \\ [D'] \end{bmatrix} = \begin{bmatrix} [D] \\ [B] \end{bmatrix} = \begin{bmatrix} [S_{22}] & [S_{21}] \\ [S_{12}] & [S_{11}] \end{bmatrix} \begin{bmatrix} [C] \\ [A] \end{bmatrix} = \begin{bmatrix} [S_{22}] & [S_{21}] \\ [S_{12}] & [S_{11}] \end{bmatrix} \begin{bmatrix} [A'] \\ [C] \end{bmatrix}. \tag{2.53}$$

Thus,

$$\begin{bmatrix} [S'_{11}] & [S'_{12}] \\ [S'_{21}] & [S'_{22}] \end{bmatrix} = \begin{bmatrix} [S_{22}] & [S_{21}] \\ [S_{12}] & [S_{11}] \end{bmatrix}. \tag{2.54}$$

This clearly shows that once the scattering matrix elements are calculated for a junction, the overall scattering matrix for either direction of propagation can be readily found, and so steps from wider to narrower waveguides, and vice versa, can be handled readily.

Finally, the effect on the modes of propagating them along a short uniform waveguide section is considered. This, combined with the treatment above, allows the modelling of a complete horn. There is clearly no reflection at either end of modes externally incident on the section, and so the reflection coefficients are zero. This gives the scattering matrix for the section a very simple form that includes the modal phase change and the modes propagate through the section. We have that

$$[S_{11}] = [S_{22}] = 0 \tag{2.55a}$$

$$[S_{12}] = [S_{21}] = [V]\,, \tag{2.55b}$$

where $[V]$ is a diagonal matrix of dimension equal to the number of modes being used in the

analysis. [46] gives the elements of the scattering matrix as

$$V_{nm} = e^{-j\beta_n L}\delta_{nm}, \tag{2.56}$$

where, as before, β_n is the propagation constant for a particular mode and is real for real propagating modes, and pure imaginary for evanescent modes. L is the length of the waveguide section. The scattering matrix for this section must be cascaded with a scattering matrix for a discontinuity in order to obtain the overall scattering matrix for a complete waveguide section. This is because the matrix for the uniform section adds the modal phase slippage that occurs over the finite length of the waveguide section. Simply cascading the step matrices would not suffice, as this would not be included. This approach could only work if the sections were vanishingly thin, which is unphysical. It is also important to note that the sections in question are generally short, and in this case evanescent modes that are excited due to scattering into higher order modes at the previous junction *will* contribute to the aperture field at the following junction. For this reason, a sufficiently high number of modes must be included in the analysis so as to model this accurately.

The final step in the process is to calculate the overall scattering matrix for the entire horn. This is done by cascading the scattering matrices for each step, and so the effects of the modal scattering at all of the junctions is accounted for. If we consider two consecutive sections of waveguide, with scattering matrices $[S^a]$ and $\left[S^b\right]$ such that

$$[S^a] = \begin{bmatrix} [S^a_{11}] & [S^a_{12}] \\ [S^a_{21}] & [S^a_{22}] \end{bmatrix} \qquad \left[S^b\right] = \begin{bmatrix} \left[S^b_{11}\right] & \left[S^b_{12}\right] \\ \left[S^b_{21}\right] & \left[S^b_{22}\right] \end{bmatrix}. \tag{2.57}$$

These two matrices are cascaded into an overall matrix for that junction, called $[S^c]$, which is found according to, [46]

$$[S^c_{11}] = [S^a_{12}]\left[[I] - \left[S^b_{11}\right][S^a_{22}]\right]^{-1}\left[S^b_{11}\right][S^a_{21}] + [S^a_{11}] \tag{2.58a}$$

$$[S^c_{12}] = [S^a_{12}]\left[[I] - \left[S^b_{11}\right][S^a_{22}]\right]^{-1}\left[S^b_{12}\right] \tag{2.58b}$$

$$[S^c_{21}] = \left[S^b_{21}\right]\left[[I] - [S^a_{22}]\left[S^b_{11}\right]\right]^{-1}[S^a_{21}] \tag{2.58c}$$

$$[S^c_{22}] = \left[S^b_{21}\right]\left[[I] - [S^a_{22}]\left[S^b_{11}\right]\right]^{-1}[S^a_{22}]\left[S^b_{12}\right] + \left[S^b_{22}\right], \tag{2.58d}$$

where $[I]$ is the identity matrix, and $[]^{-1}$ is the inverse. By using an iterative process, the complete scattering matrix for one whole section is cascaded with the complete scattering matrix for all previous sections, resulting in an overall scattering matrix that describes the entire horn. Using this, the field patterns associated with the horn can be calculated. This will now be considered.

2.3 Waveguide/Horn Antenna Radiation Patterns

The scattering matrices calculated via the mode matching approach can be used to compute the radiation patterns of the waveguide/horn under consideration, where a horn antenna is simply a flared waveguide section. It is usually assumed that for electromagnetic horns, [37,47,48], the horn aperture acts as the radiation source. Therefore, the aperture field must be calculated, and for this, the field equivalence principle is employed, which is a more rigorous formulation of the Huygens principle. This requires that the antenna's aperture be replaced with a sheet of equivalent currents that produce radiation fields that are similar to those produced by the aperture.

Consider an aperture plane with the real sources of radiation located behind it, and an observation point, P, located some distance in front of the aperture. Outside of the aperture extent, the fields are assumed to be zero on the aperture plane. The tangential components of the field are known in the aperture, and are assumed to be as a result of equivalent current sources within the aperture, despite being as a result of the real sources behind the aperture that have propagated through the waveguide structure. This definition of the aperture field strengths allows the corresponding electric and magnetic source densities, $\mathbf{J}$ and $\mathbf{M}$, to be written as

$$\mathbf{J}_a = \hat{\mathbf{n}} \times \mathbf{H}_a \tag{2.59a}$$

$$\mathbf{M}_a = -\hat{\mathbf{n}} \times \mathbf{E}_a, \tag{2.59b}$$

where $\mathbf{E}_a$ and $\mathbf{H}_a$ represent the electric and magnetic tangential field components in the aperture, and $\hat{\mathbf{n}}$ is the normal to the aperture plane on which the currents are defined. For a homogeneous isotropic system, the associated retarded potentials for the electric ($\mathbf{A_a}$) and magnetic ($\mathbf{F_a}$) current sources are given respectively by [49] as

$$\mathbf{A}_a = \frac{\mu}{4\pi} \int_{S'} \frac{\mathbf{J}_a e^{jkr}}{r} ds' \tag{2.60a}$$

$$\mathbf{F}_a = \frac{\epsilon}{4\pi} \int_{S'} \frac{\mathbf{M}_a e^{jkr}}{r} ds', \tag{2.60b}$$

where s' is the aperture, and the integral is over the aperture area. It follows from this, that if the aperture fields are known, then the corresponding fields in the farfield of the aperture can be found using these potentials. In the farfield, the electric and magnetic components of the field are orthogonal, and form a TEM (combined TE/TM) wave. The electric field is given by $\mathbf{E}_{FF} = \eta_0 \mathbf{H}_{FF} \times \hat{\mathbf{u}} = -j\omega\eta \mathbf{F} \times \hat{\mathbf{u}}$, where η is the wave impedance, given by $\eta = \sqrt{\mu/\epsilon}$ and $\eta_0 = \sqrt{\mu_0/\epsilon_0}$. In the farfield, the total electric field is given by [49] as $\mathbf{E}_{FF} = -j\omega\left(\eta_0 \mathbf{F} + \mathbf{A} \times \hat{\mathbf{u}}\right)$. The field distribution in the farfield is given by the Fourier Transform of the corresponding aperture field, and vice versa. The aperture plane is two-dimensional, and the two-dimensional

Fourier Transform of this field (to give the farfield) is given by

$$\mathbf{E_{ff}}(\theta,\phi) = \iint_{S'} \mathbf{E}_a(x',y')\, e^{-jkr} dS', \tag{2.61}$$

where $\mathbf{r} = |\mathbf{r}_0 - \mathbf{r}'| \approx r_0 - \hat{\mathbf{r}}_0.\mathbf{r}' = r_0 - r'\cos\psi$, and r_0 and $\mathbf{r}$ are respectively the distance from the aperture centre and some point in the aperture to the observation point, P. Spherical coordinates are now introduced and substituted for r in the phase term for the farfield observation plane, according to the transformation

$$r = r_0 - r'\sin\theta\cos(\phi - \phi'). \tag{2.62}$$

Using this, the Fourier Transform of the aperture field can be split into its Cartesian components,

$$f_x(\theta,\phi) = \int_{x'}\int_{y'} \mathbf{E}_{ax}(x',y') \exp\left[jk\left\{r'\sin\theta\cos(\phi-\phi')\right\}\right] r' dr' d\phi' \tag{2.63a}$$

$$f_y(\theta,\phi) = \int_{x'}\int_{y'} \mathbf{E}_{ay}(x',y') \exp\left[jk\left\{r'\sin\theta\cos(\phi-\phi')\right\}\right] r' dr' d\phi'. \tag{2.63b}$$

If it is assumed that the farfield field distributions correspond to those of a TEM field in the aperture, then we have that $E_{ax} = \eta_0 H_{ay}$ and $E_{ay} = \eta_0 H_{ax}$. Inside the horn, both axial components of H_z and E_z exist, so this is only an approximation, however at the aperture the field strengths of these components tend to zero. If the electric fields are written in terms of their spherical polar components, the radiated electric fields take on the form

$$E_\theta \propto (1+\cos\theta)(f_x\cos\phi + f_y\sin\phi) \tag{2.64a}$$

$$E_\phi \propto (1+\cos\theta)(-f_x\sin\phi + f_y\cos\phi). \tag{2.64b}$$

This method is known as the modified **E**-field and **H**-field model. An alternative model that is commonly used is known simply as the **E**-field model. In this case. an imaginary, conducting surface is placed behind the current sources, such that it does not alter the current densities themselves. Using image theory, [36], it is predicted that image currents are produced such that $\mathbf{J} = 0$ and $\mathbf{M} = \hat{\mathbf{n}} \times 2\mathbf{E}_a$. In the farfield, this model yields the following radiation distributions

$$E_\theta \propto (f_x\cos\phi + f_y\sin\phi) \tag{2.65a}$$

$$E_\phi \propto \cos\theta(-f_x\sin\phi + f_y\cos\phi). \tag{2.65b}$$

When the aperture field is approximately linearly polarised, and the aperture is of a large radius, the models agree closely over the extent of the main beam. The differences between the models only become clear for horns with narrow apertures and for large off-axis angles, [46].

When presenting farfield radiation patterns, it is normal to do so in terms of the principal planes of polarisation of interest. If a reference plane of polarisation is selected, then the co-polar field is the field that is parallel to this plane of polarisation. The zero point of the azimuthal angle, ϕ, is usually placed so as it corresponds to the **E**-plane, as this is the principal co-polar direction. The **H**-plane is obtained by setting $\phi = 90°$. The orthogonal component of this field is the cross-polar component. All of these definitions are dependent on the direction with respect to which one measures the polarisation, and to this end, Ludwig's third definition is used, [50]. This states that the reference polarisation is taken to be the polarisation of a Huygen's source (electric and magnetic tangential field components that lie in the aperture plane, radiating equal, equiphase, fields along the z-axis). The cross-polarisation direction is obtained by rotating the source through 90° in the aperture plane. This definition gives the co-polar (x-direction) and cross-polar (y-direction) components respectively, as

$$E_{co} = E_\theta \cos\theta \cos\phi - E_\phi \sin\phi \approx E_\theta \cos\phi - E_\phi \sin\phi = (1 + \cos\theta)\, f_x \tag{2.66a}$$

$$E_{xp} = E_\theta \cos\theta \sin\phi + E_\phi \cos\phi \approx E_\theta \sin\phi + E_\phi \cos\phi = (1 + \cos\theta)\, f_y, \tag{2.66b}$$

The overall radiated field (including orthogonal components), expressed in spherical components, is given by equation 2.67, where the plane that is being analysed is specified by setting ϕ to an appropriate value.

$$\mathbf{E}\,(r, \theta, \phi) \cong [(1 + \cos\theta)\,(f_x \cos\phi + f_y \sin\phi)]\,\hat{\theta} + [(1 + \cos\theta)\,(f_y \cos\phi - f_x \sin\phi)]\,\hat{\phi}, \tag{2.67}$$

where f_x and f_y are given in equations 2.63a and 2.63b. For these expressions, the trigonometric functions that multiply the integrals are simply obliquity factors. The $\cos\theta$ term attenuates the diffracted waves as a function of θ, and in doing so accounts for the amplitude dependence on the direction of boresight. The standalone trigonometric functions account for the projection effect onto the farfield, of aperture equivalent surface current densities. If the aperture is several wavelengths wide, which it typically is, then the obliquity factor does bot significantly attenuate the main beam and the first few sidelobes, however it is still included within the modelling technique to guarantee accuracy.

2.3.1 Farfield Radiation Distribution of Cylindrical Waveguide Modes

The previously obtained forms for the aperture field of conical horns in terms of TE and TM modes must be transformed from cylindrical to Cartesian form in order to calculate the f_x and f_y components that are required for the calculation of the farfield.

Using the standard recurrence relations that allow the rewriting of the Bessel functions, and by expressing the cylindrical unit vectors $\hat{\mathbf{r}}$ and $\hat{\phi}$ in terms of the Cartesian unit vectors $\hat{\mathbf{i}}$

and $\hat{\mathbf{j}}$, allows the *TM* field to be written as

$$\begin{aligned}\mathbf{e}_{nl}^{TM} \propto \left[J_{n-1}\left(p_{nl}\frac{r}{a}\right) - J_{n+1}\left(p_{nl}\frac{r}{a}\right)\right]\begin{pmatrix}\cos n\phi \\ \sin n\phi\end{pmatrix} \\ \left(\cos\phi\hat{\mathbf{i}} + \sin\phi\hat{\mathbf{j}}\right) - \left[J_{n-1}\left(p_{nl}\frac{r}{a}\right) + J_{n+1}\left(p_{nl}\frac{r}{a}\right)\right]\begin{pmatrix}-\sin n\phi \\ \cos n\phi\end{pmatrix} \\ \left(-\sin\phi\hat{\mathbf{i}} + \cos\phi\hat{\mathbf{j}}\right).\end{aligned} \tag{2.68}$$

Combining the trigonometric functions and isolating the Cartesian components, this expression can be re-written as

$$\begin{aligned}\mathbf{e}_{nl}^{TM} \propto &\left[J_{n-1}\left(p_{nl}\frac{r}{a}\right)\cos\left(n-1\right)\phi - J_{n+1}\left(p_{nl}\frac{r}{a}\right)\cos\left(n+1\right)\phi\right]\hat{\mathbf{i}} \\ &-\left[J_{n-1}\left(p_{nl}\frac{r}{a}\right)\sin\left(n-1\right)\phi + J_{n+1}\left(p_{nl}\frac{r}{a}\right)\sin\left(n+1\right)\phi\right]\hat{\mathbf{j}}.\end{aligned} \tag{2.69}$$

Similarly, the TE modes can be written as

$$\mathbf{e}_{nl}^{TE} \propto \left[J_{n-1}\left(p'_{nl}\frac{r}{a}\right)\cos\left(n-1\right)\phi + J_{n+1}\left(p'_{nl}\frac{r}{a}\right)\cos\left(n+1\right)\phi\right]\hat{\mathbf{i}} - \left[J_{n-1}\left(p'_{nl}\frac{r}{a}\right)\sin\left(n-1\right)\phi - J_{n+1}\left(p'_{nl}\frac{r}{a}\right)\right. \tag{2.70}$$

The f_x term can be calculated by considering the $\hat{\mathbf{i}}$ component of the TM modes, according to

$$\begin{aligned}f_x \propto \int_0^a\int_0^{2\pi}&\left\{J_{n-1}\left(p_{nl}\frac{r'}{a}\right)\cos\left(n-1\right)\phi' - J_{n+1}\left(p_{nl}\frac{r'}{a}\right)\cos\left(n+1\right)\phi'\right\} \\ &e^{jkr'\sin\theta[\cos(\phi-\phi')]}r'dr'd\phi',\end{aligned} \tag{2.71}$$

where the prime refers to the horn aperture. The Fourier Transform of this expression can be performed analytically. Performing the inner integral with respect to ϕ yields, [36]

$$\begin{aligned}f_x \propto \int_0^a &\left\{J_{n-1}\left(p_{nl}\frac{r'}{a}\right)J_{n-1}\left(kr'\sin\theta\right)\cos\left(n-1\right)\phi\right\} \\ &-\left\{J_{n+1}\left(p_{nl}\frac{r'}{a}\right)J_{n+1}\left(kr'\sin\theta\right)\cos\left(n+1\right)\phi\right\}r'dr'.\end{aligned} \tag{2.72}$$

This integral can be evaluated using the standard integral forms of the Bessel functions, resulting in the x component of the TM modes being given by

$$E_{xFF}^{TM} \propto \left(1+\cos\theta\right)\sqrt{\frac{2}{\pi J_{n+1}^2\left(p_{nl}\right)}}\left\{\frac{1}{2}\left(\jmath_{n-1}^{FF}\left(p_{nl}\right)\cos\left(n-1\right)\phi + \jmath_{n+1}^{FF}\left(p_{nl}\right)\cos\left(n+1\right)\phi\right)\right\}, \tag{2.73}$$

where $\jmath$ is the result of the integration performed on equation 2.72, and is given by

$$\jmath^{FF}_{n\pm1}(\theta) = \frac{r\left\{p_{nl}J_{n\pm1}(kr\sin\theta)\,J'_{n\pm1}(p_{nl}) - (kr\sin\theta)\,J_{n\pm1}(p_{nl})\,J'_{n\pm1}(kr\sin\theta)\right\}}{(kr\sin\theta)^2 - p^2_{nl}}. \tag{2.74}$$

Similarly, the remaining components of the electric field are found to be

$$E^{TM}_{yFF} \propto -(1+\cos\theta)\sqrt{\frac{2}{\pi J^2_{n+1}(p_{nl})}} \left\{\frac{1}{2}\left(\jmath^{FF}_{n-1}(p_{nl})\cos(n-1)\,\phi - \jmath^{FF}_{n+1}(p_{nl})\cos(n+1)\,\phi\right)\right\} \tag{2.75}$$

$$E^{TE}_{xFF} \propto (1+\cos\theta)\sqrt{\frac{2}{\pi J^2_{n+1}(p_{nl})}} \left\{\frac{1}{2}\left(\jmath^{FF}_{n-1}(p_{nl})\cos(n-1)\,\phi - \jmath^{FF}_{n+1}(p_{nl})\cos(n+1)\,\phi\right)\right\} \tag{2.76}$$

$$E^{TE}_{yFF} \propto -(1+\cos\theta)\sqrt{\frac{2}{\pi\left(1-\left(\frac{n}{p'_{nl}}\right)^2\right)J^2_n(p'_{nl})}}\times \left\{\frac{1}{2}\left(\jmath^{FF}_{n-1}(p_{nl})\cos(n-1)\,\phi - \jmath^{FF}_{n+1}(p_{nl})\cos(n+1)\,\phi\right)\right\}. \tag{2.77}$$

The modes orthogonal to those listed above can be obtained by moving ϕ through $\frac{\pi}{2}$. These orthogonal modes, combined with those listed above, can be used to precisely recreate the fields that are calculated in Scatter via the mode matching technique.

The implementation of the theory discussed to this point in a single piece of software in order to model the performance and radiation patterns of arbitrary waveguide structures of circular cross-section will now be discussed.

2.4 SCATTER Implementation

SCATTER was originally developed in Mathematica, a technical, symbolic environment that is based on the programming language C, [36, 37]. Versions also now exist in Matlab and in raw C. Syntax differences aside, each implementation is equivalent, and so a brief description of the general implementation will now be given. Scatter provides an excellent platform for the calculation of the properties of waveguide based structures and their radiation patterns, along with any associated performance metrics. SCATTER currently exists for both rectangular and circular geometries and each version is in constant development, being made more efficient and having new functionality added.

SCATTER consists of five sections, which can be summarised in the following sequential

steps:

- Reading in of a data file containing the data that describes the horn shape in terms of the length and radius of each section that the horn has been broked down into (geometry file)
- Scattering matrices calculation
- Phase and intensity calculation for the horn aperture
- Phase and intensity calculation in the farfield of the horn
- Output of all data to a .txt file for additional analysis/records

These steps will now be briefly discussed.

2.4.1 Geometry File Input

The geometry file is typically of . txt or . i format, and consists of two columns, as illustrated in table 2.2. It is assumed that a given section has length ΔL and radius a. The second column,

Line Number	Geometry File Column 1	Geometry File Column 2
Line 1	Calculation Frequency (ν) /GHz	
Line 2	Maximum Azimuthal Order, n	
Line 3	Number of Horn Sections, N	
Line 4$\rightarrow$ N + 3	Length of i^{th} section, ΔL /mm	
Line N + 4 $\rightarrow$ 2N + 3	Radius of i^{th} section, a /mm	Number of Modes for Junction

Table 2.2: Layout of SCATTER geometry file

which is empty outside of the radius portion of the geometry file, contains the number of modes required to describe the scattering at the junction. This is present for historical reasons, and is not used by SCATTER. SCATTER assumes (currently) that the number of modes is constant for each section, and is defined at the start of the program. N is determined by the requirement to have at least ten sections per wavelength in a smooth-walled horn. In the case of a corrugated horn, the number of corrugations determines the number of sections, however the number of corrugations must satisfy the criteria previously outlined for the number of slots per wavelength. Using this value for N, suitable values for ΔL can be calculated. SCATTER reads the geometry in, assigning the first three rows to their appropriate variables, ν, n and N respectively. The lengths and radii are read into separate arrays, with the number of modes at each junction ignored. Knowing the frequency allows the roots of the Bessel functions and

their derivatives to be calculated. These are defined once for each execution of SCATTER, at the execution frequency, and used throughout the calculation. Having these values pre-defined and saved into a list eliminates the requirement to constantly calculate them, thus reducing the computation time significantly.

2.4.2 Scattering Matrices Calculation

The number of modes that the radiation can be contained within as it propagates is defined by *modenum*. This is split evenly between TE and TM modes, with each containing d modes, where $d = \frac{\text{modenum}}{2}$. Using this number of modes and the geometry of the horn, the scattering matrices are calculated for each section, in a given azimuthal order. This allows the overall scattering matrices for the horn to be determined for each azimuthal order. Again, in the circular case there is no need to consider interactions between equivalent matrices in different azimuthal orders as scattering between modes of different azimuthal orders is not permitted.

Starting from the input waveguide end, where all modes to be considered are present in equal amounts initially, the scattering matrices for each section are calculated. Each section is considered to have a uniform radius throughout, and to be of the length given in the geometry file. In each section, the output from the previous section is taken as the input to the current section, and the coupling coefficients are used to determine which modes propagate to the output of the current section, and in what proportions, resulting in the scattering matrices for that section being calculated. These matrices are cascaded with the matrix VV, which is unique for that section. It adds the appropriate phase to each of the complex modes in the output scattering matrices, to account for the propagation of each mode along the finite length of that section. For each section, its set of scattering matrices is cascaded with the previous matrices for the structure resulting in an overall set of four scattering matrices that define the structure up to and including the current section. This is repeated for all sections, at which point all corresponding matrices are cascaded together, resulting in the overall set of matrices for the horn. This is repeated for all n azimuthal orders, resulting in four (modenum $\times$ modenum $\times$ n)-dimensional matrices for the horn, one for each scattering parameter. These arrays are called S_{XY}total, where X and Y indicate which of the scattering parameters is under consideration and the arrays effectively represent the four S parameters of the system over each of the azimuthal orders. The separate calculation for each azimuthal order results in different orders not mixing.

2.4.3 Phase and Intensity Calculation for the Horn Aperture and Farfield

The required fields are calculated by using the modal expressions previously discussed for the TE and TM fields as a basis set, and expanding them in terms of the coefficients of the scattering matrix that is being used to calculate the required field. If for example the horn is

illuminated at the entrance aperture and the field at the exit aperture is desired, then since SCATTER assumes illumination at port 1 and the exit aperture in this case is at port 2, the S_{21} matrix should be used for the field calculation. Another example is the illumination of a horn at its exit aperture, where the horn feeds a cavity. It may be desirable to analyse the field of the horn with respect to changes in the cavity structure. Since the aperture whose field is to be examined is the exit aperture of the horn, and this aperture is also the aperture which the horn is excited via, then in this case the S_{11} matrix should be used.

Effectively, the amount of each mode present is found from the scattering matrix that is in use, and this is used to scale the individual waveguide mode expressions which are then added in quadrature in order to obtain the overall field. The difference between aperture and farfield calculations lies in the mode set used to expand the transmission matrix. Two different sets of expressions are used, the farfield mode set being a Fourier Transform of the aperture mode set. The transmission matrix remains the same in each case. The intensity is calculated in either case by using equation 2.78

$$I = \left| \sum_{i}^{\text{modenum}} D_i \Psi_i \right|^2 , \tag{2.78}$$

where Ψ_i represents the modal field being expanded, which varies for the aperture and farfields, and D is given by $[D] = [S_{\text{xy}}]\,[A]$, where $[A]$ represents the input fields, as the scattering matrix merely forms a link between these and the output, and so the input must be included in order for the scattering matrices to have meaning.

A gives each mode an initial intensity (unity power), and the scattering matrix gives the corresponding intensity of this mode at the output (relative to the initial intensity). All individual modal contributions are summed in order to give the overall field. For example, consider a single-moded horn, supporting only the TE_{11} mode at the narrow waveguide end, the throat of the horn. Only azimuthal order $n = 1$ is considered, as the TE_{11} cannot excite any modes in other azimuthal orders, nor can the modes that are excited by the TE_{11} mode. Assuming *modenum* modes are used in the calculation, A will be of length *modenum*, with each entry equal to unity, as shown in equation 2.26. At the output side, it will be found that the transmission scattering matrix column representing this mode will have non-zero entries corresponding to all of the modes that have been excited by the TE_{11} mode scattering into them and exciting them. They will all have an amplitude of ≤ 1.

The resulting intensity pattern is stored in an array for future plotting. To calculate the phase of each of the field distributions at either the aperture or in the farfield, the argument of the inner sum in equation 2.78 is taken and stored in an array. The intensity and phase of the fields at the aperture are plotted against off-axis distance, over a user defined range. The corresponding farfield plots are obtained by plotting the data against the off-axis angle, θ. It is possible to plot the farfield and aperture field distributions using combinations of the x, y and orthogonal components of each of the waveguide modes. In SCATTER, using only the x or y

components, and not including the orthogonal modes, will result in the co-polar or cross-polar fields being plotted respectively. By using both of these components and by including the orthogonal modes as well, it is also possible to plot the full unpolarised symmetric beam. By choosing different values for the azimuthal angle ϕ, it is possible to plot cuts of these fields, or by sampling cuts in ϕ at several points in the range $\phi \in [0, 2\pi]$ it is possible to produce a full two dimensional representation of the field in question.

2.4.4 Output of All Data to a .txt File for Additional Analysis/Records

Each of the data sets above are outputted to a .txt file in column form, where one column is the calculated data, and the other is the corresponding off-axis angle/distance, as required. This allows further analysis and the possibility for the fields to be imported to software such as GRASP, and subsequently propagated through an optical system.

In this section, the implementation of the theory behind the mode-matching in the SCATTER program was discussed. SCATTER has been extensively tested and verified by comparing its results to both experimental results and results obtained using other software and modelling methods, for a variety of horns. It was found to be extremely accurate and computationally efficient. SCATTER is used as the basis for much of the computational work described in this thesis.

2.5 MODAL

In this section, the Maynooth Optical Design and Analysis Laboratory (MODAL) software package will be considered. MODAL is a piece of software being developed in order to model and analyse quasi-optical systems efficiently and accurately by means of combining optimised physical optics and mode matching routines. It features a three dimensional CAD design environment, making it easy to define and place the various components within the system. The visual environment can be explored by means of several interactive camera angles, and many physical parameters can be measured including cuts and full 3D contour plots of the desired fields at any given plane in the system.

The key design attributes for MODAL, as outlined in [38–40], are as follows.

1. **Specifically designed for the terahertz portion of the electromagnetic spectrum.** This allows the use of approximations and techniques that apply only to this region of the spectrum, resulting in more efficient calculations relative to standard software packages.

2. **User friendliness.** The graphical user interface allows the user to easily select required features and components, with the most common features readily available.

3. **Ease of system implementation.** The 3D CAD environment allows the optical systems to be designed quickly and accurately, as the system can be visually inspected at every step. It is also possible to visual beams by means of a best Gaussian fit, which is extremely useful when investigating system performance or when a system is being designed.

4. **Efficient modelling and analysis techniques.** MODAL offers a variety of methods for analysing systems, often allowing the same system to be analysed via several different methods. These methods include both a full physical optics (PO) treatment, as well as modal approaches to representing beams (Gaussian beam modes and plane wave decomposition) and horn antennas (mode matching). This selection of techniques allows systems to be modelled that a one size fits all approach (one method of analysis for the entire system) cannot be applied to, for example a horn feeding an optical system. It also allows the ability to trade accuracy for speed. An example of this would be using the modal technique to analyse a system as opposed to a full PO simulation.

5. **Cross platform.** MODAL is currently supported by both Windows and Linux system architectures, and so users of either system can use MODAL without the need for virtual boxes and the associated computational overheads. The underlying code uses industry standard software such as OpenGL for system visualisation and Fast Light Tool Kit for the user interface. This simplifies the process of porting as it does not require more esoteric software.

The functionality of MODAL is constantly being upgraded with current features refined and new features being added. Currently, when dealing with horn antennas, it is limited (to the extent of verified functionality) to waveguide sections of circular cross-section, however an efficient implementation of the rectangular mode matching technique is being worked on at present, which will allow the consideration of waveguide sections of rectangular cross-section. MODAL is also capable of utilising distributed computing, a feature which is critical when it is desired to analyse large systems accurately.

A complete description of the analysis techniques used in MODAL is given in [38], however a brief description of the principal techniques will be given, apart from the mode matching technique which has already been considered.

2.5.1 Physical Optics (PO)

For a field propagating through an optical system, the use of PO allows a more accurate description than what may be encountered in Gaussian beam mode optics in that it does not make the paraxial approximation and does not represent the field as a scalar. This vectorial representation of the field allows a more accurate description of the field in each direction of

propagation. PO uses the assumption that the surface currents flowing on a conductor are related to the electromagnetic field that is incident on the conductor. An example of this would be propagating a field from a horn antenna to a mirror and looking at the reflected field. The field present at the aperture of the horn (calculated using perhaps the mode matching technique) is analysed, and the equivalent currents flowing at the surface of the horn at the exit aperture are synthesised. This is the input plane to the PO portion of the system. Assuming that the horn presents a perfectly electrically conducting surface at the input plane, then a boundary condition is imposed such that the electric fields tangential to this plane are zero at the boundary and so only the tangential magnetic field needs to be calculated across the input plane. The induced currents are re-radiated into this field, and it is propagated to a mirror which is also treated as a perfect conductor. Using the same boundary conditions as before, currents are induced in the mirror. These are known as reflector currents, as when re-radiated they give rise to the reflected field which can then be propagated to a measurement plane or the next optical component in the system.

By considering the field as a series of magnetic sources which can be propagated by means of vector propagators, any system can be modelled. The accuracy is limited ultimately by the number of physical optics points that are considered in the calculation. Using more points increases the accuracy at the expense of execution time and memory requirements. If it is desired to reduce execution time and the use of computational resources, then a common approach that can be employed is to employ a scalar approach rather than using vectors. This scalar diffraction technique will now be briefly summarised, with full details again available from [38].

2.5.2 Scalar Diffraction

To implement scalar optics, one must represent the **E** and **H** vectors that are considered fully in the PO approach as a single complex scalar, U, that satisfies the Helmholtz equation

$$\left(\nabla^2 + k^2\right) U = 0. \tag{2.79}$$

This approach only works if the propagation medium being considered is isotropic, homogeneous and non-magnetic. As free space is typically considered, this is valid. Introducing optical components to such a medium corresponds mathematically to imposing boundary conditions on a wave that is propagating in such a medium. These boundary conditions result in coupling between the scalar components of the **E** and **H** vectors. As a consequence of this, the scalar approach introduces a degree of error at the edge of the aperture describing the optical components. For aperture areas that are large when expressed in terms of λ, these coupling effects and therefore the error introduced, will be small. This can be thought of as being an equivalent system to the paraxial approximation. In such a system, the diffraction angles caused by

the aperture are considered to be small. The two scalar approaches that follow this principle that were implemented in MODAL are the Fresnel Transform and Plane Wave Decomposition, which again are analysed in detail in [38].

2.6 Conclusions

In this chapter, the concepts of waveguide modes and mode matching were introduced for systems of circular cross-section. The theory and underlying mathematics were examined, along with the process of using the theory to determine a set of scattering matrices that can be used to relate an arbitrary input field consisting of TE and TM waveguide modes to the output field of a waveguide based structure. This allows the determination of the modal content of such structures, subject to a given set of input modes. Methods to represent the radiation patterns of waveguide structures were also introduced, along with how the calculated scattering matrices can be used to obtain the radiation patterns of the structure under test in the in-house developed mode matching code, SCATTER. MODAL, the quasi-optical design and analysis that is currently being developed in-house was also discussed.

Chapter 3

SCATTER Modifications

In this chapter, the further enhancement of SCATTER will be considered. The purpose of this was to reduce the execution time of the code which would give rise to the more efficient calculation of the scattering matrices and the radiation patterns of the horn antennas under consideration. Increasing the efficiency of the code can be critical in cases where the horn is large with respect to the wavelength being considered, and thus many sections are required to fulfil the minimum geometric sampling requirement of typically ten sections per wavelength. It is also important when the code is to be run several thousand times in the course of optimising the design of a horn antenna by means of an optimising algorithm. In this case it is clear that even the saving of a second or two per execution is large when it is considered that the code will run thousands of times.

There are several areas one could examine when looking to increase the performance of the code. These range from examining the physics and mathematics of the situation (for example recognising symmetries within the various integrals or identifying recursive relations) to altering how the code is written. For example in Mathematica one could compile functions that are called often, as this will result in Mathematica not having to re-evaluate each expression in full each time it is called. This brings clear time saving benefits. It would also be possible to rewrite the code to remove certain loop statements that are time consuming to execute, or to implement the SCATTER code in a programming language that is designed for speed of execution and to use an efficient compiler. For the purposes of the work presented here, two modifications will be discussed in this chapter:

1. The use of a suitable number of modes for each section based on the radius of the section and the wavelength of the radiation (i.e. considering the cut-off frequency of the modes), as opposed to a larger, fixed number of modes throughout the geometry

2. The use of singular value decomposition (SVD) to increase the speed of the radiation pattern calculations

These modifications will now be presented.

3.1 Dynamic SCATTER

This section deals with the implementation of the first proposed modification, as listed above. Originally, SCATTER used a fixed number of modes throughout the calculation, with this number being defined at the start. This number was found by examining the section of the horn with the largest radius and calculating the number of modes supported by this section. Several additional modes with cut-off wavenumbers beyond that supported by the horn would also be included in order to account for evanescent modes within the system. As stated in chapter 2, in order to accurately model a section of waveguide (which means accounting for evanescent modes), a number of modes equal to approximately ten times the ratio of the waveguide diameter to the wavelength, [41], is required. This requirement is then applied to the section of waveguide with the largest diameter, and therefore this number of modes will accurately represent the other sections in the system. This means that when the scattering matrices are being calculated at each junction, the same number of modes are always used regardless of the dimensions of each junction and the wavelength of the radiation. It should be noted that the requirement for the number of modes outlined above is 'overkill' in most situations, and excellent agreement with experimental results can often be found using less modes, in addition to no real change being noted in the simulated field distributions as more modes are added in simulations.

It is not necessary to consider this number of modes throughout the horn. Each section can support a fixed number of modes, and since during one execution, SCATTER considers one wavelength, this number is determined by the radius of each section. The larger the radius, the greater the number of modes that the section can support. If SCATTER is executed several times to consider a band, then a fixed radius will require more modes the lower the wavelength. To this end, when selecting the number of modes to used, it is currently necessary to consider both the wavelength and the radius of the largest section of the horn being modelled. If too few modes are used then it is likely that inaccuracies will manifest themselves in the calculation. This is due to the possibility that modes which are not being considered may physically contain a non-negligible amount of power. Is it also possible that evanescent modes may have an effect on the system. These modes do not propagate as they are not physically supported by the section, however in order to include their effects a few additional modes beyond those supported by the section should still be considered.

Assuming that a sufficient number of modes is used, then no issues will arise from an accuracy point of view, even if many additional modes that are not necessary are considered, as these will simply not be supported and so will have no impact on the system. Their consideration in the calculation will however result in a longer execution time, which is to be avoided. This is due to the fact that the 'correct' number of modes will be based on the largest section, and so the other sections will have an excess number of modes included in the

calculation. The method of fixing the number of modes is used as it simplifies the mathematics and implementation of the power coupling integrals and the cascading of the scattering matrices for successive junctions. This is due to the fact that the matrices remain square, and so it is not necessary to track the dimensions of the matrices as the calculation progresses.

The issue with keeping the number of modes fixed at a value that is sufficient (neither too large nor too small) for the largest section of the antenna is that this number is extremely likely to be highly excessive for large portions of the horn antenna. For many sections, a number far smaller than this fixed number is sufficient. For example, the aperture of the horn may require the consideration of 30 modes, however the throat may only require 2 modes, with the increase in the number of modes as one goes from section to section being only gradual. In this instance, the original implementation of the code will consider 30 modes throughout the horn, when in reality many sections will require the consideration of far less modes. This means that the calculation will take considerably longer than it needs to, as many unnecessary modes are being considered. To address this situation, the number of modes supported in a given section must be calculated. If this is done for each section, then for each junction, the number of modes supported either side of the junction will be known and so the number of modes being considered will be capable of changing dynamically for every junction considered. This process will now be examined.

3.1.1 Determining the Number of Modes Supported by a Section

To determine the number of modes supported by a section of arbitrary radius r, illuminated by a monochromatic source of wavelength λ and supplying uniform, unity power to each mode being considered, it is necessary to calculate the cut-off wavelength of a particular set of modes. One could equally deal in terms of cut-off frequencies and wavenumbers, but wavelengths will be used here. The set of modes examined is an arbitrary choice. The calculation of cut-off wavelength is a rapid process, and so an extremely large number of TE and TM modes spread over many azimuthal orders can be considered. Typically, hundreds of modes over ten to twenty azimuthal orders are considered, as reducing these numbers offers little impact from the point of view of execution time on a scale that is useful to the calculations involved. Using such a large number of modes means that it is virtually guaranteed that all supported modes will be included in this list, and so all relevant modes will be picked out and used in the calculation for each section. Using a selection of modes that only just cover the upper limit of modes that will be used to represent a horn runs the risk of leaving certain modes out which could potentially negatively affect the accuracy of the results, for example evanescent modes may end up being ignored. To this end, padding with additional modes is essential and as stated previously, it does not adversely impact upon the execution time when compared with the alternative of using a fixed number of modes throughout the entire geometry.

The cut-off wavelength for the TE_{mn} and TM_{mn} modes are given in equations 2.21 and 2.9

respectively, however these equations are now highlighted again.

$$\lambda TE_{mn} = \frac{2\pi r}{p'_{mn}} \tag{3.1}$$

$$\lambda TM_{mn} = \frac{2\pi r}{p_{mn}}. \tag{3.2}$$

p'_{mn} and p_{mn} refer to the root of the derivative of the Bessel function and the Bessel function respectively, of degree m and order n. It is clear that as r increases, i.e. the section gets wider, the cut-off wavelength of the modes increases. A mode is only supported if its cut-off wavelength is larger than the wavelength of the radiation propagating through it, and so as the cut-off wavelengths increase, more and more modes will satisfy this condition and so more modes will be supported within the section under consideration.

To determine the number of modes supported by a section, the radius of the section is taken from the list of radii that is read into SCATTER from the geometry file of the horn antenna and initially entered into equation 3.1. All TE modes that are being considered for testing (for whether or not they are supported in the section in question) then have their cut-off wavelengths generated based on this radius and this value is compared with the wavelength of radiation being used to illuminate the system. If the cut-off wavelength exceeds the wavelength of the radiation then this mode is supported and a tracker variable is incremented to reflect this. If N azimuthal orders are being considered then the tracker variable is a $1 \times N$ matrix, and each azimuthal order has the number of TE modes that are supported tracked separately. For example, if the matrix was of the form $[0, 2, 1, 0, 0, \dots]$, then this would indicate that only the first and second azimuthal orders contained TE modes that were supported in the section. There are two points to note here. The first point is that the first entry in this matrix corresponds to modes of azimuthal order zero and so the second entry corresponds to $n = 1$. Therefore the i^{th} entry in the matrix refers to the number of TE modes supported with azimuthal order $i - 1$. The second point to note is that within an azimuthal order the modes turn on sequentially, i.e. the TE_{1m} modes turn on in the order $m = 1, 2, 3, \dots$. With these points in mind, it is clear that in this case the section supports the following three modes: TE_{11}, TE_{12} and TE_{21}.

In order to account for the possibility that some modes may begin to turn on in sections that only barely don't support them, and to account for the possibility of evanescent modes, an additional number of modes are added to the calculation for each section. These additional modes must be included, as in general they will have an impact on horn performance, and so to recover the correct scattering matrices and therefore the correct radiation patterns, as calculated by using a sufficient fixed number of modes. The effect of the number of additional modes considered will be examined in section 3.1.3. The same process is repeated for the TM modes, and in this way the precise number of TE and TM modes supported in the section is

found.

3.1.2 Calculation of Scattering Matrices and Radiation Patterns

With the number of modes supported in each section calculated, it is now possible to calculate the scattering matrices that govern each junction in the horn antenna following the same principle as in the original version of SCATTER, albeit with some slight modifications related to the potentially unequal number of TE and TM modes that are now encountered. The immediate difference between the dynamic and original versions of SCATTER is that the matrices are now rectangular as opposed to square. It is also extremely likely that the number of TE modes will not equal the number of TM modes. This brings with it an additional layer of complexity, requiring careful monitoring and tracking of the numbers of each mode present in each section.

The scattering matrices govern the junction between any two consecutive sections, therefore for each junction two sections are considered. If the first section (the section to the left of the junction) is section a and the second section (to the right of the junction) is section b, then the number of TE and TM modes in section a are assigned the variables ea and ma respectively, with the total number of modes supported in this section known as $modna$, where $modna = ea + ma$. For section b, the equivalent variables are eb, mb and $modnb$.

For a given azimuthal order, the P, Q and R matrices can now be calculated, as outlined in chapter 2, where again the matrices represent respectively the power coupled between the modes across the junction and the self-coupled power of the modes on the right and the left of the junction. In the formulation used in SCATTER, columns of the scattering matrix represent the input modes to the section, and the rows represent the output modes for the operation corresponding to the matrix. Thus for the S_{11} matrix of a junction, then if the first azimuthal order is considered, the entry in row 2 of column 1 will correspond to the amount of power coupled from the TE_{11} mode back into the TE_{12} mode in that section, and so forth. R and Q will therefore be square matrices of dimension $modna \times modna$ and $modnb \times modnb$ respectively. P, since it deals with both sides of the junction, will be a rectangular matrix of dimension $modnb \times modna$.

Populating these matrices is a relatively straightforward process. The Q and R matrices only require the consideration of TE-TE and TM-TM coupling and so will turn out to be diagonal matrices. Taking the left of the junction (the R matrix), a list of length $modna$ is therefore created and populated by evaluating (see chapter 2) TE_{nl}-TE_{nl}, where n is the azimuthal order, for $l \in [1, ea]$ and $TM_{n(l-ea)}$-$TM_{n(l-ea)}$ for $l \in [ea + 1, modna]$. This list is then transformed into a diagonal matrix, resulting in the R matrix. The Q matrix is calculated in the same manner by considering the number of modes supported in the section to the right of the junction.

The P requires the calculation of TM-TE and TE-TM coupling, in addition to the coupling outlined above. This matrix will not be diagonal. The process is best described by examining the code used, which is given in the code fragment below, where *rsmall* and *rbig* assume their normal meanings according to the original version of SCATTER.

```
for j=1:1:modnb

    if(ii<((ea)+1) && j<((eb)+1))
        Piij(ii,j)=ptete(n,ii,j,rsmall,rbig);
    end

    if (ii<((ea)+1) && j>(eb))
        Piij(ii,j)=ptetm(n,ii,(j-deb),rsmall,rbig);
    end

    if(ii>(ea) && j<((eb)+1))
        Piij(ii,j)=ptmte(n,(ii-ea),j,rsmall,rbig);
    end

    if(ii>(ea) && j>(eb))
        Piij(ii,j)=ptmtm(n,(ii-ea),(j-eb),rsmall,rbig);
    end


end
```

Ultimately, the columns correspond to the modes on the input (left) side of the junction and the rows correspond to the output (right of junction) modes. In the original SCATTER, the transpose of the required matrix is calculated, consistent with the structure of the loop. This matrix is then transposed in order to obtain the required matrix. To remain consistent with this approach, the input modes are therefore the rows in the Piij matrix calculated in the code fragment above.

This fragment is part of a larger loop that cycles the iterator ii from 1 to *modna*. Bearing in mind the structure of the scattering matrices, if $ii < ea + 1$ and $j < eb + 1$, then we are considering TE modes on both sides of the junction and so TE-TE coupling is used. If j remains in this range, but ii leaves it, then on the input side of the junction we are considering TM modes, but still considering TE modes on the output side and so TM-TE coupling is used. By careful tracking of which quadrant of the matrix the entry in question is located in, the TM-TM and TE-TM coupling can also be found, as outlined in the code fragment. Once the matrix has been fully populated, it is transposed in order to obtain it in a form consistent with the R and Q matrices.

To complete the description of the section, it is necessary to calculate the V matrix that accounts for the phase shift of the modes as they propagate along the section to the right of the junction, towards the next junction that will be considered. This is again a diagonal matrix, and is populated by calculating the phase shifts of the required number of TE and TM modes for the section to the right of the junction in precisely the same manner as in the original version of SCATTER, simply taking care to account of the likelihood that there will not be an even split between TE and TM modes.

When the V matrix is calculated it is then combined with the P, Q and R matrices in the standard manner in order to calculate the four S matrices for that portion of the horn. Although the matrices are no longer square, the conditions for matrix multiplication are still met due to the order in which the matrices are multiplied. This process is carried out for each section, and as in the original version of SCATTER, after each section has had its four S matrices calculated, they are cascaded with the preceding scattering matrices in order to obtain a new set that give a description of the behaviour of the horn up to the last section considered. This is carried out until all sections have been accounted for. Again, the conditions for matrix multiplication remain satisfied during the cascading process due to the fact that the output section of one set of matrices is the input section for the next set, and so the dimensions of the matrices will match as required, considering the order that they are multiplied in.

Following the calculation of the S matrices, it is possible to obtain the radiation patterns of the horns. Assuming that a horn has been analysed with the throat located at port 1 and the exit aperture at port 2, and that the radiation pattern at port 2 is required, then the S_{21} matrix is required to be analysed. If the throat supports m modes in total, consisting of ea TE modes and ma TM modes, and the exit aperture supports n modes consisting of eb TE modes and mb TM modes, then the S_{21} matrix will be of dimension $n \times m$.

It is necessary for the purposes of the analysis of this matrix to store the number of each mode supported by the throat. With this information and the same information for the throat, then it is possible for each element of the matrix to calculate whether it is a TE or TM mode coefficient, and to then call the appropriate Bessel function, as outlined in chapter 2. In this way, the partially coherent fields in either the aperture or farfield planes can be calculated for a given azimuthal order and added in quadrature. If multiple azimuthal orders are being considered, then this procedure is carried out for every order of the matrix, with the field resulting from each azimuthal order being added incoherently in order to yield the total field.

3.1.3 Verification of Dynamic Mode Number Approach

In order to verify the approach to calculating the scattering matrices outlined above, a profiled smooth-walled horn was analysed at 100 GHz ($\lambda = 3$ mm). The horn in question will be presented in detail in chapter 5. Details of the geometry are given in column 4 of table

5.1, however the key specifications are a total length of 49 mm requiring 151 sections in order to accurately model it in SCATTER, with entrance and exit apertures of 1.604 mm and 12.969 mm respectively. At this frequency, the horn, for the purposes presented in chapter 5, is approximately single-moded, with all excited modes being considered located in the first azimuthal order. If the process can be verified for one azimuthal order, then as extension to multiple azimuthal orders merely requires additional summations, then the approach can be assumed to hold for all other azimuthal orders.

Since the radiation patterns of a horn are found using the scattering matrices, then the most straightforward method to illustrate the agreement between a scattering matrix calculated using the dynamic approach and one found using the original approach is to calculate the radiation pattern using each one. Since the matrices should communicate the same information, then if the method has been implemented correctly then the field patterns should agree. Therefore, the E-cut (0°) of the co-polar field of the horn will be found using both methods and compared in order to validate the dynamic approach.

Figure 3.1 shows a comparison between the E-cuts found using the original (30 modes, with the simulated results produced using this number of modes validated experimentally to a high degree of accuracy, as will be shown in chapter 5) and dynamic matrix calculations, for different amounts of padding within the dynamic calculation. Again, this padding corresponds to 'extra' modes considered in each section, beyond those that are explicitly supported by the section. It is clear from figure 3.1(a) that it is not sufficient to use only the supported number of modes in each section. The additional modes that are ignored in doing this may not have a huge effect on a single section, but the knock on effect of ignoring them can lead to significant inaccuracies arising as more and more of the horn is considered. For example, ignoring an evanescent mode in one section may not alter the performance of that section by a great deal, however this evanescent mode may excite other modes which propagate through subsequent sections of the horn, and so by not including the evanescent mode these modes will not be excited, resulting in an accurate representation of the true performance of the horn. It is clear from examining figure 3.1 that as more 'extra' modes are included in each section, the performance of the horn is more accurately modelled relative to the original approach. Figure 3.2 shows the number of modes supported in each section plus ten, so by subtracting ten from each value, the number of modes supported by each section of the horn can be readily found. Adding only 4 additional modes to the number supported gives rise to excellent agreement with the original approach immediately, as can be seen in figure 3.1(c) where the agreement is excellent down to approximately -55 dB. Adding further modes increases agreement at levels below this, matching the sidelobes extremely well. With the addition of 10 extra modes beyond the number supported by each section, the disagreement becomes negligible. Although this seems like a large number, a large time saving is still made, as will be shown below. These results imply that the scattering matrices agree very well, validating the dynamic calculation.

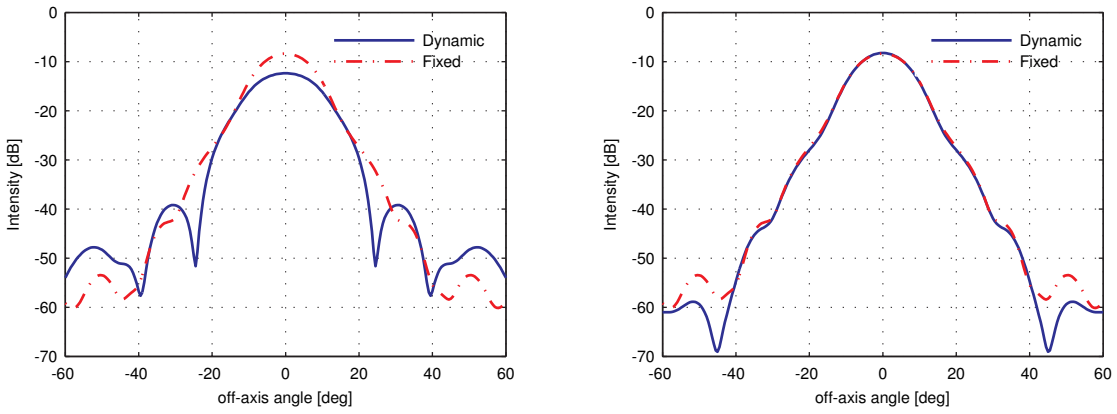

(a) 0 extra

(b) 2 extra

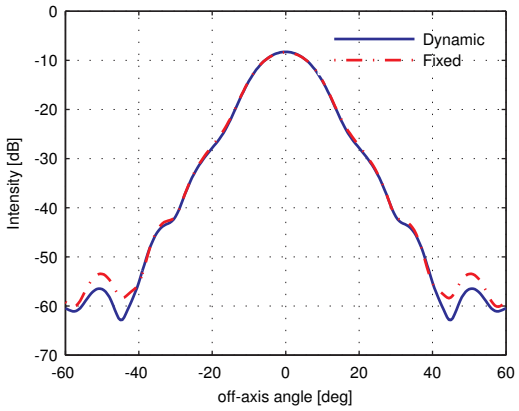

(c) 4 extra

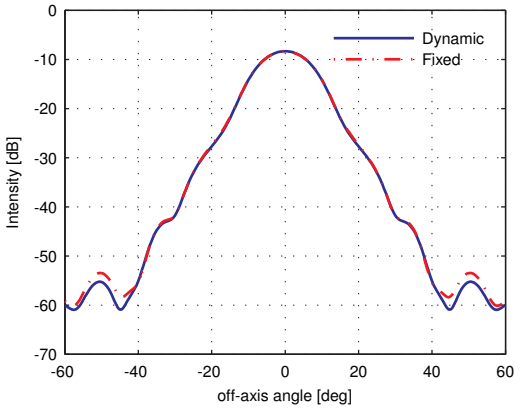

(d) 6 extra

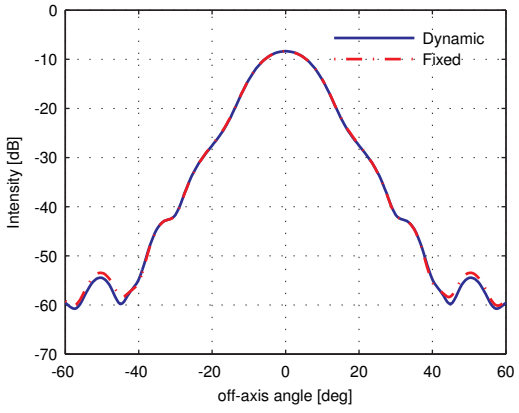

(e) 8 extra

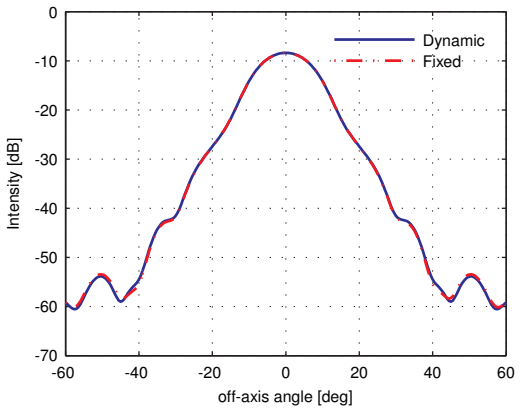

(f) 10 extra

Figure 3.1: Co-polar (farfield) cuts, calculated using original and dynamic (for different amounts of padding) matrix calculation methods

The purpose of using a variable number of modes in the scattering matrix calculations is to speed up the process. To illustrate the effectiveness of the dynamic approach in achieving this goal, table 3.1 shows the execution time of the scattering matrix calculation portion of the code for the dynamic calculation for varying amounts of additional modes added. Using the original method and 30 modes, the calculation took 56.29 s, yielding an S_{21} matrix (the matrix used to calculate the field in this case) of size 30× 30. It is clear from table 3.1 that utilising the dynamic

Extra modes	**S_{21} dimensions**	**Execution time [sec]**
0	16×1	3.05
2	18×3	4.47
4	20×5	5.92
6	22×7	7.93
8	24×9	10.24
10	26×11	13.05

Table 3.1: Execution times and overall matrix dimensions for dynamic scattering matrix calculations

approach yields a significant saving in the execution time of the scattering matrix calculation. This is true even when 10 additional modes beyond the minimum required are utilised in each section. Thus, a sufficient number of modes can be used to achieve the accuracy indicated in figure 3.1 but with significant time savings. This is useful in all calculations, particularly in ones which require several evaluations of a horn design, for example in an analysis of the horn across a band or in an iterative optimisation process.

An examination of the dimensions of the S_{21} matrices calculated using the dynamic method (table 3.1) shows that as more modes are added, the dimensions of the matrices appear to become large. The number of rows in the overall matrix is determined by the number of modes supported by the largest section, the aperture, while the number of columns is determined by the smallest section, the throat. Due to this, the number of rows in the matrix, when ten extra modes are considered, is comparable to the matrix obtained using the original approach. Although the number of columns is significantly less in the dynamic case, and this would account for some of the significant time difference, it is worthwhile to consider why the execution time in the dynamic case is not even half of that of the original case, despite the relative sizes of the matrices.

The final matrix, as previously stated, gets its dimensions from the number of modes supported by the exit and entrance apertures. It is not indicative of the number of modes used throughout the horn, with this number gradually scaling up to the final figure used. From section to section, the number of modes used is actually less than this. This is clearly shown in figure 3.2, where the number of modes considered in each section in both approaches is shown. It is clearly seen that for most sections, the number of modes considered in the fixed approach

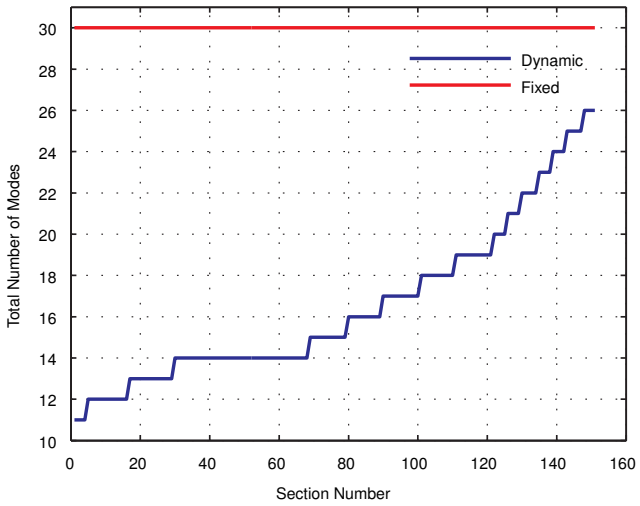

Figure 3.2: Number of waveguide modes considered in each section for both the dynamic and the fixed number of modes approaches

is far greater than the number used in the dynamic calculation. The average number of modes considered per section in the dynamic approach is 17, which is far less than the 30 considered in the fixed approach. This clearly illustrates where the large saving in time originates.

The above discussion shows that the scattering matrices for a horn antenna can be calculated in a significantly quicker manner without losing any accuracy, by calculating the required number of modes in each section of the horn antenna as opposed to using a fixed number of modes throughout. This decreases the execution of the SCATTER code significantly. In the next section, a method to enhance the speed of the field calculation will be discussed. This, combined with the more efficient matrix calculation code, will enhance the overall speed of SCATTER, giving a powerful optimisation tool. This dynamic implementation of SCATTER has also been applied to other horn geometries, and using ten additional modes was found to give similar levels of agreement to that presented above. This means that the approach can be applied with confidence to any horn geometry, and not just to the example presented herein. If one were to use the original version of SCATTER, it would be possible to save time by examining the largest section and using a fixed number of modes equal to the number of modes supported by this section plus ten. This is clear from the fact that 30 modes were used in the fixed approach in the study above, and this produced results which yielded excellent experimental agreement.

3.2 Singular Value Decomposition Based Radiation Pattern Calculations

In this section, the implementation of the second proposed modification will be discussed; the use of the mathematical method of singular value decomposition (SVD) in order to reduce the time taken to calculate the radiation patterns of the waveguide structures under analysis in SCATTER. This is important as the majority of the time when analysing a horn antenna, one requires the field patterns to be calculated either to examine them directly or to extract a related performance metric from them such as coupling, Gaussicity, plane symmetry etc. The field pattern calculation can take a large amount of time and to this end it is therefore a reasonable area to focus on when attempting to reduce the overall execution time of the calculation. A portion of this thesis deals with the issue of horn optimisation, and if the horn is being optimised with respect to a parameter that requires the calculation of field patterns, then the benefits of reducing the execution time of an individual iteration of the code are clear. It is also beneficial for example, in the bandwidth analysis of a horn antenna.

3.2.1 Singular Value Decomposition and Associated Field Calculations

The mathematical technique of singular value decomposition can be used to determine the modes present in a system. These 'singular' modes are not necessarily TE or TM modes, despite these modes being used to represent the horn, but can be combinations of the two. For example in a corrugated horn that supports only the HE_{11} mode, there would be one singular mode of that shape. These singular modes are a re-expression of the original mode set used (the TE/TM waveguide modes) in a more compact form with no loss of complex information. It reduces the original mode set used to its most compact form. The singular modes are mutually orthogonal and are independent of one another, meaning they form a new basis set with which to represent the system. Re-expressing the waveguide mode set in terms of there modes allows the system to be represented in a highly compact manner, using a few independent modes that contain all of the complex information as opposed to using many partially coherent modes where many modes make only a small contribution. The original mode set defined by the original scattering matrix typically contains many entries that are effectively zero, representing the negligible contribution of the higher order modes to the output fields. Although they make a contribution of zero or close to it, time is still required to process them. It is not possible to simply ignore these modes as they may make a small contribution, however SVD allows a compact mode set to be formed that takes into account any slight contributions from higher order modes.

If $\mathbf{S}$ is an $m \times n$ matrix, then the singular value decomposition (SVD) of $\mathbf{S}$ is defined as

$$\mathbf{S} = \mathbf{U}.\mathbf{\Sigma}.\mathbf{V}^{\dagger}, \tag{3.3}$$

where $\mathbf{U}$ is an $m \times m$ unitary matrix, $\mathbf{\Sigma}$ is an $m \times n$ diagonal matrix with non-negative real numbers on the diagonal. and $\mathbf{V}^{\dagger}$, the complex transpose of $\mathbf{V}$, is an $n \times n$ unitary matrix, [51]. The diagonal entries of $\mathbf{\Sigma}$, Σ_{ii}, are known as the singular values of $\mathbf{S}$, hereafter represented by σ_i, where $\sigma_i = \Sigma_{ii}$, and are arranged in order of decreasing value. It is usual that only the first β diagonal elements are non-zero. In this case, $\mathbf{S}$ is said to have β non-zero singular values, and if we only consider σ to consist of these entries, then the set σ is of length β. The m columns of $\mathbf{U}$, $(\mathbf{u_i})$ are known as the left singular vectors of $\mathbf{S}$, and the right singular vectors of $\mathbf{S}$ are given by the n columns of $\mathbf{V}$, $(\mathbf{v_i})$. In more familiar terms, a simple calculation shows that the left singular vectors are simply the eigenvectors of $\mathbf{S}.\mathbf{S}^{\dagger}$, and the right singular values are eigenvectors of $\mathbf{S}^{\dagger}.\mathbf{S}$. The (non-zero) singular values are also found to be the square roots of the non-zero eigenvalues of $\mathbf{S}^{\dagger}.\mathbf{S}$ and $\mathbf{S}.\mathbf{S}^{\dagger}$. Thus, there is an equivalence between the singular modes and eigen modes of a system.

The unitary nature of $\mathbf{U}$ and $\mathbf{V}$ means that their columns form a set of orthonormal vectors, which can be taken to be basis vectors for a set representing the basis set given by $\mathbf{S}$. Using equation 3.3 as a starting point, it is found that

$$\begin{aligned} \mathbf{S} &= \mathbf{U}.\mathbf{\Sigma}.\mathbf{V}^{\dagger} \\ \Rightarrow \mathbf{S}.\mathbf{V} &= \mathbf{U}.\mathbf{\Sigma}.\mathbf{V}^{\dagger}.\mathbf{V} = \mathbf{U}.\mathbf{\Sigma} \\ \Rightarrow \mathbf{S}^{\dagger}.\mathbf{S}.\mathbf{V} &= \mathbf{V}.\mathbf{\Sigma}^{\dagger}.\mathbf{U}^{\dagger}.\mathbf{U}.\mathbf{\Sigma}.\mathbf{V}^{\dagger}.\mathbf{V} = \mathbf{V}.\mathbf{\Sigma}^{\dagger}.\mathbf{\Sigma}, \end{aligned}$$

due to the fact that $\mathbf{V}^{\dagger}.\mathbf{V} = \mathbf{I}_n$ and $\mathbf{U}^{\dagger}.\mathbf{U} = \mathbf{I}_m$, where $\mathbf{I}_r$ is a square identity matrix of dimension r. The singular modes of the horn are found from the basis sets contained within the columns of $\mathbf{U}$ and $\mathbf{V}$. Given these singular modes, the input and output fields of a horn can be reconstructed by using them as a basis set for expanding the waveguide modes. From the theory of singular value decomposition, it is known that $\mathbf{S}\mathbf{v}_i = \sigma_i \mathbf{u}_i$, where σ_i is the i^{th} singular value of $\mathbf{S}$ and $\mathbf{v}_i$ and $\mathbf{u}_i$ are the i^{th} columns of $\mathbf{V}$ and $\mathbf{U}$ respectively. Bearing in mind that $\mathbf{S}$ is a scattering matrix that relates the input of a system to some output, it is clear that the $\mathbf{v}_i$ can be used as a basis set to represent the input modes of the horn, and that the $\mathbf{u}_i$ can be used to represent the output modes, when scaled by the corresponding singular value. This observation is now used to analyse the output field of a horn from an SVD point of view.

By noting that the $\mathbf{v}_i$ can be used to expand the input field, and by assuming an appropriately dimensioned matrix of coefficients, $\mathbf{A}$, it is observed that the input field is given

by

$$\begin{aligned} \mathbf{E}_{in} &= \Sigma_i A_i \mathbf{v}_i = \mathbf{V}.\mathbf{A} \\ \Rightarrow \mathbf{E}_{in}^{\dagger}.\mathbf{E}_{in} &= \mathbf{A}^{\dagger}.\mathbf{V}^{\dagger}.\mathbf{V}.\mathbf{A} = \mathbf{A}^{\dagger}.\mathbf{A} = \Sigma_i \left|A_i\right|^2, \end{aligned}$$

which is simply the total power in the input beam. The field on the output side is found by inserting the scattering matrix between the coefficient term and the singular vector in the field input calculation, as this transforms the singular vector from the input side to the output side due to the fact that the scattering matrix controls this operation and is now acting on the input modes. The output field is thus given by

$$\begin{aligned} \mathbf{E}_{out} &= \Sigma_i A_i \mathbf{S}.\mathbf{v}_i = \mathbf{S}.\mathbf{V}.\mathbf{A} = \mathbf{U}.\mathbf{\Sigma}.\mathbf{V}^{\dagger}.\mathbf{V}.\mathbf{A}. = \mathbf{U}.\mathbf{B} \\ \Rightarrow \mathbf{B} &= \mathbf{\Sigma}.\mathbf{A} \\ \Rightarrow B_i &= \sigma_i A_i. \end{aligned}$$

The matrix $\mathbf{B}$ represents the output mode coefficients, which, as shown above, are related to the input coefficients simply by the singular values of $\mathbf{S}$. Calculating the power in the output beam gives

$$\begin{aligned} \mathbf{E}_{out}^{\dagger}.\mathbf{E}_{out} &= \mathbf{A}^{\dagger}.\mathbf{\Sigma}^{\dagger}.\mathbf{U}^{\dagger}.\mathbf{U}.\mathbf{\Sigma}.\mathbf{A} \\ &= \Sigma_i \left|\sigma_i\right|^2 \left|A_i\right|^2 = \Sigma_i \alpha_i \left|A_i\right|^2. \end{aligned}$$

It is clear from the above discussion (in particular the form of the output coefficients), that if the input field is considered to be expanded in terms of $\mathbf{v}_i$, the input mode coefficients, then the only modes that propagate to the output side with a non-zero amount of power are the ones whose singular values are non-negligible. The individual $\mathbf{v}_i$ are scattered into the corresponding individual $\mathbf{u}_i$ by these singular values, with no inter-modal scattering, and so there is no scattering between singular modes. The individual modes are therefore independent, and so are incoherent, meaning they propagate independently with no phase interaction. This relationship is governed by the diagonal matrix, $\mathbf{\Sigma}$, which by virtue of it's diagonal shape, only allows input-output scattering between corresponding modes. Modes with a corresponding singular value in this matrix that is non-negligible propagate through the system, and these are the modes of the horn that carry power independently of each other. These are the independent channels of power supported by the entire horn. The number of non-negligible singular values a horn has is therefore equal to the number of modes propagating in the horn. Using this number, β, the field on the output side can be found by using the first β $\mathbf{u}_i$ that have these non-negligible singular values, i.e. the output modes that are supported at the aperture of the horn. This means that the entire matrix need not be considered when calculating the fields, but rather only column subsets of the matrix. This brings significant time savings in the field

calculation, as will be shown in section 3.2.3.

3.2.2 Implementation of SVD based Field Reconstruction

An SVD based method of calculating the field patterns of any horn simulated in SCATTER has been implemented in both Matlab and Mathematica versions. This will now be described. The example used to illustrate the implementation assumes that the S_{21} matrix is to be used to calculate the radiation pattern, however the method equally applies to any of the matrices. The matrix was calculated using the standard mode matching method, in both original and dynamic forms. The code was developed for multi-moded horns, however the method will be outlined for calculating the radiation patterns within a single azimuthal order. The extension to multi-moded horns merely involves calculating the patterns for each azimuthal order and adding them in quadrature.

The azimuthal order under consideration in this example is $n = 1$, however due to MATLAB numbering conventions and the method in which the scattering matrices for each order are stored, it is necessary in the code to set $n = 2$, with the various functions compensating for this internally. Specifically, this is due to the presence of the $n = 0$ azimuthal order and the indexing requirements of Matlab. As such, when a value for n (azimuthal order) is referenced in this description, it follows the formula $n =$ azimuthal order $+ 1$. In MATLAB, the scattering matrices are three dimensional arrays, and scattering matrices are calculated for all azimuthal orders from zero to the required number, i.e. n runs from one to the maximum azimuthal order plus one. In this case, we therefore consider the scattering matrix where $n = 2$.

This matrix is calculated as outlined previously, and decomposed using the SVD command. The diagonal entries from the $\mathbf{\Sigma}$ (singular value) matrix are isolated in a list, with all negligible values (< 0.000001) then removed. This cut-off figure is a variable, and once it is set to be reasonably small then the correct results will be achieved as modes with a small amount of power in them will not result in a significant contribution to the field pattern. The singular values represent the amplitude of the modes, and so the deleted terms represent the even smaller physical contribution in terms of power once they are squared, and so can clearly be ignored due to their negligible contribution to the field distribution. The indicated cut-off is sufficient to achieve this. The number of remaining singular values is stored, and this is the number of modes supported by the horn in the $n = 1$ azimuthal order. The matrix $\mathbf{U}$ is also stored, which contains the singular vectors associated with the mode set of the output of the scattering matrix, which in this case corresponds to the mode set at the exit aperture of the horn.

The actual field calculation follows the same structure as the original version of SCATTER, the only real difference being in the matrices that are called. In the case of the farfield, a range of off-axis angles is defined, and at each point the desired electric field is calculated, using

Fourier transformed Bessel functions, either as a one dimensional cut or as a series of cuts to yield a full two dimensional plane. In the aperture plane, the points are simply defined by an off-axis distance and found using Bessel functions. For the time being, it is assumed that all polarisations are included in the field calculation, i.e. x, y and orthogonal modes. If other cuts such as co or cross-polar are required, then the orthogonal modes are simply removed, and one of the x or y components, as required. With the polarisation definition used in SCATTER, the x-components yield the co-polar plane, and the cross-polar plane utilises the y-components.

The script for calculating the radiation pattern reads in the list of singular values that contain a non-negligible amount of power, the number of these values that there are, and the **U** matrix. It sets n equal to the required number (i.e. to the azimuthal order in question plus one, so in this case 2), and uses this as an input to the functions that calculate the shapes of the waveguide modes that are being expanded by using the singular vectors as a basis set. This new basis set still works in terms of TE and TM modes, it merely expresses all of the contributions to these modes in a compact form, so the same Bessel functions are still valid in this approach.

The script considers one set of basis vectors at a time, i.e. one column of **U** at a time. Each of these columns corresponds to each of the $\mathbf{u}_i$ that were introduced above, so $\mathbf{u}_1$ corresponds to column 1 of **U** etc. If the horn contains β propagating modes, then the script starts with the first column of **U**, and saves this in a temporary one dimensional matrix known as *evo* (see the code fragment below), that is weighted by the singular value corresponding to that mode as this normalises the singular vector to have the correct amplitude. If *modenum* is the total number of waveguide modes being considered, and $d = \frac{modenum}{2}$, then the first d elements of *evo* correspond to the complex weighting coefficients associated with the *TE* modes, and the remaining d entries weight the *TM* modes. This architecture assumes that a fixed number of modes were used throughout the calculation. To integrate this with the dynamic matrix calculation method, the code fragment outlined in section 3.1.2 is modified to call the **U** matrix and the list of singular values, as is carried out below.

In the standard calculation, the S_{21} matrix replaces *evo*, and each column represents a complex weighting of the corresponding waveguide mode. By considering all columns, and adding all columns in quadrature, the correct amount of each mode will be included in the radiation pattern. In *evo*, the same information is represented using fewer columns, for example, all of the information pertaining to the TE_{11} mode is contained in only a few entries, as opposed to many. Each entry in either matrix effectively weights the corresponding waveguide mode. For example, the first entry in *evo* (or S_{21}) weights the x, y and orthogonal components of the TE_{11} mode, the $d+1$ entry weights the equivalent components of the TM_{11} mode and so on. With *evo* fixed, a for loop cycles the iterator *ik* from 1 to d, or the equivalent value in the dynamic case. The following is the loop that is executed over this cycle, which calculates the amount of each component of each waveguide mode that is present at the off-axis angle in

question, and then calculates the intensity at that point. This therefore relates to a farfield calculation, assuming a fixed number of modes.

```
for jj=1:jjj %jjj is number of singular values
evo(1:2*d)=U(1:2*d,jj)*sval(jj)

for ik=1:d
p1=p1+exteffn(h,ik,rr,phi,b)*evo(ik);
q1=q1+extmffn(h,ik,rr,phi,b)*evo(d+ik);
p2=p2+eyteffn(h,ik,rr,phi,b)*evo(ik);
q2=q2+eytmffn(h,ik,rr,phi,b)*evo(d+ik);
p3=p3+extefforthn(h,ik,rr,phi,b)*evo(ik);
q3=q3+extmfforthn(h,ik,rr,phi,b)*evo(d+ik);
p4=p4+eytefforthn(h,ik,rr,phi,b)*evo(ik);
q4=q4+eytmfforthn(h,ik,rr,phi,b)*evo(d+ik);
end

ans1=ans1+abs(p1+q1)^2;
ans2=ans2+abs(p2+q2)^2;
ans3=ans3+abs(p3+q3)^2;
ans4=ans4+abs(p4+q4)^2;

end
```

The p and q terms are set to zero at the start of the calculation for each point (during the calculation they track the total amount of the relevant component present in the field). h assumes the value of $n+1$, rr is the distance off-axis associated with the off-axis angle, phi is the azimuthal angle and b is the aperture radius. This is repeated β times, in each case evo is the jj-th column of $\mathbf{U}$, scaled by the corresponding singular value. The p and q terms ensure that a running total is kept of the total amount of each waveguide mode present, in terms of the x and y components of both TE and TM modes, as well as their orthogonal polarisations.

At the end of this section of the overall loop, the total amount of each component present a distance rr off-axis is known, and so the components must be added correctly in the last section of the loop that is illustrated. Corresponding TE and TM components are spatially coherent, and so they are added together directly, allowing the complex components to mix. The script calculates the intensity directly, so the sum of each corresponding pair of components is squared. This is done for each column of $\mathbf{U}$ being considered, with the values obtained for each column all being added together.

At the end of the overall loop, we are left with $ans1$, $ans2$, $ans3$ and $ans4$, which respectively correspond to the intensity of the x, y, x-orthogonal and y-orthogonal components of the

combined TE and TM modes. To obtain the value of the total field (or component thereof) at the spatial point in question, the required components are added directly. This is repeated for the entire set of points under consideration to retrieve the desired field.

3.2.3 Verification of SVD Calculation

The SVD approach offers a quicker computational time than the standard calculation for the following reason. In the standard calculation, all columns of the scattering matrix (realistically, upwards of 20) are considered. In reality, most of these columns result in a contribution that is zero, however the calculation is still carried out. This unnecessary computation adds significant time to the execution of the script, however it must still be carried out. In the SVD approach, a more compact mode set is used, whereby all contributions are condensed into the first β columns of a reduced unitary matrix, meaning that only these columns need be considered, as the remaining columns contain negligible power, as indicated by their singular values. In the worst case scenario, β will be limited to the number of columns in $\mathbf{U}$ (the number of modes under consideration in the calculation) and so there is no saving in time, but there will be no increase in the time taken either.

Typically, for a single-moded horn, β is one, and so only one column is considered in the field calculation, as opposed to every column. This results in a significant saving in computational time. Studies carried out on multi-moded horns further emphasise this point. A corrugated horn was used in this analysis, with an aperture radius of 7.1345 mm and of length 69.9 mm, at a frequency of 100 GHz. Eleven azimuthal orders were considered with twenty waveguide modes in each, resulting in the field calculation using 220 columns when the fields were calculated using the standard method. The SVD approach for this horn yielded only three propagating modes across the entire 11 azimuthal orders (when a cut-off of 0.0001 was used), and so only three columns had to be considered in total, giving precisely the same result as in the standard calculation, but in a fraction of the time.

Some sample times are illustrated in table 3.2 to demonstrate the reduction in computation time when the SVD method of field calculation was applied to the corrugated horn. The calculation was carried out on a Dell Optiplex 780 PC, with an Intel core 2 duo, 2.93 GHz processor, running 64 bit Windows 7 with 3 GB of RAM. In this case, where there are so many

Field Pattern	**Original method** [sec]	**SVD method** [sec]
Co-Polar ($\phi = 0°$)	171.26	3.90
Cross-Polar ($\phi = 0°$)	171.83	4.01
E (with orthogonal modes, $\phi = 0°$)	348.10	8.35
H (with orthogonal modes, $\phi = 90°$)	344.07	8.12

Table 3.2: Comparison of computation times for various field patterns of a multi-moded horn

azimuthal orders to be checked, the SVD method brings about significant time savings, taking only ≈3% of the time taken by the traditional approach. To calculate the field including all polarisations takes significantly longer, of the order of twice as long. Whilst this brings with it a large increase in execution time using the original method, using SVD allows the calculation to still be carried out in under nine seconds. The large savings in time are clearly of benefit in cases where multiple executions of the code are required, for example in bandwidth analysis or as part of an optimisation process.

The time saving does not come at the expense of accuracy. The results obtained using the SVD method agree excellently with those obtained using the original calculation method. Figure 3.3 shows a comparison between both approaches for different cut-off levels used in the SVD calculation. The lower the cut-off threshold, the more singular modes are included

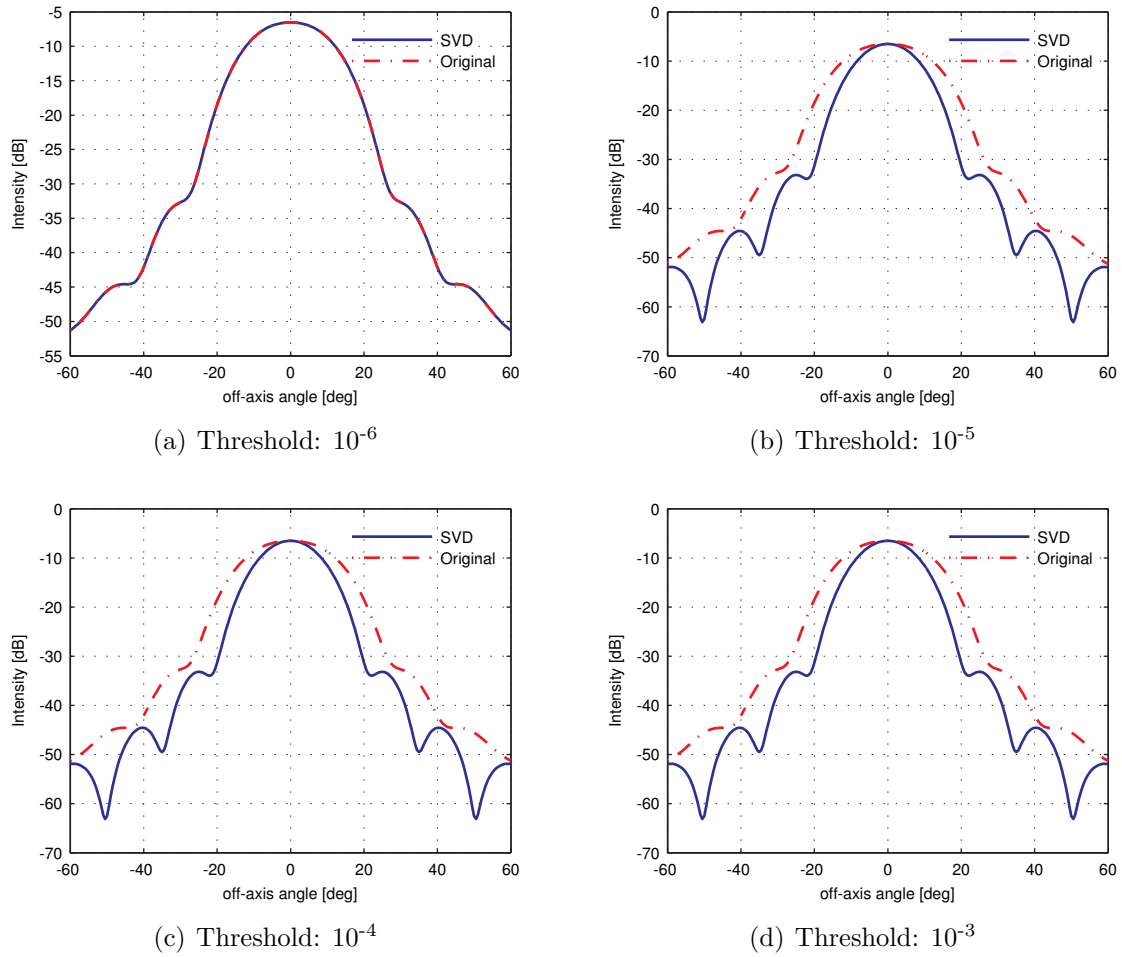

(a) Threshold: 10^{-6}

(b) Threshold: 10^{-5}

(c) Threshold: 10^{-4}

(d) Threshold: 10^{-3}

Figure 3.3: Co-polar fields, calculated using original and SVD based methods for different SVD thresholds

in the field calculation. If the cut-off is set too high it may result in excluding modes that make a meaningful contribution to the radiation pattern. It is clear from figure 3.3 that setting the threshold level to the order of 10^{-6} is sufficient to retain accuracy. This figure has been observed to be sufficient in all situations encountered, and so will be used in the work

presented. Considering that the power contained within a mode is indicated by the square of the corresponding singular value, modes with a singular value of less than this value will contain an amount of power that is so negligible it can be safely ignored. This is consistent with observations. When the threshold is raised to 10^{-5}, the field deviates from what is expected. This is due to the fact that the mode contributing the smallest amount of power, located in $n = 3$, is now no longer included. It is clear that this mode makes a significant contribution to the field, particularly off-axis, and so this illustrates why care should be taken to set the threshold suitably low. As the threshold is further increased, the beam pattern remains constant as the modes that continue to contribute (from $n = 1, 2$) both have singular values of >0.99. These dominant modes are never in danger of being ignored, however it is to be noted that they too could be removed by setting a suitably high threshold.

3.3 Conclusions

In this chapter, two modifications made to SCATTER in order to increase its efficiency were presented. They were discussed and their performance verified relative to the original version of SCATTER. The alterations in question were the use of a variable number of modes throughout the horn geometry being analysed as opposed to a fixed number of modes throughout, and the use of the singular value decomposition matrix operation to represent the scattering matrices in a more compact form. The use of a variable number of modes decreases the execution time of the scattering matrix calculation portion of SCATTER, whilst the use of SVD allows a decrease in the execution time of the radiation pattern calculation. Each modification was shown to significantly reduce the execution time of the corresponding portion of SCATTER compared to the original version, and so when combined they give rise to a much more efficient overall modelling process for horn antennas.

The time savings have clear benefits in several areas. For single instance executions of SCATTER, results will be obtained more quickly. This is particularly true for horns which are electrically large (large in terms of wavelength) and/or for high frequencies. In these cases, the sampling requirements of the feedhorns require that a significant number of sections are used, and so the execution time can be high. With these modifications, the execution time is significantly reduced. The savings are also useful in cases where a single execution may not be time consuming, but where several executions may be required. Examples of this are the bandwidth analysis of a horn where the entire process is repeated for each frequency, and also in an optimisation algorithm where the process must be repeated for each successive horn design ,perhaps numbering in the thousands. This latter application will be of relevance in the coming chapters.

Chapter 4

Optimisation Algorithms and Implementation

The aim of the optimisation process described in this thesis is to design a horn whose performance is optimised with respect to the parameters required for future CMB missions, using a standard desktop PC and standard MATLAB functions, within a reasonable timeframe (less than a day). In this chapter, the underlying process of designing a feedhorn with its performance optimised for a particular application (future CMB/Far-IR missions) will be examined. This includes analysing the various optimisation algorithms that could be used, along with seeing which optimisation metrics produce horn designs which meet the given design goals best, the limitations of the various horn designs and algorithms, etc. To this end, this chapter will analyse the limitations of using each algorithm/optimisation metric to see which results in the best overall design performance. In order to make valid comparisons, a Gaussian profiled horn antenna and spline based design will be optimised with respect to (individually), cross-polar levels, return loss, and co-polar beam symmetry. Each of these optimisations will be carried out using genetic and simulated annealing global optimisation algorithms. It should be noted that the starting point of each of these algorithms is not a specific horn design, with this design being adjusted to a more optimum configuration to improve performance in a certain area. Rather, the starting point is selected by the optimiser using a range of user defined values for each of the parameters that will be used to define the profile of the horn. As such, the output of the optimiser, the optimised horn, is not an optimised version of a particular horn that was used as a starting point, but rather a close approximation to the horn that gives the best performance possible with respect to a particular performance characteristic, within the range of values allowed for the horn parameters.

These optimisation routines were selected as they are easily available on the Matlab platform in an efficient implementation, and are global as opposed to local, meaning that they will attempt to locate the best design as opposed to the best design in the immediate neighbourhood of the starting parameters chosen. The metrics selected above relate to the overall goal of this

design process, as will be presented in chapter 5, namely to design a horn antenna for future Far-IR and CMB missions. Excellent horn performance is required in terms of these horn performance characteristics, and so this chapter will investigate which optimisation algorithm gives the most efficient result, along with which horn performance metric is best to optimise with respect to, in order to retain good overall performance in all required areas. The horn designs that were chosen were selected due to the fact that, as will be shown when the horn geometries are examined, they have yielded good results for other groups when designed using more brute force optimisation techniques requiring large amounts of computing power. The challenge faced in this project is to achieve good optimisation results using standard computing power, but still no design tricks, so these horns were a natural starting point.

The ability to design a horn subject to various performance requirements is extremely important. This is clear, as different applications will require strength in different areas. For example, if a horn was required to measure a given signal distribution, then the horn could be designed so as to maximise the coupling of its radiation field to this distribution. If the application required high directionality then a narrow beam and/or extremely low sidelobe levels would be required. Another application may require that the horn retain various characteristics over a large frequency band as opposed to a narrow band or spot frequency. Often, a design will not exist for the desired application, or it will exist but be too expensive or some other unacceptable penalty such as size or weight. In this case it would be necessary to design the horn oneself, however this can be time consuming when it comes to consulting the literature in order to discover the design tricks required in order to obtain a satisfactory design. The other alternative is to use a brute force computational approach, however this requires significant amounts of time and computational power as it would require that the horn geometry be sampled vary finely with all parameters allowed to vary.

An ideal solution would be the ability to design a horn for a given application without having to use super computers or wait days/weeks for the result and without having to use any tricks in the design process. The purpose of the optimisation work presented in this thesis is to achieve exactly that, namely to design a horn for a given purpose that has a simple and easy to manufacture geometry, that relies only on standard desktop computing power and standard, easily obtainable, generic MATLAB algorithms.

The scientific application of the horn whose design is the ultimate goal of the optimisation work presented in this portion of the thesis is in future far-IR and CMB radiation missions. As outlined in chapter 1, in order to further constrain the various inflationary models that currently exist and to better understand the origin and evolution of galactic structure and the Universe as a whole, it is necessary to study the polarisation of the B-Mode component of the CMB anisotropy. Such measurements are polarisation sensitive and take place at a very low level, as the desired signal is extremely weak. Naturally, this places significant demands on the performance requirements of any measurement systems.

There are various designs proposed for such sensitive systems. One such design is that which was proposed for the SAFARI instrument, as discussed in section 1.4. It involved an array of feed horns located in the focal plane of the system optics that feed a Transition Edge Sensor (TES) detectors located near the aperture of a bounded waveguide cavity. The number of horns is necessary in order to couple as much power into the system as possible, as due to the low signal levels involved, the alternative is a long integration time which places time constraints on any missions due to the finite supply of coolant for the optical components that would be located in a cryostat. For a CMB mission, these horns would require for example extremely low levels of cross-polarisation, a highly symmetric co-polar beam with low sidelobe levels, low levels of return loss etc. The obvious solution to this is to use a corrugated horn, as this offers excellent performance in these areas in addition to many more, and all over a relatively large bandwidth, [45]. Corrugated horns are however relatively expensive, and carry mechanical and thermal penalties rendering them far from ideal for an array. A better alternative would be to design a smooth-walled profiled horn that is easier to manufacture, and therefore cheaper, that gives the desired performance in the required areas, at the expense of performance in the areas where it is not required.

The ultimate goal of the part of this thesis dealing with the optimisation process is to design a horn for the above applications (in the specific case of a CMB mission), whilst using standard computing power (Dell Optiplex 780 PC, with an Intel core 2 duo, 2.93 GHz processor, running 64 bit Windows 7 with 3 GB of RAM.), in a reasonable timeframe, and without using any tricks or known solutions in advance, and to see how well the horn achieves the design goals despite the simple and computationally economic approach taken. To this end, the various optimisation algorithms will now be examined.

4.1 Optimisation Algorithms

In this section, a brief description of the two algorithms listed below is given, along with details on the implementation of these algorithms in MATLAB. The two algorithms to be considered (assuming that the function being optimised is one of several variables) are

1. Genetic Algorithm
2. Simulated Annealing

These algorithms were selected for several reasons. They both solve the function that is being optimised for the global minimum. The optimum horn design is being sought in each case, and so if a local solver such as the Nelder-Mead algorithm were used, then the optimal solution may not be found, as the solver would get stuck at a local solution. This algorithm was used with reasonable success in early optimisation attempts where the system being optimised was simple and the approximate solution was known in advance, with an optimisation being carried

out on such systems in order to assess the effectiveness of the mode matching code, as modified for use with optimisation algorithms. In order to avoid the situation whereby the algorithm gets stuck at a local solution, it would be necessary to know the approximate solution in advance, which negates the objective of this design process, which is to calculate a suitable horn geometry using a generic starting point that only specifies the lower and upper limits for each of the horn parameters so as to keep the horn geometry within certain physical limitations. Additionally, due to the vast number of factors involved in determining the performance of a feedhorn from its geometry and the complex relationships and interactions between the different design parameters in this context, the optimisation process is a highly complex and non-linear problem. These algorithms lend themselves well to such problems as they are ideal for multi-variable, non-linear, chaotic functions. Both algorithms are also efficient optimisers, which assists in the requirement that the process be carried out in a reasonable amount of time on a desktop computer.

In each case, the goal of the optimisation algorithm is to minimise the value of some 'cost' function. The cost function is a function of one or more variables, in this case more than one, that returns potentially several values, although only one return value will be used in order to minimise the computational load. In the case of horn optimisation, the value returned by the cost function represents the performance of the horn with respect to some metric, for example peak cross-polar power, beam symmetry or return loss. The input variables that are passed to the cost function are what describe the horn geometry. By altering the values of these variables, the horn geometry changes, thus each iteration of the algorithm amounts to a different horn design being evaluated with respect to the metric that the horn is being optimised for. The optimisation variables may define the horn in terms of the radius and length of several sections, which form a horn when combined, or some design parameter such as flare angle or slant length. The precise parameters and optimisation metrics used depend on the horn design being used as a starting point and the desired result of the optimisation process in terms of the performance of the horn. The parameters are varied by the algorithm until the horn assumes a shape that minimises the cost function that represents the value of the relevant metric. The structure of the cost function itself is simple, involving the generation of the horn geometry from the input variables, the calculation of the scattering matrices using the SCATTER code with modifications described in chapter 3, and then from this the required metric is calculated and outputted as the return value of the cost function.

The two algorithms listed above are now examined.

4.1.1 Genetic Algorithm

Introduction

The first optimisation algorithm to be examined is the genetic algorithm. This is a global optimisation method that uses rules based on natural selection in order to solve both constrained and unconstrained optimisation problems. It works by taking a pool of possible values for the input values and repeatedly modifies them in order to obtain the minimum value of the cost function, i.e. the optimum solution, [52]. It selects individual sets of values for the input parameters at random from a given pool of possibilities and uses the best performing ones to produce the next pool of values. Each new pool of values is known as a generation, and each generation slowly approaches an optimum solution, as the best candidates from each generation are chosen to populate the next generation. The algorithm allows cost functions which are discontinuous, non-differentiable, stochastic or highly non-linear to be optimised, which is precisely what is required in the context of horn optimisation problems. It can also be applied to mixed integer optimisation, where some components are restricted to be integer values.

Classical Algorithm	Genetic Algorithm
Each step generates a single point, with the sequence of points approaching the optimum solution.	Each step generates a set of points, with the best point in each set approaching the optimum solution.
A deterministic calculation determines the next point.	The next set of points is found using a calculation involving random number generators.

Table 4.1: Some differences between classical and genetic optimisation algorithms

Some differences between the genetic algorithm and the more standard derivative based optimisation algorithms are illustrated in table 4.1.

Terminology

Some terminology used in relation to genetic algorithms is now given.

Fitness Functions

Fitness function is the name given to the cost function in the context of genetic algorithms. The algorithm searches for the minimum of this function. The fitness function is a function of one or more variables, which correspond, in the context of this work, to the parameters which control the physical geometry of the horn. The parameters used depend on the horn type in question, as outlined earlier. The output of the fitness function corresponds to the

value of some measurable quantity associated with the horn, for example cross-polar power or Gaussicity (the coupling of the horn to a particular Gaussian beam).

Individuals

If a test point can have the fitness function applied to it, it is known as an individual. If for example, the fitness function were a function of five real variables, then a vector containing five real numbers would be an individual, with the output of the fitness function having evaluated the individual being called the score of that individual. Individuals are also sometimes known as genomes, with the vector entries known as genes.

Populations and Generations

The set of individuals in any one iteration of the algorithm is known as a population. Typically, if the fitness function contains m input variables, and the population size is selected to be n, then the population can be represented by a $n \times m$ matrix. The same individual is permitted to appear multiple times in a given population. Each iteration of the genetic algorithm (GA) produces, by computation involving the current population, a new population. Each successive population is known as a new generation.

Diversity

In a given population, the average distance between the individuals is known as the diversity. A large average distance is said to have high diversity, and so on. Diversity is an essential parameter in the GA, as it determines the size of the parameter space that it searches. If the diversity is set too low, it may miss the global minimum and settle at a local minimum.

Fitness Values and Best Fitness Values

Each individual has a fitness value or score, which is the value of the fitness function evaluated for that individual. The best fitness value in a population is the lowest fitness value for a given population.

Parents and Children

When creating the next population, the GA selects certain individuals which become known as parents. Usually, the algorithm selects parents with good fitness values compared to the other individuals within the population. Using various rules, which will shortly be discussed, the parents are used to create the individuals for the next generation, known as children.

Overview of Algorithm

The main steps involved in the genetic algorithm will now be discussed. Below, a summary of the major steps is given, and each of these steps will subsequently be expanded upon.

1. Algorithm initiation: A random initial population is created.
2. The algorithm iterates repeatedly, creating a new population each time. At each iteration, the best performing individuals of the current generation are used to create those for the next population. To determine the new population, the following steps are observed:
    - Each individual in the current generation is scored by determining it's fitness value.
    - The fitness values are scaled so as to make them span over a more computationally usable range.
    - By using various selection rules, based on their fitness scores, certain individuals are chosen as parents.
    - Individuals with particularly low fitness values are chosen as elite, and are passed directly to the next generation.
    - The remainder of the next generation is filled by producing children from these parents. Children are produced by making random changes to a single parent (mutation), or by combining the vector entries of a pair of parents (crossover).
    - This new generation populates the next iteration of the algorithm.
3. This process continues until the stopping conditions for that particular execution of the GA are met.

The individual steps are now examined.

Initial Population

This is the first step of the algorithm, in which it creates an initial random population. The population size is set by the user, and in MATLAB it is twenty by default. The user should supply an initial range to the algorithm that determines the range of values which the genes in any population can assume. If the minimum of the function is approximately known, then clearly the initial range should include this point, with it ideally lying approximately in the middle of the initial range. The GA should still function correctly with a less than ideal selection of the initial range, however the computational effort required will be far greater.

Creation of the Next Generation

In every iteration of the algorithm, the current population is used to create the children that make up the next generation. A group of individuals are selected as parents, and their genes are

used to create the new children. It is usually the case that the parents are the individuals with better fitness values. The precise function that the GA uses to select the parents is variable. Three types of children are created, and together these form the new population. Physically, in this context, each generation corresponds to a pool of parameters that define the horn, and so each individual in the generation corresponds to a horn geometry. The performance metric in question is analysed for each horn, and the optimal solutions are retained and used to produce the next set of horns. The different types of children are now listed.

- **Elite children** are the individuals with the best fitness values, and get passed onto the next generation unchanged.
- **Crossover children** are created by selecting two parents and combining their genes. The child vector is defined, and for a given coordinate within this vector, the crossover function randomly selects the corresponding gene from one of the parents assigned to the child.
- **Mutation children** are created by making random changes to the genes of one parent. This is accomplished by adding a random vector from a Gaussian distribution to the parent.

The number of each child created can be specified by the user or determined by MATLAB.

Stopping Conditions

The GA uses a number of conditions when deciding on the stopping point. The options available in MATLAB for example, are given below. The algorithm can be set to stop:

- **Generations:** when the number of generations exceeds a certain value.
- **Time Limit:** after having run for a certain number of seconds.
- **Fitness Limit:** when the best fitness value in a population is less than or equal to the fitness limit.
- **Stall Generations:** when the weighted average change in the fitness function value over "stall generations" number of generations is less than some function tolerance.
- **Stall Time Limit:** if after a stall time limit number of seconds, there is no improvement in the value of the best fitness value.

If any one of these conditions is met, the algorithm stops. [52] gives details on how to customise all parameters relevant to the GA.

MATLAB Implementation

Using the default options for the GA, it can be run from the command line with the syntax

```
[x fval]=ga(@fitnessfun,nvars)
```

fitnessfun is the fitness function that is to be minimised. The @ symbol creates a function handle that directs the GA function to the desired cost function, which is a separate file. *nvars* specifies the number of independent variables that are required as an input to the fitness function. For the output, x is the final point that gives the best fitness value for the population in which the GA was stopped, and *fval* gives the value of the fitness function at x. This allows the user to immediately see the result of the optimisation process. It is possible to generate additional output arguments which display the reason for the algorithm being terminated (which stopping condition was met etc.), information about the performance of the algorithm during each iteration, the final population, the final scores etc. By including a variable named *options*, the values of the previously discussed parameters can be changed from the default values.

The second optimisation algorithm listed, simulated annealing, will now be examined.

4.1.2 Simulated Annealing

Introduction

The simulated annealing (SA) algorithm is a method that solves unconstrained and bound constrained optimisation problems, [52]. It is based on the physical process of heating a material and then allowing the temperature to slowly reduce, thereby decreasing defects in the material and minimising the total system energy. Each iteration of the algorithm involves a new test point being randomly generated. The distance between this point and the current point (the extent of the search), is based on a probability distribution that has a scale which is proportional to the 'temperature' of the algorithm. If the new point lowers the cost function value, it is accepted. Points which raise the cost function value are also accepted to within a certain tolerance based on the probability distribution. This latter point means that the algorithm avoids being trapped in a local minimum. An annealing schedule is implemented that decreases the temperature as the algorithm runs, which physically corresponds to the algorithm reducing the extent of the search as the number of iterations increases, so as it converges to a minimum. Again, the cost function value corresponds to some physically measurable quantity that is associated with the horn, such as the cross-polar power. Each test point consists of a set of parameter values that define the horn geometry. Based on the result of each horn, appropriate changes are made to the parameters to form a new horn geometry, which is then tested.

Terminology

Some terminology used in SA will now be introduced.

Objective Function

Objective function is the name given to the function whose global minimum is to be found. This is the same as the cost function in the GA. This term will be used from this point on, when the SA algorithm is in use.

Temperature

This is a parameter which has an impact on two parts of the algorithm;

1. The distance of a test point from the current point (size of the search space).
2. The probability of a trial point being accepted that yields a higher objective function value.

The temperature tends to be a scalar initially, however as the algorithm proceeds, it can generally be represented as a vector with a different value for each independent variable. As the algorithm proceeds, the temperature also begins to decrease. This shrinks the search space of the algorithm, and so converges towards a minimum. Reducing the temperature also reduces the probability that a test point which increases the value of the cost function will be accepted. The rate of decrease of the temperature can be controlled manually or handled internally by MATLAB. A slower rate leads to more accurate results, however this takes longer.

Annealing Parameter

This is a proxy for the iteration number of the algorithm. If the algorithm deems it necessary to raise the temperature, it can do so by setting the annealing parameter to a value that is lower than the current iteration value.

Reannealing

This process refers to the temperature being raised after the algorithm has accepted a certain number of new points and has raised the temperature to start the search again, thus avoiding getting caught in a local minimum.

Overview of Algorithm

A description of the overall steps performed by the SA algorithm will now be given.

1. The algorithm begins by generating a random trial point relative to the starting point. It decides how far this point is from the current (starting) point by using a probability distribution whose scale is determined by the temperature for that iteration. The scaling of the probability distribution can be selected in the *AnnealingFcn* option. @*annealingfast* is the default option and the step length equals the current temperature, with the direction of the step being uniformly random. @*annealingboltz* invokes a form of the Boltzmann distribution, where the step length is equal to the square root of the temperature, with the step being in a uniformly random direction again. It is also possible to define a custom annealing function.

2. By using the trial point to evaluate the objective function, the algorithm determines if the trial point is better than the current point (lower value of the objective function). If this is the case, it becomes the new point. If the trial point is worse, the algorithm can still set it as the new point, subject to an acceptance function. The acceptance function is selected using the *AcceptanceFcn* option. The default option is @*acceptancesa*. The acceptance probability is given by

$$\frac{1}{1+\exp\left(\frac{\Delta}{\max(T)}\right)} \tag{4.1}$$

   where Δ = new objective-old objective, T_0 is the initial temperature of variable i and T is the current temperature. We have that Δ and $T \geq 0$, so the probability of acceptance (p) lies in the range $0 \leq p \geq \frac{1}{2}$. The smaller the temperature, or the larger Δ is, the smaller the acceptance probability. It is also possible to define a custom acceptance function. If a point is accepted, the algorithm proceeds. If not, it generates a new point and starts at step one again.

3. As the algorithm progresses, it gradually lowers the temperature, storing the best point found so far. The rate at which the temperature changes is defined by various temperature functions which control how quickly the temperature decays. This can be changed in the *TemperatureFcn* option. The default setting is @*temperatureexp*, which gives $T = T_0 \times 0.95^k$, if k is the annealing parameter, which is the same as the iteration number until re-annealing commences. Other options are

   - @*temperaturefast*: $T = \frac{T_0}{k}$
   - @*temperatureboltz*: $T = \frac{T_0}{\log(k)}$
   - A custom temperature function

4. After examining a certain number of points, the algorithm begins to re-anneal. This process sets the annealing parameters to values lower than that of the iteration number. This raises the temperature for each variable. The annealing parameters are based on

the values of the estimated gradients of the objective function in each variable dimension. The formula for evaluating this is given by

$$k_i = \log\left(\frac{T_0}{T_i}\frac{\max_j(s_j)}{s_i}\right) \tag{4.2}$$

where k_i , $T_{0,i}$, T_i and s_i are respectively the annealing parameter, initial temperature, current temperature and gradient of the objective in the direction i times the difference of bounds in the direction of i, for component i.

5. The algorithm continues until one of the stopping conditions is met, usually set to be the default whereby it stops when the change in objective function is small relative to the set tolerance, which is set to 1e-6 by default. Other possible stopping conditions are

   - **MaxIter:** Stops when the number of iterations of the algorithm exceeds a maximum number of iterations, set to infinity by default.
   - **MaxFunEval:** Stops when the number of objective function evaluations exceeds some number, set to 3000 times the number of variables, by default.
   - **TimeLimit:** Stops when the algorithm has been running for longer than some maximum time. The default value is infinity.
   - **ObjectiveLimit:** Stops when the best objective function value is less than or equal to a set target. Set to negative infinity by default.

MATLAB Implementation

To run the SA algorithm from the command line, utilising the default options, the following syntax is used

```
[x, fval]=simulannealbnd(@objfun,x0)
```

The first input argument is a function handle for the objective function. The second input argument is the vector of initial guesses for all of the independent variables involved in the calculation. The first output argument, x, is a vector containing the values of the final point tested, i.e. the point that gave the best value for the objective function. The second output argument is the value of the objective function at this point. Additional output arguments can be included, which for example, can display the stopping condition that resulted in the termination of the algorithm or the performance of the algorithm after every iteration.

The most common additional input argument is a variable which points the SA function towards a customised options block for that algorithm. Here, the parameters listed above can be changed, as can the stopping conditions and the functions used to determine the rate of cooling and the acceptance probability. This can all be important, as the algorithm is

stochastic and so can produce slightly different results each time. Information about the cause of the stoppages and the algorithm performance can therefore be used to examine the algorithm to confirm correct operation.

4.2 Horn Types and Optimisation Metrics

In this section, a brief description of the horn types and optimisation metrics that were used in the optimisation trials will be given.

4.2.1 Horn Types

For the optimisation trials, two horn designs were selected as the 'skeleton' around which the design process was to be based. These horns can be defined with respect to various parameters, and altering the values of these parameters result in the precise profile of the horn changing. In this way, the performance of the horn can be readily controlled, and a particular set of values for the parameters will result in optimal performance with respect to some horn performance characteristic.

The two horns selected were a version of the standard spline profile horn [53, 54] and a mixture of a Gaussian profiled horn antenna (GPHA) and a sine profile horn [55]. The studies cited have shown that these horn designs give excellent flexibility in terms of performance (specifically also in the areas important for future Far-IR and CMB missions, as will be outlined later) and can achieve this using relatively few degrees of freedom, something which has obvious importance for optimisation algorithms. Other horns were studied when assessing which designs to apply the optimisation process to, for example Picket-Potter and Potter horns, but the degree of control over the performance from an optimisation algorithm point of view was questionable. Also, critically, since it was desired to come up with a geometry that was easy to manufacture, the profiled horns are significantly easier to machine as there are no large jumps in radius, so the profiled horns were selected.

Piecewise conical profile horn

As indicated in the references above, a spline profile horn is one which is defined by i points, and in turn i linear sections when these points are joined. Each point is separated from the last by some defined distance, which gives rise to the length of that section. An example of a spline profile horn consisting of six points is shown in figure 4.1. Six points were selected, as this provides a good compromise between the number of degrees of freedom in the system and the computational load that would be experienced during the optimisation calculation. Since i is typically large for a spline profiled horn ($i >> 6$), this design was named a piecewise conical profile horn, as it is constructed of several pieces 'stitched together' in the same vein

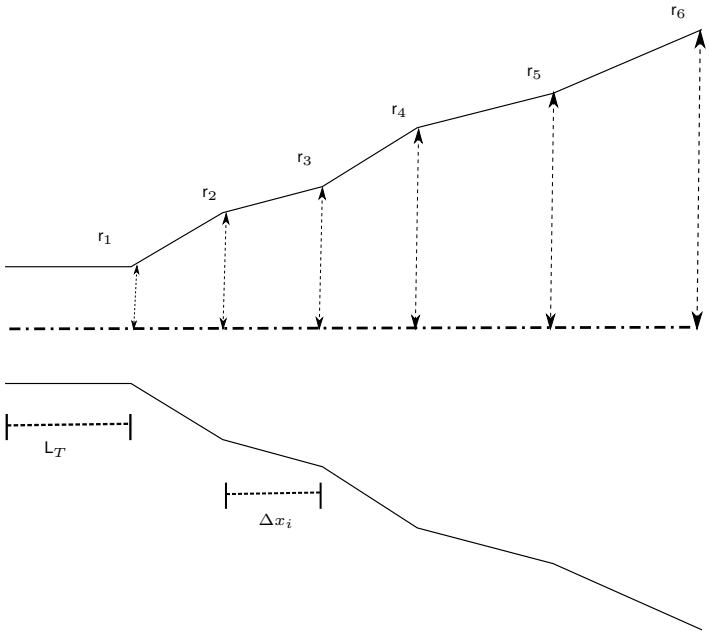

Figure 4.1: Six point spline profiled horn

as a piecewise mathematical function. In order to reduce the computational load further, the throat radius is kept at a fixed value, defined such that it acts as a filter to only allow the desired modes to be supported throughout the horn. This parameter, the length of the throat, is known as L_T. In the examples considered here, the throat length will be fixed at 5 mm, which is sufficiently long to remove evanescent modes. The separation between any two particular points is given by Δx_i. In general, this will have different values for each section and so choosing enough sections and making each section sufficiently small, will allow any profile to be obtained. In order to further reduce the computational load, Δx_i will be a variable controlled by the algorithm, however it will be the same for each point, so each section will be of uniform but variable length. In doing this, the length of the horn (minus the throat section) can be readily controlled by defining an upper and lower bound on the desired length (determined by the physical limitations on the dimensions of the desired horn). The algorithm can then vary this parameter and calculate the resulting different value for Δx_i in each iteration. In a given optimisation run, $\Delta x = \Delta x_i, \forall i$.

The radius of each section after the throat, r_i, is unique to a given section. Each of these radii is left as an independent optimisation variable. Doing so allows the profile of the horn to be adjusted so as it can assume the optimal configuration in order to achieve the best performance in terms of the optimisation metric. Changing the radii independently means that the flare angle for each section can be controlled. This changing flare angle can be used to excite various modes in certain relative proportions. The range allowed for the radius of each section is such that the slope of each section does not have to always be positive. i.e. the horn does not always have to open out. It is permitted for the radius of a section to be less than the radius of the

previous section, and so the horn can have a negative slope for a given section. This allows the optimiser to explore a wider range of parameter space, as such configurations of the horn could be what is necessary to control the modal content in such a way as to optimise the horn with respect to the metric being used in the cost function. This flexibility combined with the ability to control the overall length of the horn means that the phasing/modal content/relative amplitudes of the modes can be used to achieve the desired performance. Although control of the lengths of the individual sections would potentially yield a better result, the implemented method should allow a good compromise between overall performance and execution time.

Using the throat length, radius, and the other variables, the complete horn profile can be constructed for analysis in SCATTER. For these examples, again, six points (sections) were used, including the throat. Since the throat length is fixed, this means that the horn profile is a function of seven variables, which is certainly reasonable from the point of view of computational load. With the throat taken to be one section of length 5 mm, for the purposes of constructing the geometry in SCATTER the radius of the throat is called $y1$. The section between $y1$ and $y2$ of length Δx must now be considered. The flared section of the horn is of length $5\Delta x$, and this portion of the horn must consist of sufficient sections so that there exists a minimum of 10 sections per wavelength, as a requirement of the modal matching technique as outlined in chapter 2. The number of sections is therefore calculated using equation 4.3, where $L = \zeta\Delta x$, with ζ being the number of piecewise sections in the conical part of the horn (5 in this case) and κ being the desired number of horn sections per wavelength, in this 10.

$$s = \text{ceil}\left(\frac{\kappa L}{\lambda}\right) \tag{4.3}$$

The result obtained from this equation is rounded up by Matlab to the next integer, and this result is increased by one until it is divisible by κ so as to have the number of horn sections evenly distributed over each piecewise section. Using $\frac{s}{\kappa}$ sections, the region between $y1$ and $y2$ (i.e. the first piecewise section) is linearly interpolated so as the radius of the horn monotonically changes from $y1$ to $y2$ over this region, of length Δx. The length of each horn section is saved into an array, as is the corresponding radius of each section. This is carried out for each piecewise section of the horn until the horn is defined, from entrance to exit aperture, as a set of lengths and radii. The arrays containing these figures are then passed to the mode matching portion of the code where they are processed using the dynamic version of SCATTER along with SVD based field calculations to increase the speed of the calculation, as described in 3.

Gaussian profile horn antenna

In this section, Gaussian profile horns are examined. The purpose of these feed horns is to use the desired mode set (in this case the Gaussian mode set) to achieve a horn profile based on

the shape of this mode set, that essentially moulds the beam into the correct shape.

As cited in [55], solutions of the Helmholtz equation can be found which yield a field distribution that correlates strongly with certain waveguide modes. For example, the HE_{11} waveguide mode strongly resembles the fundamental Gaussian mode, which is denoted by Ψ_0^0. The issue encountered in horn design is that of how to efficiently excite these Gaussian like waveguide modes. This is where the GPHA is used.

A GPHA has a profile that is tapered in a similar way to the Gaussian beam modes, so for the fundamental Gaussian, this results in the radius of the GPHA being tapered according to the equation

$$r(z) = r_0\sqrt{1 + \left(\frac{\lambda z}{\pi w_0^2}\right)^2}, \tag{4.4}$$

where r_0 is the input radius, λ is the design wavelength of the horn, and w_0 is the waist radius of the fundamental Gaussian beam that the horn is being designed to generate. This is based on the waist being placed at the input to this tapered section, and so the output will be Gaussian, but with improved Gaussicity, sidelobe and cross-polarisation levels relative to the input. These qualities make such a horn a good candidate for missions that require the performance qualities sought by this investigation, as will be outlined later.

This technique requires that a good approximation to the desired aperture mode is placed at the input of the horn, with the horn itself merely refining it. Numerous methods exist for generating such mode mixtures, for example

1. Changing the corrugation depth of waveguides from $\frac{\lambda}{2}$ to $\frac{\lambda}{4}$ or from smooth walled to $\frac{\lambda}{4}$, depending on whether or not the HE_{11} mode is being generated from the TE or TM mode sets respectively. In this case, the input and output radii are equal.

2. A step in radius between adjacent waveguide sections, as described for Potter horns.

3. The use of other conical horn antennas to control the coupling between the various modes, acting as an input to the GPHA

Since to reduce design time and manufacturing cost/complexity it is desired to not use corrugations or to design an additional horn to excite the GPHA, the mode will be generated by inserting a step between the throat of the horn and the flared section, the depth of which can be controlled by the optimiser.

The overall formula for the profile of the GPHA is given by equation 4.5

$$r\left(z\right) = r_{in} + \left(r_{out} - r_{in}\right)\left\{\left(1 - A\right)\frac{z}{L} + A\sin^2\left(\frac{\pi z}{2L}\right)\right\}, \tag{4.5}$$

where L is the length of the horn, and r_{in} and r_{out} are the input and output radii respectively. A is a parameter that determines the severity of the profile added to the linear taper, and

in general, a value of between 0.7 and 0.9 produces good overall results, [46], however values outside of this range are encountered if one parameter wishes to be maximised at the expense of others, which is a situation that the optimisation process naturally lends itself to. It was shown in [55], that given a reasonably pure input to the horn, sidelobe and cross-polarisation levels could be readily controlled by varying the horn parameters, showing a large improvement over the initial input to the horn. The limitations of this horn profile, subject to the optimisation algorithms and performance metrics in use, will be investigated. The algorithm will be allowed to vary the depth of an initial step from the throat into the flared section in order to excite the mode that will be shaped by the flared section.

4.2.2 Optimisation Metrics

The three optimisation metrics that were tested will now be presented. These metrics were chosen with the requirements for a future Far-IR or CMB mission in mind, as will be outlined in the following chapter. The chosen metrics all either directly influence the characteristics important to meet the performance requirements, or indirectly influence them.

Co-polar beam symmetry

It is desired that the co-polar beam be largely circularly symmetric with respect to azimuthal angle. This means that the performance of the antenna will not be heavily dependant on azimuthal orientation which is important if the measurement system is to be robust and mobile. This metric can be quantified by carrying out an overlap integral between the 0 and 90 degree cuts of the complex co-polar field, as in equation 4.6.

$$c_{eh} = \frac{\int_0^b E_E E_H^* 2\pi x \mathrm{d}x}{\sqrt{\int_0^b E_E E_E^* 2\pi x \mathrm{d}x}\sqrt{\int_0^b E_H E_H^* 2\pi x \mathrm{d}x}}, \tag{4.6}$$

where E_E and E_H represent the 0° and 90° cuts of the complex co-polar beam respectively, each containing the peak intensity value, and b is the off-axis distance corresponding to the maximum angle considered in the farfield representation of the field. The farfield is considered, as this is where the horn will couple to the incident radiation. In this form, a coupling coefficient of one would indicate perfect coupling between both orthogonal planes in terms of phase and amplitude and so would indicate a perfectly circularly beam, thus the higher the value for the coefficient, the more symmetric the beam. This measure of symmetry, defined in this manner, can assume a value between zero and unity.

Return loss

In order to maximise the throughput of the horn, i.e. to maximise the power emerging from it into the instrument (a very important quality when low level signals such as those associated

with the CMB are in question), it is necessary to minimise the power being reflected within the horn. If the horn is assumed to be illuminated at port 1 in SCATTER, then the power contained within the S_{11} matrix indicates the total reflected power, i.e. the power that enters through port 1 that arrives back at port 1. Since power is conserved in the mode matching technique, any power that enters through port 1 that does not arrive back at port 1 must exit through port 2. Thus, if the power contained within the S_{11} matrix (the integrated return loss) is minimised, then the horn throughput is maximised. Since this metric only considers a scattering matrix, no fields need be calculated and so this will reduce execution time.

As outlined in chapter 2, the columns within a given scattering matrix represent partially coherent modes, thus the total power is found by adding the entries in a given column coherently and adding the results from all the columns in quadrature (incoherently). This figure is then minimised by the optimisation algorithm. This is evaluated according to equation 4.7.

$$\text{ReturnLoss} = \sum_{j=1}^{N} \left(\left| \sum_{i=1}^{N} S_{11,ij} \right|^2 \right), \tag{4.7}$$

assuming that N modes are considered throughout the entire horn and that only one azimuthal order is being considered. For multiple azimuthal orders, this figure is calculated for the S_{11} matrix in each order and then the figures are totalled.

Maximum cross-polar level

If figure 5.22 is examined, for example, the familiar quad-lobe structure that is associated with the cross-polar field is observed. This is the standard distribution of the cross-polar field for a horn antenna. By considering a cut along the $\frac{\pi}{4}$ azimuthal direction, it is clear that this cut will contain the maximum level realised by the cross-polar field. Thus, if the horn is designed so as to minimise the maximum value reached within this cut, then the maximum level contained within the entire cross-polar field will be minimised due to symmetry. The field is therefore calculated, with the maximum value being extracted and passed to the optimiser which has the goal of minimising this value.

4.3 Optimisation using the Genetic Algorithm

In this section, the genetic algorithm (GA) will be used to optimise the performance of the piecewise conical profile horn and the GPHA. As outlined earlier, each horn will be optimised with respect to three horn characteristics,

1. Co-polar beam symmetry
2. Return loss

3. Minimising the peak level in the cross-polar beam pattern

Since the optimisation algorithm will attempt to solely optimise the performance of the horn with respect to the metric in question (at the design frequency of 100 GHz in each case), it is expected that it will do this at the expense of other horn metrics, i.e. there will be a performance trade off. The purpose of optimising with respect to each of the different parameters is to observe which parameter still allows the horn to perform well with respect to other metrics, in addition to the one that was optimised with respect to. It will also be of interest to observe which metric allows good bandwidth performance to be achieved, despite the fact that the horn was designed at a spot frequency.

For each of the optimisation runs illustrated that relate to the piecewise conical profile horn, the variables were limited to the values shown in table 4.2.

Parameter	**Lower bound [mm]**	**Upper bound [mm]**
Overall length (without throat)	30.00	40.40
r1	1.60	2.00
r2	1.61	4.00
r3	3.40	6.00
r4	5.50	8.00
r5	7.50	10.00
r6	9.50	12.10

Table 4.2: Lower and upper bounds for the horn design parameters used in each of the piecewise conical horn optimisation runs

The values presented in table 4.3 were used to define the range of values that each of the parameters pertaining to the GPHA horn could assume. These values were chosen so as the overall defining properties of the horn (such as throat and aperture radii and overall length of the flared section) were limited to the same maximum values, allowing a valid comparison to be made between the performance of the two resulting designs, subject to the same physical constraints. The values given in both tables were selected by placing a reasonable constraint

Parameter	**Lower bound**	**Upper bound**
Aperture radius	11.70	12.10
Length	19.80	44.00
A	0.50	0.95

Table 4.3: Lower and upper bounds for the horn design parameters used in each of the GPHA optimisation runs. Lengths are quoted in mm.

on the physical size of the horns (for example, length $\approx 14\lambda$ and aperture radius $\approx 4\lambda$) and by considering the fact that ultimately single-moded horn designs are required for the application

of a horn designed by this performance, as will be outlined in the next chapter. This requirement has an impact on the throat radius. The values for the intermediate radii in the case of the piecewise conical profile horn were chosen so as to keep the horn profile physically reasonable and to therefore ensure relative ease of manufacture. This was best achieved by having the upper and lower bounds of consecutive sections monotonically increasing and by placing the lower bound of the radius of the i^{th} section midway in the range of values permitted for the radius of section $i-1$. Allowing the radii of each section to vary freely between the aperture and throat radii resulted in highly unphysical designs that performed poorly. Limiting the search space to more physical designs resulted in good horn designs, the results of which will be shown below. Similarly for the GPHA, the values allowed for the various parameters were determined from the literature, as per the sources cited in section 4.2.1, with the throat length again fixed at 5 mm and with the throat of the horn fixed to be of radius 1.2 mm in order to reduce the size of the search space.

4.3.1 Co-Polar Beam Symmetry

In this section, the optimisation of the two horn designs with respect to maximising the symmetry of the co-polar beam will be investigated. In either case, the parameters defining the geometry of the horn under test during a given iteration of the algorithm were used to generate a horn profile which was then analysed using the mode matching code. The 0 and 90° cuts of the farfield co-polar field were then calculated and the symmetry of the overall beam evaluated as outlined in section 4.2.2. As stated previously, the value returned for this metric will be in the range $[0, 1]$. Since it is desired to maximise this value, but the genetic algorithm seeks a global minimum, it is necessary to return the negative of the obtained value to the optimiser. In this way the optimiser will adjust the horn parameters in order to achieve a value for the metric that tends towards -1, and in doing so will physically design a horn that has as close to unity coupling as is possible, as minimising this negative value maximises the absolute value, and it is this figure that has physical significance for the horn design.

Using the values quoted in table 4.2 as the starting point, a piecewise conical profile horn was optimised with respect to the symmetry of the co-polar beam. The resulting horn geometry is shown in figure 4.2, with the optimised values for the horn profile parameters shown in column two of table 4.4. The optimisation process took approximately 10-11 hours. The throat length was fixed at 5 mm and the optimisations all carried out at a spot frequency of 100 GHz. This timescale is extremely acceptable in terms of one of the goals of this investigation; namely to carry out the design process using standard desktop computing power on the order of hours.

From figure 4.2, it is clear that there are no sudden jumps in radius and that the radius is monotonically increasing. This means that the horn can be readily manufactured. Figure 4.3 shows the 0 and 90° cuts of the co-polar plane and the 45° cut of the cross-polar field for this horn at the middle and ends of the band, along with the co-polar beam symmetry as a

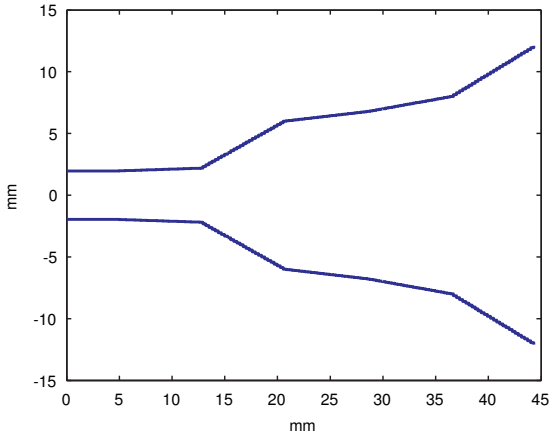

Figure 4.2: Piecewise conical profile horn optimised for maximum co-polar beam symmetry

Parameter	**Co-polar symmetry**	**X-polar level**	**Repeat of X-polar run**	**Return loss**
Δx	39.60	39.71	39.75	39.62
r1	1.96	1.66	1.69	1.62
r2	2.19	3.63	3.60	1.66
r3	6.00	3.40	3.44	3.40
r4	6.78	6.88	6.91	6.77
r5	7.99	7.50	7.46	9.30
r6	12.10	11.59	11.57	9.73

Table 4.4: Optimised values for the profile parameters of the piecewise conical profile horn, optimised using a genetic algorithm with respect to the indicated horn performance characteristics. Quoted figures are in units of mm

function of frequency.

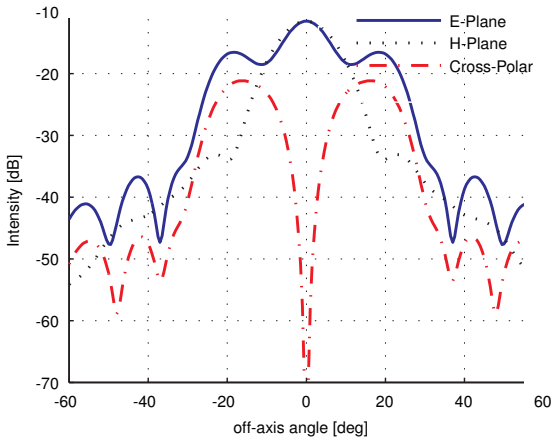

(a) 85 GHz

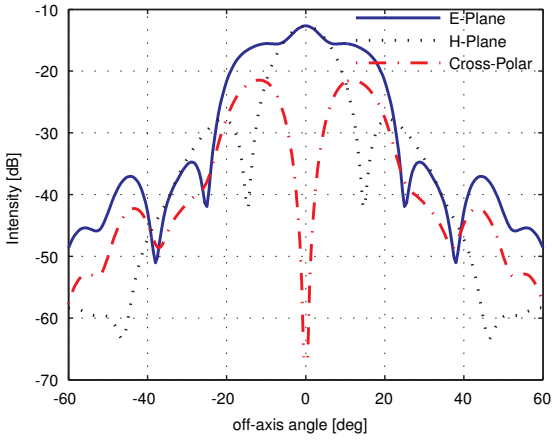

(b) 100 GHz

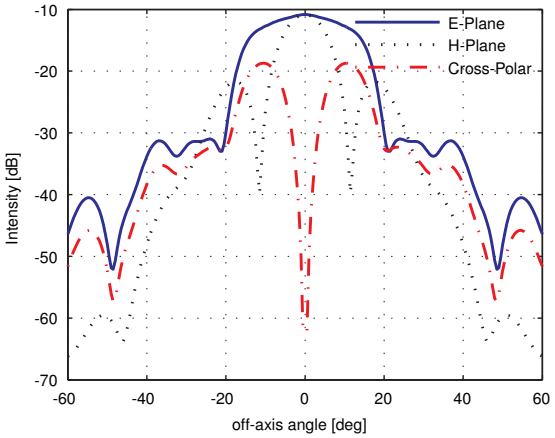

(c) 115 GHz

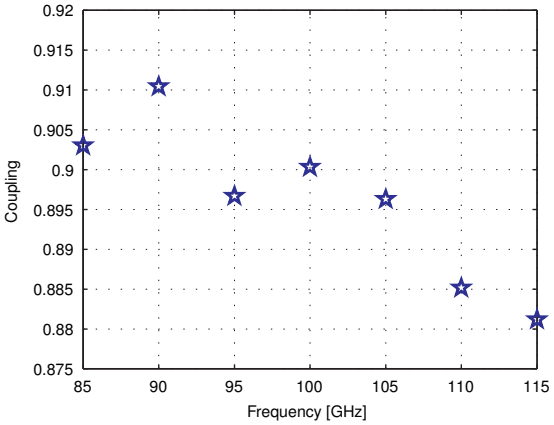

(d) Co-polar symmetry

Figure 4.3: Cuts and co-polar beam symmetry for the piecewise conical profile horn optimised with respect to co-polar beam symmetry

Inspecting the beam patterns show high sidelobe levels and significant shouldering effects at all frequencies. In terms of field distribution, the co-polar cuts are not hugely symmetric, meaning this horn would not be ideal for applications where the azimuthal orientation of the horn can vary, as the changing nature of this co-polar distribution means that the horn is sensitive to azimuthal variation. Although the co-polar levels are usably high, the cross-polar levels are highly elevated too. This would make this horn design less than ideal for polarisation sensitive measurements, such as those required when measuring the polarised CMB anisotropy. The difference in power levels between the two polarisations is clearly shown in figure 4.4a, which shows the ratio of the maximum level in the cross-polar field to the maximum level in the co-polar field. This figure is above -10.5 dB at its lowest and peaks at just under -7.5 dB, remaining between these values across the band. For a polarisation sensitive measurement, this is not ideal. Examining the return loss, shown in figure 4.4b, the horn is seen to perform very

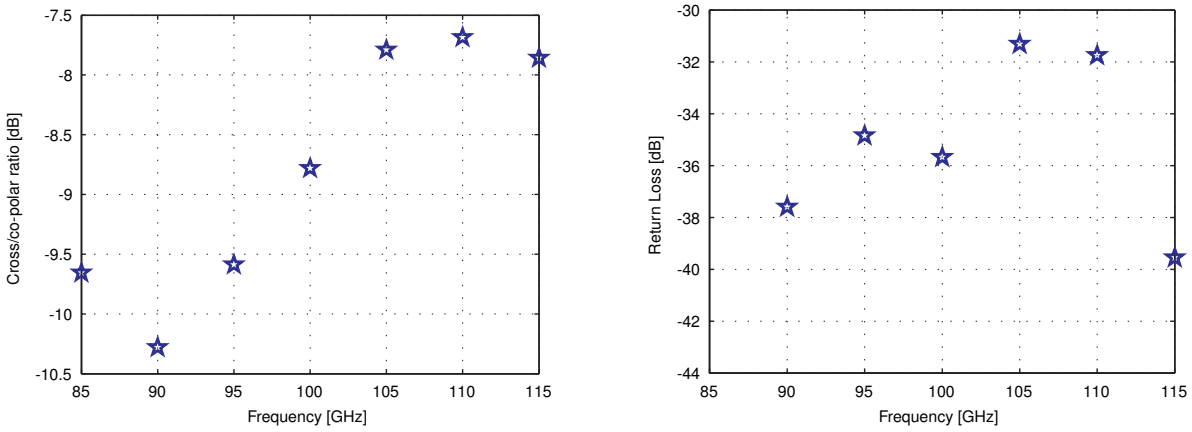

(a) Cross/co-polar ratio

(b) Return loss

Figure 4.4: Cross/co-polar ratio and return loss for the piecewise conical profile horn optimised with respect to co-polar beam symmetry

well in this regard, achieving a return loss of below -30 dB across the band. Returning to figure 4.3d, we see the coupling between the principal planes of the co-polar field (the symmetry of the co-polar field). Despite the fact that this was the metric optimised with respect to, the value is lower than might be expected. The field patterns of the cuts illustrate the source of this low level of symmetry. The aperture radius has achieved the largest permissible value. The larger the aperture radius, the lower the mismatch between the aperture and free space. This implies that the algorithm was attempting to minimise this impedance mismatch by making the aperture radius as large as possible, however it was limited by the physical requirements of the horn size. A low mismatch would reduce cross-polar levels and this would result in higher co-polar symmetry, so maximising the aperture radius would indicate that this was the strategy adopted by the optimiser, implying that the other degrees of freedom within the horn geometry were insufficient to increase the symmetry. As a result of the optimiser not being able to minimise the impedance mismatch, cross-polar levels remain high as observed and the coupling remains low. Of note also is that despite the fact that the horn was being optimised with respect to symmetry at 100 GHz, the highest levels of symmetry across the band were not realised at 100 GHz. This is explained by considering the fact that no other frequencies were considered in the optimisation process. Given the degrees of freedom available to the optimiser, and within the limits of the process identified by the optimiser to increase the symmetry, the value realised at 100 GHz was the highest possible value. It is entirely plausible that the horn geometry which gives rise to this may perform even better at other frequencies, and this would seem to be the case. The optimiser makes no allowance for other frequencies, so the fact that the symmetry is not at its maximum at 100 GHz is not an anomaly.

The above results show that although the horn profile is physically simple and therefore less costly to manufacture, using this type of horn as the skeleton for the optimisation process and optimising with respect to beam symmetry results in unacceptable performance where

symmetric co-polar beams with low sidelobe levels are required, along with low cross-polar levels. The corresponding GPHA design will now be examined.

Prior to examining the results of the optimisation process for the GPHA horns, it is necessary to briefly examine some issues that occurred whilst applying the mode matching technique to this horn profile. The optimisation algorithm initially failed numerous times, due to some of the matrices encountered during the mode matching stage being close to singular (a matrix A is singular if $\det(A) = 0$, so in this case $\det A \approx 0$) for certain values of the horn parameters, i.e. certain horn geometries. Since it is necessary to carry out several matrix inversions during the calculation and cascading of scattering matrices, this resulted in the inverses of these matrices potentially being inaccurate and containing significant numerical errors, and so the overall results contained large errors, or in the worst case scenario it caused the optimiser to crash. This issue was addressed by calculating the Moore-Penrose (pseudoinverse) matrix [56] as opposed to the direct inverse. This was found to remove the mathematical instability in the calculation and when applied to systems for which the scattering matrices were well known, returned the correct solution.

A further issue was in the conditioning of such matrices. MATLAB assigns a conditioning number, *rcond*, to a matrix. This is known as the reciprocal condition number of the matrix and is normalised to unity such that if *rcond* is near 1.0 the matrix is said to be well conditioned, and if it is near 0.0 then the matrix is said to be poorly conditioned. A matrix A, is said to be poorly conditioned if when applied to the linear equation $Ax = b$, a small change in b gives rise to a large change in x. Thus, when calculating the inverse of a matrix, which is done by considering linear equations, if *rcond* has a very low value, then the resulting inverse of the matrix can be inaccurate. If it is NaN, then the inverse in undefined, causing the code to crash or return null matrices. In the context of the optimiser, when the NaN condition is encountered, it can become stuck in the section of parameter space corresponding to this, and the cost function will never return a value, resulting in the programme crashing or hanging. To help to eliminate this, and to avoid inaccurate solutions corresponding to the low value condition of *rcond*, a statement was coded into the cost function that returns a large value for the cost function when *rcond* is either small or NaN. In this case, the optimiser receives a value from the cost function, so it does not crash, however it is so large compared to the values obtained when *rcond* is a sensible number, that the optimiser is guided away from parameter values that yield poor values for *rcond*. This does not forbid these situations occurring, however it reduces the likelihood significantly, and puts in place a mechanism to deal with it.

It was also necessary to include similar conditions to steer the optimiser away from the portions of parameter space that yield certain combinations of parameters which are not physically useful in terms of the resulting horns, despite the individual parameters being restricted to physical values. This was not an issue in the piecewise conical horn, as the parameters correspond directly to physical quantities. An example of this issue occurred when one run of the

optimiser to minimise the cross-polar power (as will be seen later) resulted in a maximum level of -126.8 dB. This is clearly an excellent result, however the resulting fields showed that this occurred due to no power propagating through the horn. The optimiser had simply reduced the throat radius to below 0.88 mm, a value at which, at 100 GHz, no modes can propagate, so no power emerges at the aperture. To eliminate this condition, and to ensure that an appreciable amount of power entered the horn, the throat radius was restricted to values of greater than 0.88 mm, and a condition inserted to ensure that horns yielding low levels of power were rejected.

During initial runs of the optimiser it was also noted that the total length occasionally meant that the number of sections per wavelength was approaching ten, the minimum required for accuracy in SCATTER. To avoid this, the number of sections for each portion of the horn became a variable, determined by the total length of each of these sections. On each iteration, the total length of each of the sections, L, was found, and the required number of sections, s, to guarantee 15 sections per wavelength and therefore an accurate calculation was found. With these issues addressed, it was possible to carry out accurate optimisations.

The results obtained using the GPHA horn as the basis for the optimisation process are shown in table 4.3. The throat length was again fixed at 5 mm. This horn profile was found to be computationally intensive to optimise, and so a further optimisation parameter was fixed. The throat radius was fixed at 1.2 mm, a sensible value as it allows the TE_{11} mode to propagate. The horn profile obtained when optimising with respect to the co-polar beam

Parameter	**Co-polar symmetry**	**X-polar level**	**Return loss**
Aperture radius	11.79	11.82	11.76
Length flared section	19.80	19.82	31.78
A	0.95	0.94	0.59

Table 4.5: Optimised values for the profile parameters of the GPHA horn, optimised using a genetic algorithm with respect to several horn performance characteristics. Lengths/radii are in units of mm

symmetry is shown in figure 4.5. What is immediately clear is that the radius of the horn decreases as the exit aperture is approached. This is a result of the high value for A. The 45° cut of the cross-polar field, along with the principal planes of the co-polar field, are shown in figure 4.6 for the centre and ends of the band, along with the co-polar beam symmetry across the band. Similar to the issues encountered with the piecewise conical profile design that was optimised according to the same metric, it is clear that the co-polar field patterns are not highly symmetric, an assertion which is confirmed by examining the symmetry across the band, which displays a similar range of values to the piecewise conical horn, although with a lower upper limit. Again, the sidelobe levels and shouldering effects are at a high level relative to the on-axis field, with the co-polar beams also exhibiting quite significant off-axis

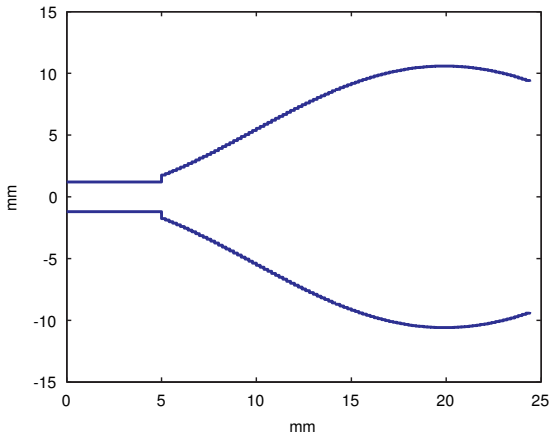

Figure 4.5: GPHA horn optimised for maximum co-polar beam symmetry

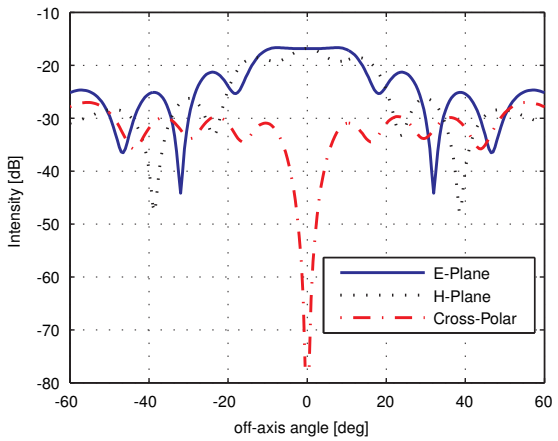

(a) 85 GHz

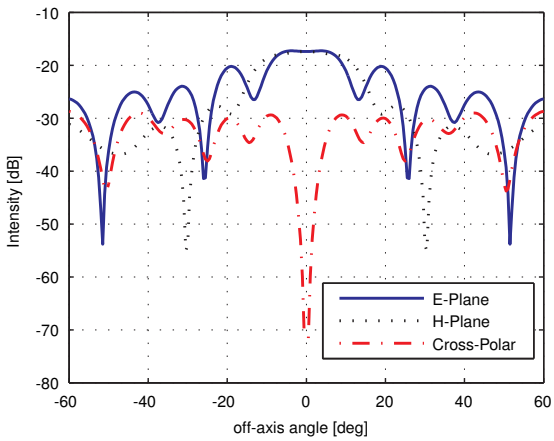

(b) 100 GHz

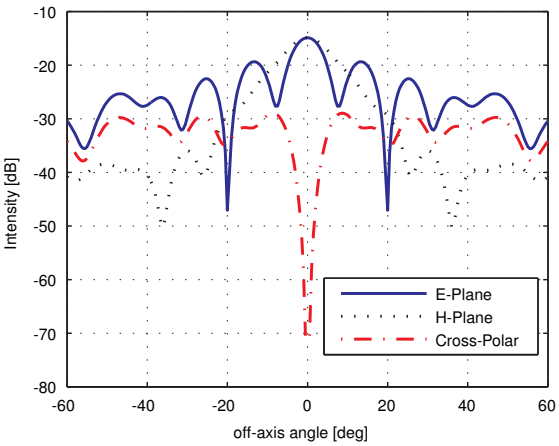

(c) 115 GHz

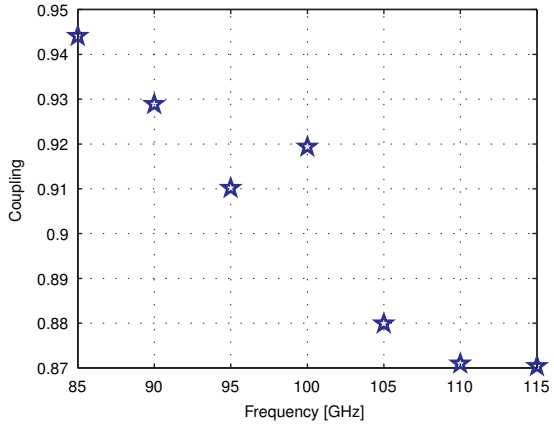

(d) Co-polar symmetry

Figure 4.6: Cuts and co-polar beam symmetry for the GPHA horn optimised with respect to co-polar beam symmetry

contributions from off-axis modes at the lower end of the band. These beam patterns would not be usable for applications that required low level sidelobes with a symmetric beam. The cross-polar fields are at a high level relative to the co-polar field, and so polarisation sensitive measurements could not realistically be carried out using this horn, although the performance in this respect is improved relative to the piecewise conical horn, as is clear from figure 4.7a, where the cross/co-polar ratio is in general lower for the GPHA. Despite this, the return loss of the horn is still low, albeit higher than in the previous example. The one exception to the general low level of return loss occurs at 100 GHz, as shown in figure 4.7. It can be

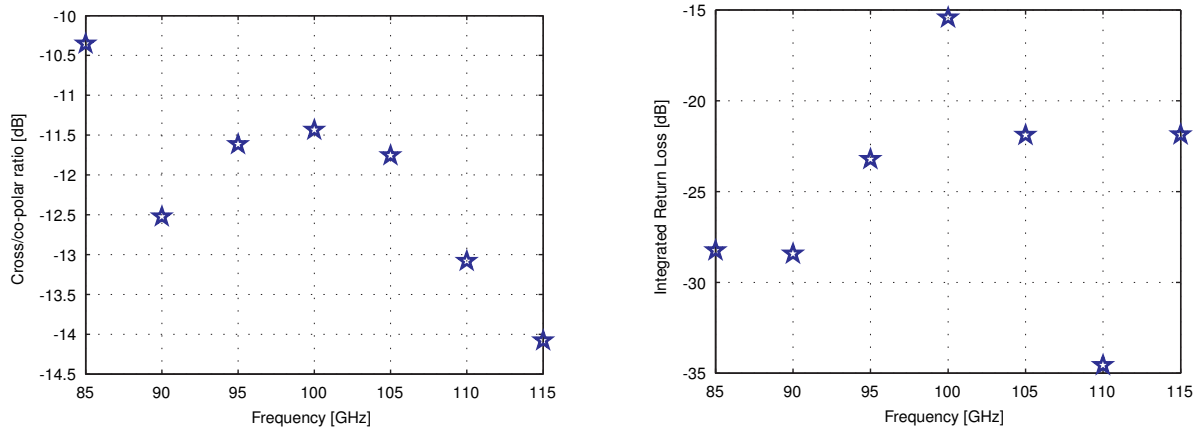

(a) Cross/co-polar ratio

(b) Return loss

Figure 4.7: Cross/co-polar ratio and return loss for the GPHA horn optimised with respect to co-polar beam symmetry

concluded from this, that using either of these horn types as design bases for the optimisation process does not yield a usable horn design when optimised with respect to the symmetry of the co-polar beam. The limitations in performance that result from the two designs are similar, implying that the optimised configuration for the horn profiles for this metric (to as high a coupling value as the optimiser can manage within the limitations of the amount of control it has on the horn profile and the process that it identifies as controlling the symmetry) can be achieved without improving the performance of the other metrics. If anything, it results in poor performance in other areas, as they are sacrificed in order to improve the performance of the metric being optimised. High performance for the other metrics does not necessarily follow from a symmetric co-polar beam, although a symmetric co-polar beam may follow naturally from some other metric being optimised. It should be noted that the values obtained for the optimisation variables for the GPHA were, for A and the length of the horn, close to the limits imposed on the search space for these parameters. This implies that the optimiser encountered difficulty in obtaining a solution in this space, although this space was selected to physically limit the horn to the same degree as the piecewise horn.

It was attempted to optimise the designs above with respect to all three metrics simultane-

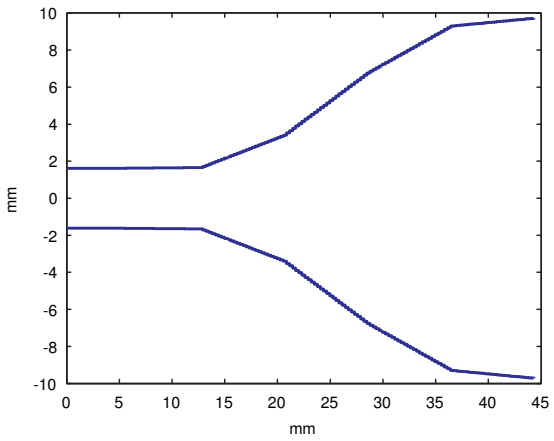

Figure 4.8: Piecewise conical profile horn optimised for minimum return loss

ously, however the process violated the requirement for a short execution time, taking several days, and the results showed no significant improvement beyond those presented above. This combined with the fact that the user has no control over the physical process that the optimiser identifies within the horn that controls the metric of interest, it was necessary to proceed and investigate the other quoted metrics individually.

4.3.2 Return Loss

The optimisation of both horn designs with respect to minimising the return loss will now be examined. The same bounding values were placed on the horn geometry parameters as outlined in tables 4.2 and 4.3. This metric was evaluated within the cost function as outlined in section 4.2.2, by calculating the overall S_{11} matrix of the horn under test and establishing the total power contained within it. With the power expressed linearly, the optimisation algorithm was allowed to directly minimise this figure.

Using the piecewise conical profile horn in the first instance, the optimisation algorithm returned the horn geometry parameters shown in column 5 of table 4.4. A cross-section of the resulting horn antenna is shown in figure 4.8. In common with the previously presented piecewise conical profile horns, this geometry is again lacking in large radius steps or complicated design features, making an excellent candidate in terms of ease of manufacture. Figure 4.9 shows again the principal cuts of the co and cross-polar fields at the ends and centre of the band. Compared to the cuts presented for either of the horn geometries that were designed with respect to maximising the co-polar symmetry, this design gives superior performance in terms of beam parameters. Although there are some relatively high level shouldering effects present in the 85 and 100 GHz cuts, the sidelobes have assumed a more regular shape and distribution. The field also remains relatively constant from 85 to 100 GHz. At the upper end of the band there is some noticeable degradation in beam quality, with the shoulders becoming sidelobes, however this beam is still of a significantly higher quality than in the case of the other

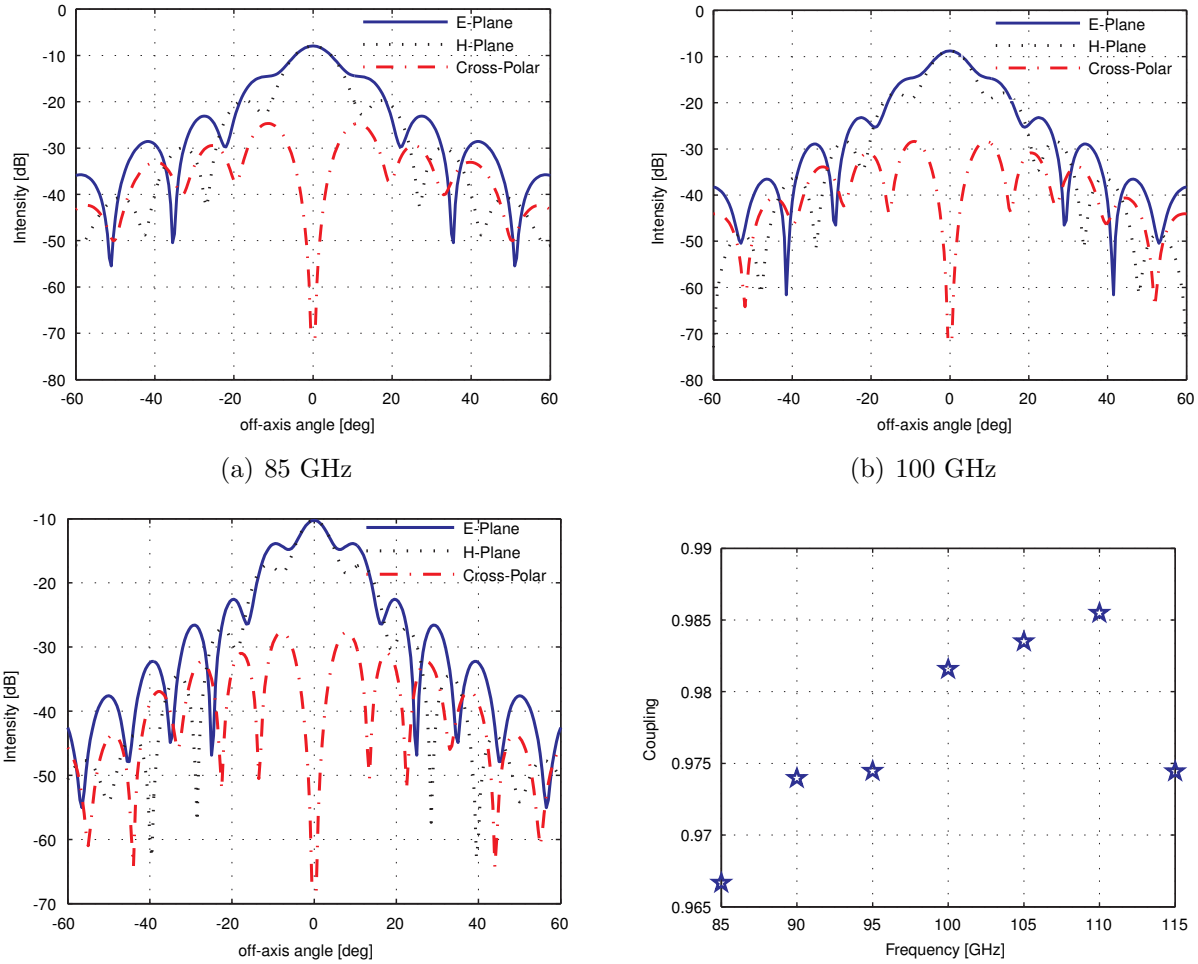

(a) 85 GHz

(b) 100 GHz

(c) 115 GHz

(d) Co-polar symmetry

Figure 4.9: Cuts and co-polar beam symmetry for the piecewise conical profile horn optimised with respect to return loss

horns thus far at this frequency. In general, the beams are a lot more consistent across the band, with less higher level sidelobes. Of note is that despite the co-polar symmetry not being optimised in this case, minimising the return loss has resulted in significantly higher beam symmetry across the band than in either of the cases where this was the metric optimised. This is clear from examining the co-polar cuts and is also quantified in figure 4.9d.

The cross-polar field is also at a much lower level compared to the other horns, both on an absolute scale and relative to the co-polar beam. The separation between the co and cross-

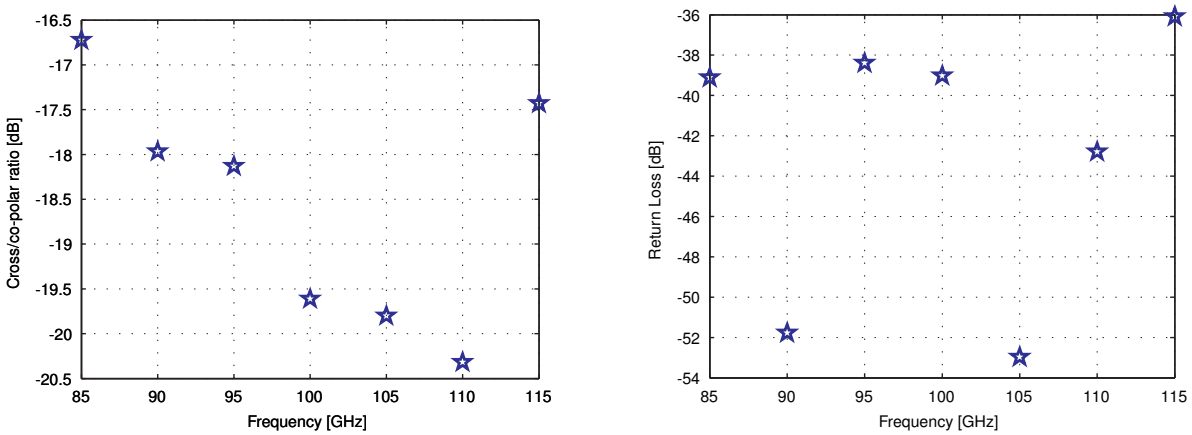

(a) Cross/co-polar ratio

(b) Return loss

Figure 4.10: Cross/co-polar ratio and return loss for the piecewise conical profile horn optimised with respect to return loss

polarisations is shown in figure 4.10a, with the lower levels than previously observed indicating that there is a significant separation between both polarisations. The return loss across the band is shown in figure 4.10b. As expected, since this was the parameter that was optimised, the return loss of the horn is extremely low. Of interest, is that despite being minimised at only 100 GHz, it remains below -38 dB across the band. All of the analysed metrics perform well across the band, considering the spot frequency nature of the design process. This is surprising, as it might be expected that focussing on one frequency would enhance performance at this frequency at the expense of others, as the horn geometry would be tailored to interact with a particular wavelength.

The same optimisation process was carried out using the GPHA horn as a basis for the optimisation process. The optimised values for the geometry parameters are given in column 4 of table 4.5. The resulting horn is shown in figure 4.11. The principal cuts of the fields are shown in figure 4.12, at the ends and centre of the band. The principal planes again show significant sidelobes and shouldering, more so than in the piecewise horn. The design is therefore unusable for applications that place emphasis on sidelobe levels. Comparing the two principal cuts shows decreasing symmetry as the frequency increases, and this is confirmed quantitatively in figure 4.12d, where the symmetry is shown as a function of frequency. The co-polar symmetry realised

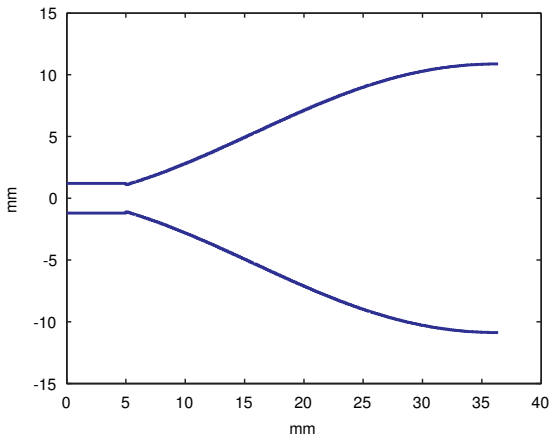

Figure 4.11: GPHA horn optimised for minimum return loss

by this horn is less than its piecewise counterpart, with this symmetry steadily dropping off to significantly lower levels as the frequency increases. Thus, the horn is not suited for applications where azimuthal orientation is variable and critical. The cross-polar field gradually approaches the co-polar field in terms of power as the frequency increases, corresponding with the loss in symmetry. This gives rise to high values for the cross/co-polar ratio, and inferior performance compared to the piecewise equivalent which displays lower levels. This is shown in figure 4.13a. It is also clear that the return loss of this horn is higher than in the piecewise case, however at 100 GHz the return loss of the horn is at the level required. Since the return loss decreases with frequency, it is clear that the mechanism that the optimiser identifies as controlling the return loss works at the expense of reducing beam symmetry and increasing cross-polar power.

From designing both horns with respect to minimising the return loss, it has been observed that the piecewise conical profile horn gives the better end result in terms of beam shape, side-lobe levels, shouldering, cross-polar power, symmetry and return loss, although at the design frequency the return loss of the GPHA is acceptable. The piecewise horn also gives excellent bandwidth performance, despite being designed specifically for 100 GHz operation, which is surprising as the changes in flare angle that are used tp ultimately control its performance are highly wavelength dependant and specific. It is clear that the range and scope of the profile parameters available to the GPHA horn are not sufficient to compete with the piecewise design when optimised according to this metric. This is logical, as the piecewise design allows significantly more scope to alter the horn geometry with more direct control of the profile as radii of each section are directly being adjusted, whereas the GPHA simply alters the shape of an overall function with no ability to customise individual portions of the horn.

When analysed according to the same metrics as the horns designed with respect to the co-polar symmetry, the piecewise conical profile horn optimised for minimising return loss performs significantly better. This implies that this parameter naturally improves performance in the

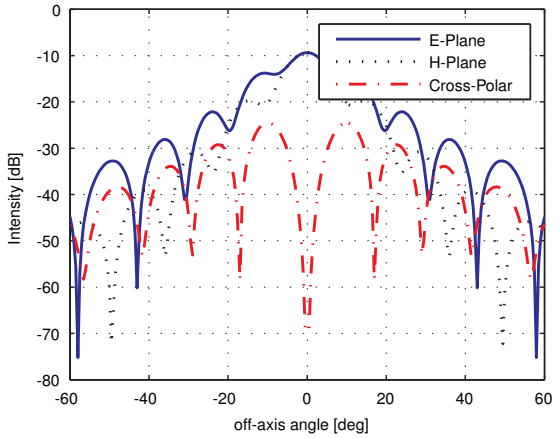

(a) 85 GHz

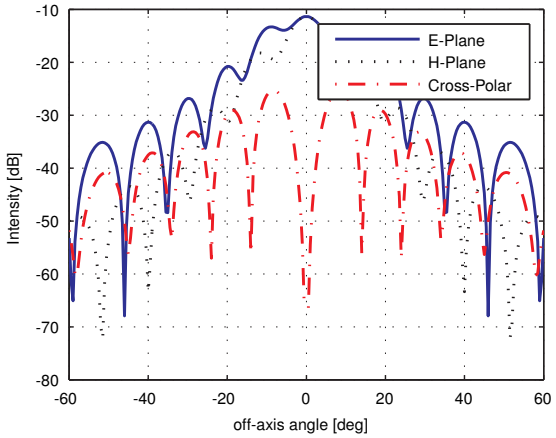

(b) 100 GHz

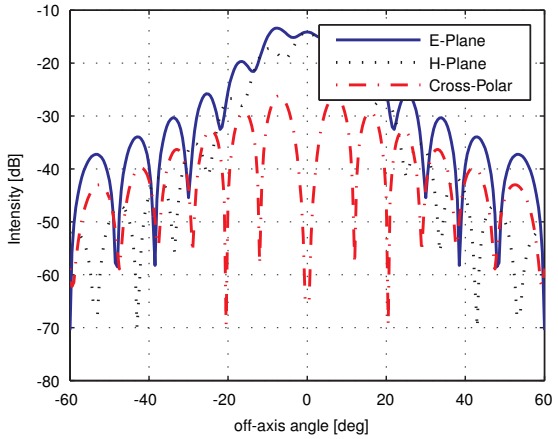

(c) 115 GHz

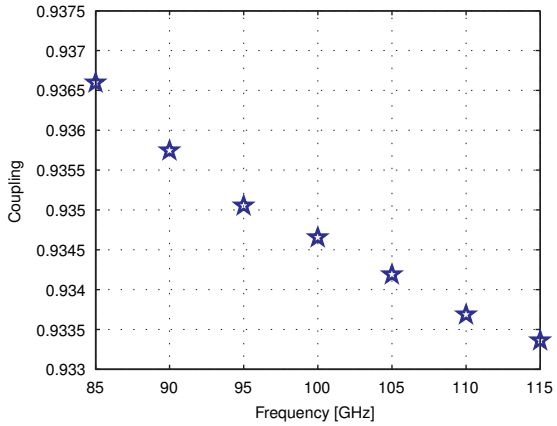

(d) Co-polar symmetry

Figure 4.12: Cuts and co-polar beam symmetry for the GPHA horn optimised with respect to return loss

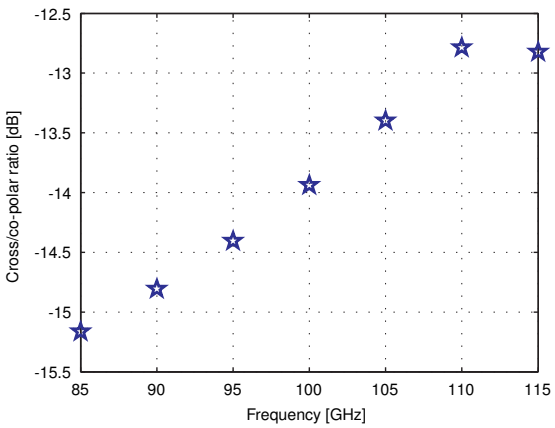

(a) Cross/co-polar ratio

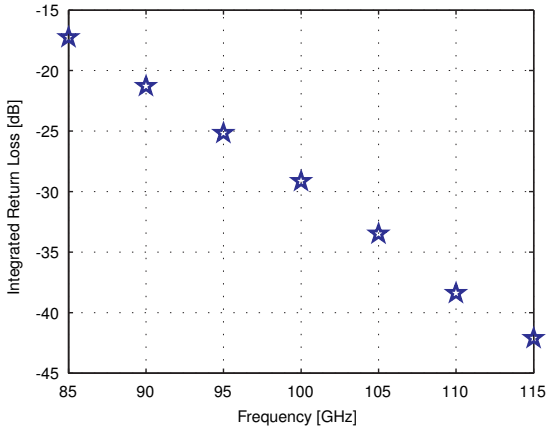

(b) Return loss

Figure 4.13: Cross/co-polar ratio and return loss for the GPHA horn optimised with respect to return loss

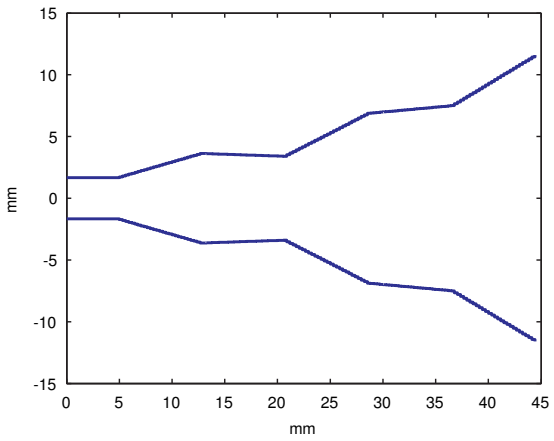

Figure 4.14: Piecewise conical profile horn optimised for minimum cross-polar levels

other areas, or that the mechanism that the optimiser identifies to control the return loss naturally improves the other parameters at the same time. In either case, using the co-polar symmetry as the optimisation metric yielded below par performance in terms of the metrics of interest, however by using the return loss as the metric of interest and a piecewise conical profile horn as the basis for the optimisation, excellent performance was obtained despite using only one metric and one frequency to drive the optimisation process, and the design was found using only one optimisation algorithm, on a standard desktop computer and on a timescale of the order of 5 hours. The corresponding GPHA, giving worse performance, took 7.7 hours despite using less horn parameters in the design process. In each case, the optimisation with respect to co-polar symmetry took approximately an additional 3.5-4 hours due to the fact that this optimisation requires the calculation of two field cuts for each horn generated by the optimiser, whereas the return loss calculation requires only the analysis of the S_{11} scattering matrix.

4.3.3 Minimising Cross-Polar Levels

In this section, the optimisation of both horn designs with respect to the maximum level realised in the cross-polar field, as outlined in section 4.2.2, will be discussed. The same bounding values were applied to the horn profile parameters as in the previous cases. Figure 4.14 shows the resulting horn in the case of the piecewise conical profile horn. Again, there are no sudden jumps in the profile or odd shapes appearing, and so this design is successful in terms of achieving ease of manufacture. Examining cuts of the field patters, as shown in figure 4.15, shows that the horn performs extremely well in terms of field pattern when compared to the other cases studied. This performance is retained across the band. Of particular note is the symmetry in the co-polar beam which is significantly higher than in previous cases, a point which is quantified in figure 4.15d, which shows coupling of above 98% across the band. The optimised horn also clearly meets its design target, as we see that the cross-polar field

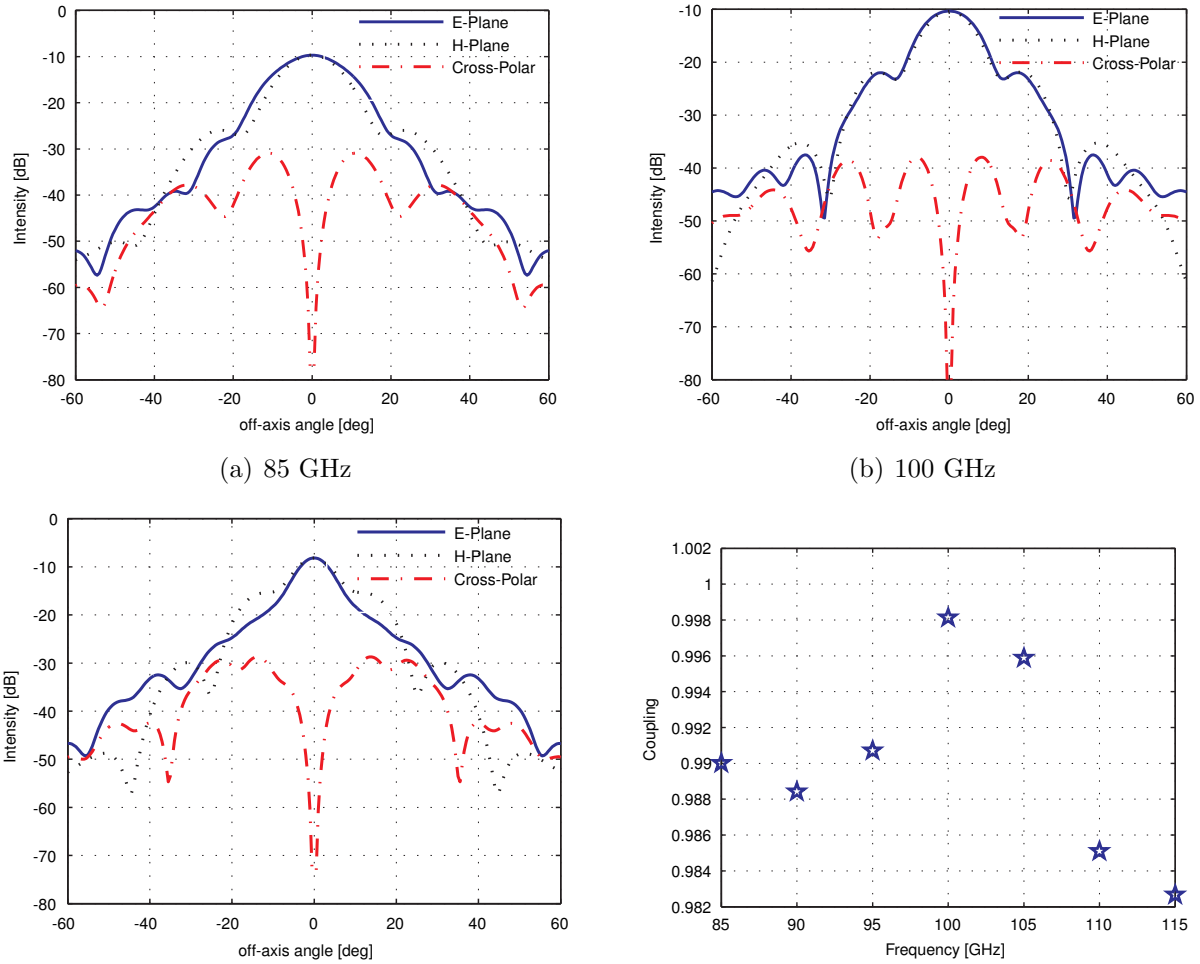

(a) 85 GHz

(b) 100 GHz

(c) 115 GHz

(d) Co-polar symmetry

Figure 4.15: Cuts and co-polar beam symmetry for the piecewise conical profile horn optimised with respect to maximum cross-polar level

is separated significantly from the co-polar field in terms of power. Previous optimised horn designs did not necessarily excel in the area that they were designed to, however optimising with respect to cross-polar power has rectified this issue. The cross-polar performance is clearly shown in figure 4.16a, where low levels of power relative to the co-polar field are observed across the band. The return loss across the band is shown in figure 4.16b, which performs extremely

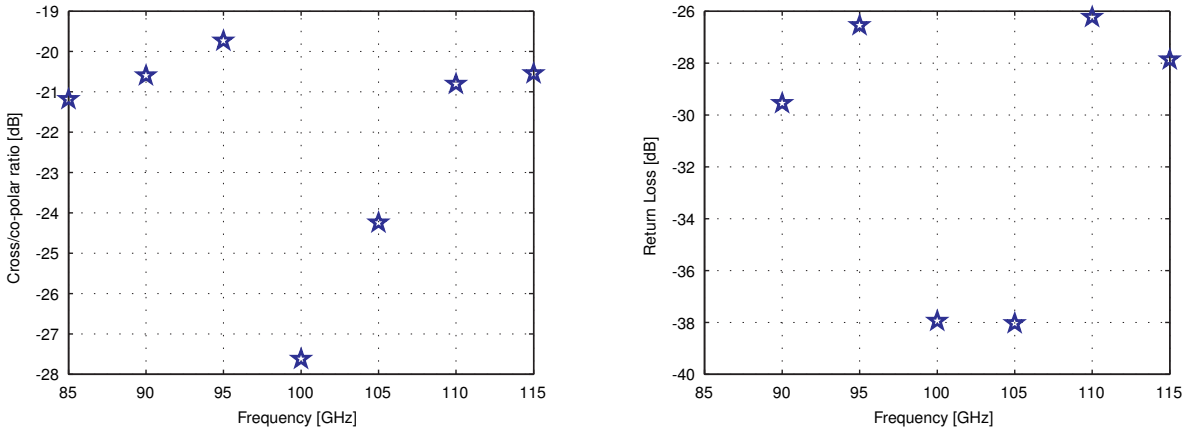

(a) Cross/co-polar ratio

(b) Return loss

Figure 4.16: Cross/co-polar ratio and return loss for the piecewise conical profile horn optimised with respect to maximum cross-polar level

well, realising extremely low levels.

In order to verify the repeatability of the genetic algorithm based optimisation process, the cross-polar optimisation process was repeated for the piecewise conical profile horn with the same starting values. The results are shown in table 4.4. Figure 4.17 gives a comparison between the profiles obtained using the original result and the result from the repeat of the process. This comparison shows no real distinguishable difference between either design, which confirms what is observed when the values of the radii and length for the two geometries are considered. In the repeated case, each parameter (when displayed to two decimal places-a reasonable figure for manufacturing tolerancing) lies within ±0.04 mm of the original value. When

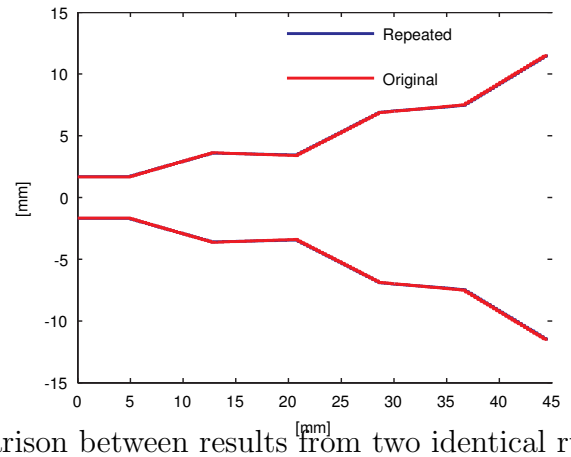

Figure 4.17: Comparison between results from two identical runs of the GA optimiser

more significant figures are used, there is some additional variation however when represented to this precision (which is consistent with manufacturing tolerances) the variation lies within this range.

Figure 4.18 compares the principal cuts at 100 GHz and the cross/co-polar ratio across the band for the original run and the repeated run in order to establish repeatability of the optimisation process. It is clear from the indicated cuts that there is a high level of agreement

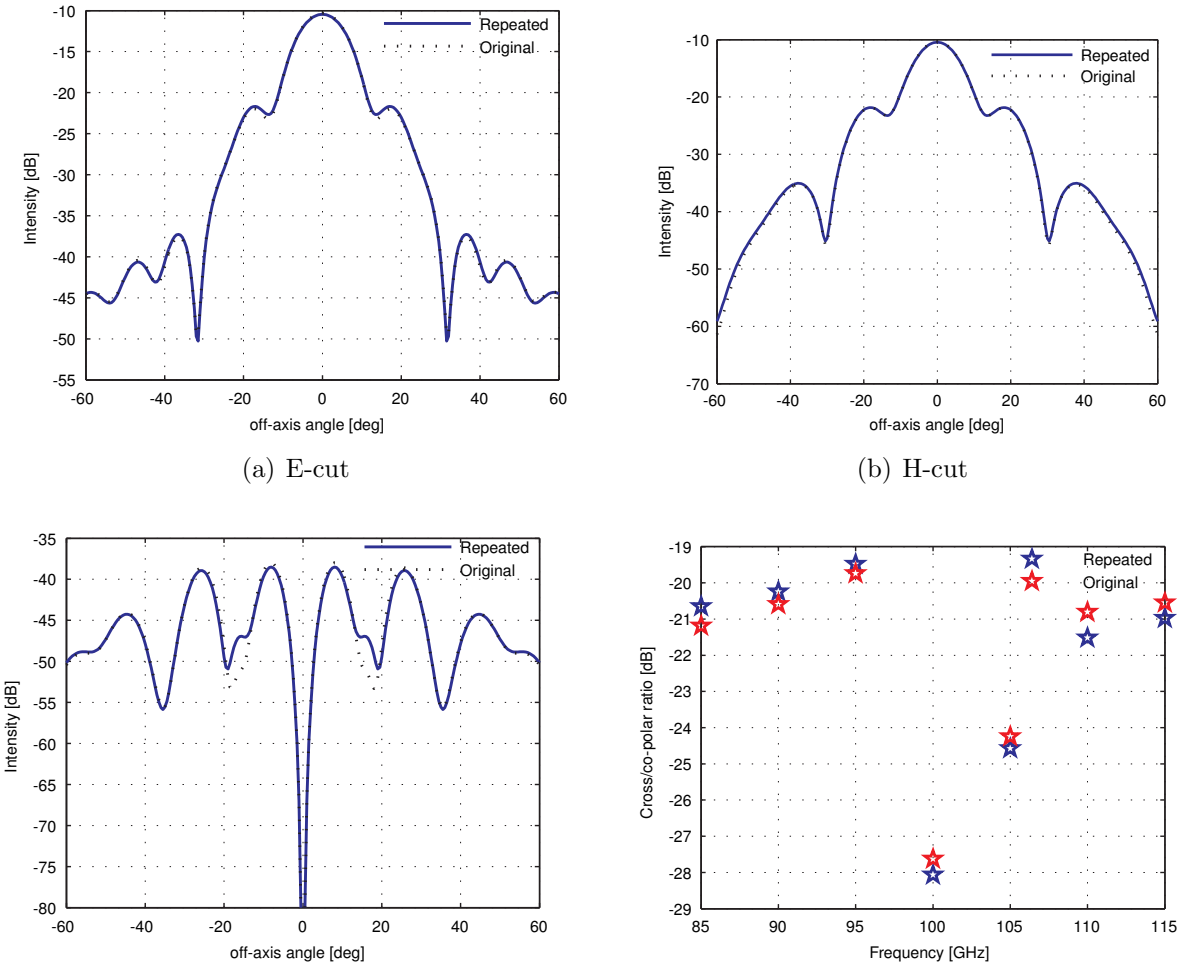

(a) E-cut

(b) H-cut

(c) Cross-polar (45°) cut

(d) Cross/co-polar ratio

Figure 4.18: Principal cuts at 100 GHz and cross/co-polar ratio when the optimisation process is repeated to test for repeatability

between the two designs, implying that the differences in dimensions do not have a detrimental effect on the performance of the horn. This is as expected, as the differences are small in terms of wavelength. Although the flare angles will change for the different parts of the horn, this does not appear to alter the performance significantly. The more noticeable disagreements that occur in the cross-polar pattern occur at extremely low levels, and so are not of concern. Examining the cross/co-polar ratio, the metric that was optimised with respect to, indicates a good level of agreement between the two, as would be expected. This shows that the optimisation process is

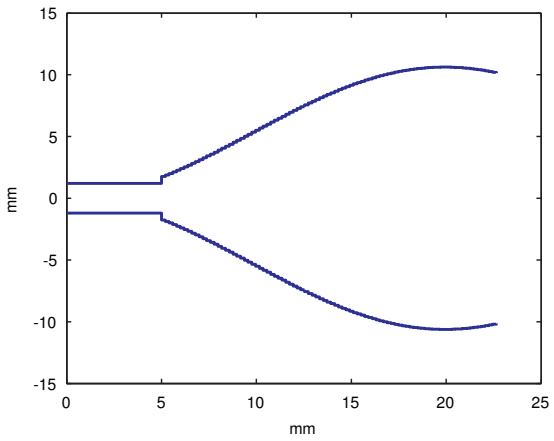

Figure 4.19: GPHA horn optimised for minimum cross-polar levels

repeatable to within the limits required for the resulting horns to physically perform similarly, and so the design process is robust.

The same optimisation process was applied to the GPHA design, with the optimised parameters shown in column 3 of table 4.5. The profile of the horn is shown in figure 4.19. The principal planes and co-polar beam symmetry across the band are again shown in figure 4.20. In comparison with the corresponding piecewise conical design, the performance in terms of beam pattern is extremely poor, with on-axis nulls appearing in both of the principal cuts of the co-polar beam until 115 GHz. The sidelobes in this case are plentiful and at an extremely high level when compared to the on-axis value, and so this is clearly not acceptable in the context of the requirements of the horn. There is also a high degree of asymmetry between the principal planes of the co-polar beam, as quantified by the low coupling between these cuts across the band. The cross-polar pattern contains a significant amount of power relative to the co-polar beam, such that their levels are comparable. This is shown in figure 4.21a, which shows that the cross/co-polar ratio is now extremely high. Figure 4.21b shows the return loss of the horn across the band. This is in general poorer than for the piecewise conical profile equivalent, particularly at the design frequency.

It can be seen from this comparison that when the GPHA and piecewise conical profile horn designs are optimised with respect to minimising the cross-polar power, that the piecewise conical profile horn gives the best results in terms of the parameters of interest, and the profile is also easier to machine which is a desirable quality. The piecewise conical horn took 6 hours to design, with the GPHA taking 8.5 hours. In both cases, the time taken is somewhere between that required for the return loss and cross-polar optimisations. It is greater than that needed for the return loss optimisation as a field calculation is required, but as only one field is necessary it takes less time than the symmetry optimisation as two fields are necessary in that case. The symmetry optimisation takes significantly longer, as the cost function is more complex due to the fact that changing the horn will affect the two fields in potentially different

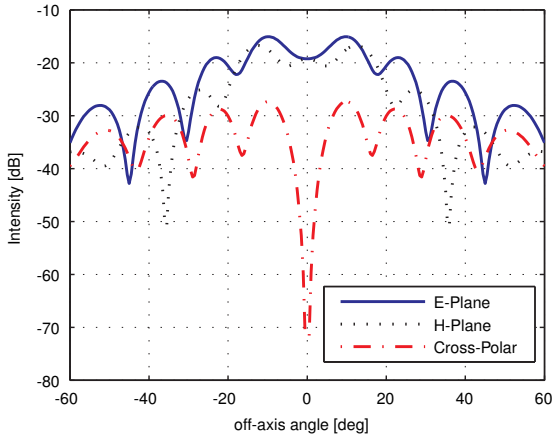

(a) 85 GHz

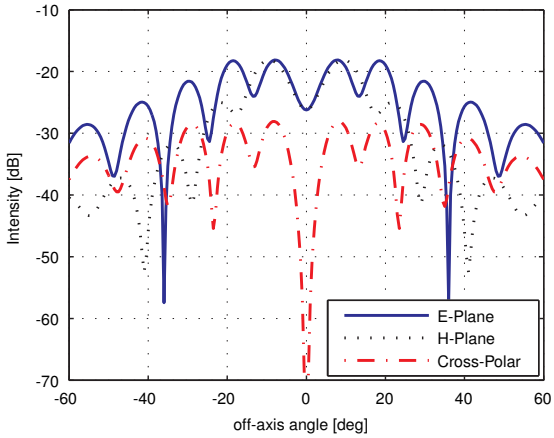

(b) 100 GHz

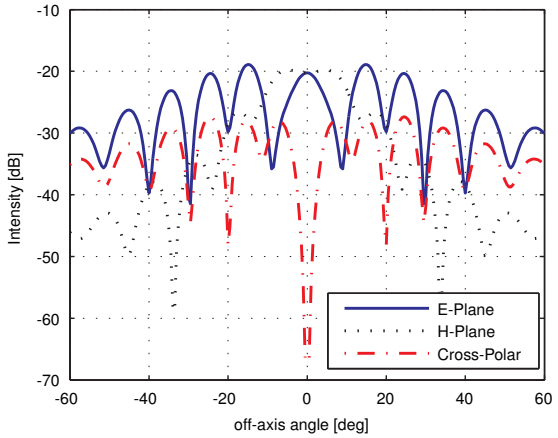

(c) 115 GHz

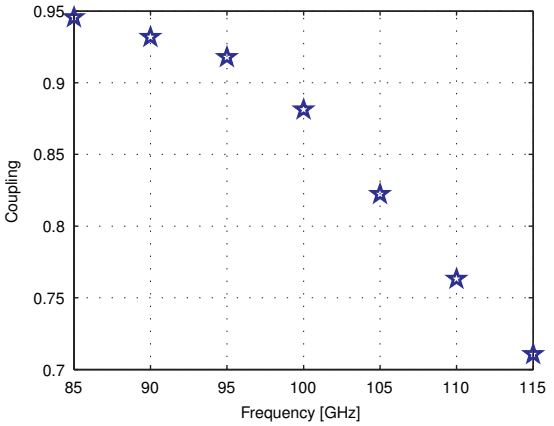

(d) Co-polar symmetry

Figure 4.20: Cuts and co-polar beam symmetry for the GPHA optimised with respect to maximum cross-polar level

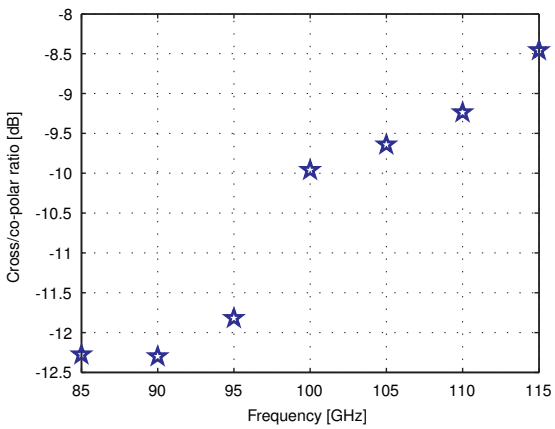

(a) Cross/co-polar ratio

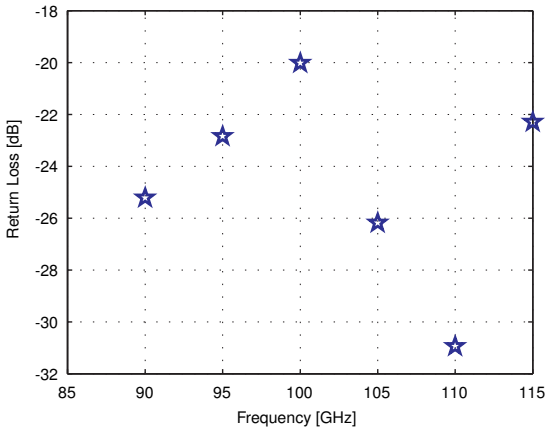

(b) Return loss

Figure 4.21: Cross/co-polar ratio and return loss for the GPHA, optimised with respect to maximum cross-polar level

ways, and so this takes significantly longer to solve.

The difference in execution time, and ultimately electromagnetic performance, between the piecewise and GPHA designs can be explained as follows. As the GPHA horn requires a single curvature term to vary the geometry file, subtle changes in the profile lead to large changes in the optical performance and requiring the algorithm to optimise the curvature term rather than radii of individual sections making up the profile does not easily allow the shape to be optimised in an efficient and robust way. To resolve this issue, it is necessary to allow the optimiser have more control over individual sections of the horn. Redefining the geometry to consist of many sections makes the optimisation space very large requiring a large computational effort to vary the radii of all sections as free parameters. This makes the optimisation of such a horn difficult on a normal PC. The presented piecewise conical profile horn provides a compromise by allowing the algorithm to control the profile of several large sections, a sufficient number so that the performance of the horn can be manipulated sufficiently to realise the required performance.

It is clear that minimising the cross-polar power naturally improves performance with respect to the other metrics analysed. Optimising with respect to the other metrics tended to produce results which at best improved performance in terms of that metric, but showed little to no improvement for the other metrics. This is a sensible result, as low cross-polarisation power tends to be a feature of horns which have symmetric co-polar beams, with little reflection (return loss) and low sidelobe levels. Therefore, if a horn is optimised for low sidelobe levels then good performance in these other areas might be expected. The reverse does not necessarily follow. A horn may be designed for example, to minimise return loss, but this does not require good performance in terms of other parameters.

By optimising with respect to several metrics which will be important for the application of a horn design which will be presented in the following chapter, the limitations of the genetic algorithm as applied to these metrics and the horn designs that were considered were investigated. It has been established that using a piecewise conical profile horn and optimising with respect to minimising cross-polar power levels gives rise to the best overall results, both at the design frequency and across the band. It was also shown that the genetic algorithm results are repeatable, producing results that agree to within acceptable limits. Further agreement could be obtained by using the genetic algorithm results as the starting point for a local optimiser to search and find the precise location of the minimum, however the purpose of this investigation is to find a method of designing a physically useful horn using the smallest computational load possible. This has been achieved, with consistent results, using only the genetic algorithm.

4.3.4 Comparison of Best Optimised Horn with a Similar Smooth-Walled Horn

As stated at the beginning of this chapter, the optimisation algorithms do not have a fixed starting point available to the user (it merely allows a range of values to be specified for each parameter), so it is impossible to examine the performance of an initial horn, and compare the performance of the horn that results from the optimisation process to this to examine the improvement in the performance of the metric that was optimised with respect to. It is however useful to compare the performance of an optimised horn to some other typical design for context. One goal of the design process was to establish a design that was simpler than a corrugated horn, and therefore cheaper to manufacture. It makes sense therefore to compare the best overall performing horn obtained from the optimisation process (a piecewise conical profile horn optimised with respect to minimising cross-polar levels) to an equivalent smooth-walled horn to examine the improvement in performance that the optimised horn brings relative to the smooth-walled horn, which is of course a viable alternative that is always worth considering.

The smooth-walled horn was modelled using the same principal dimensions as the piecewise conical horn that was described in section 4.3.3, giving rise to an equivalent smooth-walled horn. These dimensions are shown in table 4.6. A profile view of the horn is shown in figure 4.22.

Throat radius [mm]	1.66
Aperture radius [mm]	11.59
Throat length [mm]	5.00
Slant length [mm]	39.71

Table 4.6: Dimensions of the smooth-walled horn used for benchmarking the relative performance of the best performing optimised horn

The simulations were carried out at a frequency of 100 GHz. Although the throat radius of the horn is sufficiently large at this frequency to be few-moded (a mode of azimuthal order $n = 0$ is supported), for the purpose of comparison the smooth walled horn is assumed to be fed by a WR10 waveguide, as was the case with the piecewise horn. This means that only modes of azimuthal order $n = 1$ were considered in the simulation. The principal planes and the 45° cut of the cross-polar plane are shown at frequencies of 85, 100 and 115 GHz in figure 4.23, along with the co-polar symmetry of the horn across the band. Examining the field patterns of the smooth-walled horn and comparing them to those of the piecewise conical horn, as shown in figure 4.15, shows that the beam patterns produced by the piecewise conical horn are superior for application in CMB/Far-IR missions. The sidelobes and shouldering structures in the piecewise horn are less than in the smooth-walled case, with these structures also initially appearing at lower levels than in the case of the smooth-walled horn. The symmetry of the co-polar field is also greater for the piecewise horn, showing values of above 98% across the

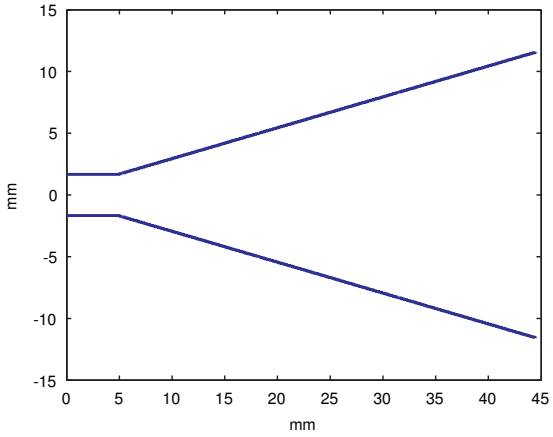

Figure 4.22: Profile view of the equivalent smooth-walled horn

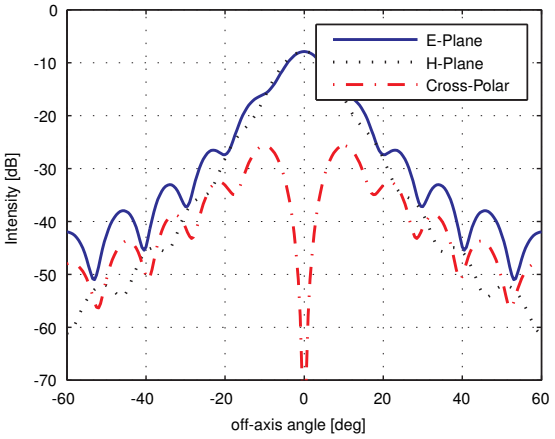

(a) 85 GHz

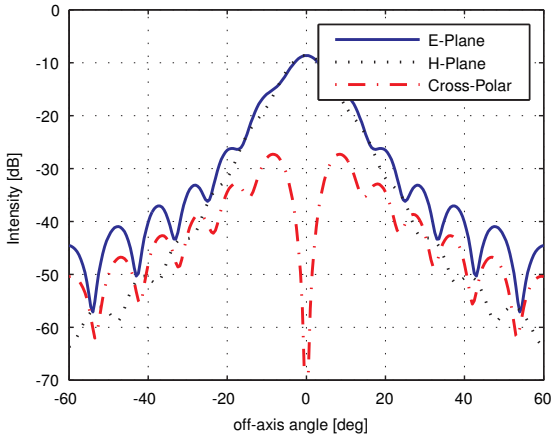

(b) 100 GHz

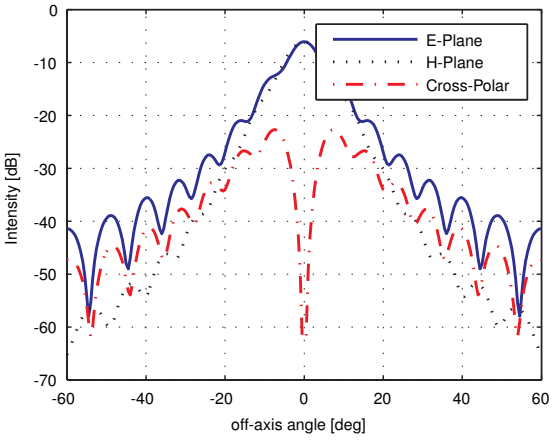

(c) 115 GHz

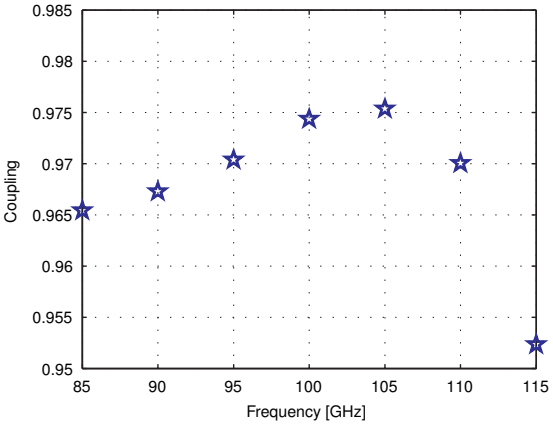

(d) Co-polar symmetry

Figure 4.23: Cuts and co-polar beam symmetry for the equivalent smooth-walled horn

band, whereas the smooth-walled horn shows values of between 95 and 97.5% across the band.

The cross-polar field for the smooth-walled horn also contains more power relative to the co-polar field, compared to the piecewise conical horn. This is quantified in figure 4.24a, which indicates that the ratio of the maximum cross-polar level to the maximum co-polar level varies between -16.5 dB and -19.1 dB across the band. The piecewise horn displays an improved cross/co-polar separation across the band, with the value of this ratio in that case varying between -19.8 dB and -27.7 dB. This is the metric that the piecewise horn was designed with respect to, so its improved performance in this regard is of no surprise. The improvement is noted across the band, despite the fact that the horn was designed specifically for 100 GHz operation, a fact which is highlighted by the superior performance of the piecewise horn at 100 GHz, realising a cross/co-polar ratio of -27.7 dB, compared to the smooth-walled horn which realises a value of -18.6 dB. The return loss of the smooth-walled horn is shown in figure

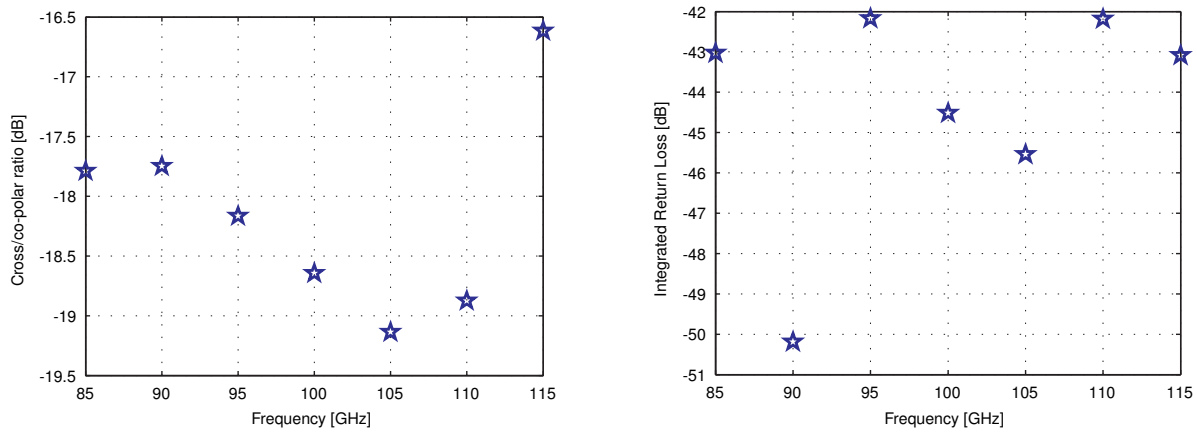

(a) Cross/co-polar ratio

(b) Return loss

Figure 4.24: Cross/co-polar ratio and return loss for the equivalent smooth-walled horn

4.24b. The performance of the smooth-walled horn in this regard is superior to that of the optimised horn. The smooth-walled horn experiences a return loss of between -42 and -50 dB across the band, while the piecewise horn experiences a return loss of -26 to -38 dB across the band, however this is still an acceptable level of return loss.

Comparing the performance of the piecewise horn and its smooth-walled equivalent showed that the piecewise horn is superior in each of the areas of interest, apart from return loss, although the return loss of the piecewise horn is still acceptable. The piecewise horn shows better co-polar symmetry with lower and fewer sidelobe/shoulder structures, in addition to an increased cross/co-polar separation. For the applications of interest to this thesis, the piecewise design is therefore a superior option to a smooth-walled horn.

4.4 Optimisation Using Simulated Annealing

The simulated annealing (SA) algorithm was integrated with the cost function that was used with the genetic algorithm, allowing the optimisation process to be carried out using SA. The goal was to repeat, using the same bounding values for each variable, the optimisation that were carried out in the previous section. It was found that while the GA returned the solutions on the order of 6-13 hours, SA required significantly longer. In each case, despite being given perhaps 8 days to run, a solution was not obtained with either the algorithm hanging or the computer running out of memory. This is due to the fact that the algorithm searches only small portions of space at a time, [52], and so if the algorithm does not commence its search in the neighbourhood of the correct solution it will spend a significant amount of time attempting to locate the correct solution, often hanging due to becoming trapped in a region that is difficult to progress from. For this reason, the role of SA is best suited to discrete optimisation problems where the approximate solution is already known, as it considers one solution at a time. The GA, considering a pool of solutions at any one time that potentially span the entire space, is better suited to problems where there is a large search space with no approximate solution known in advance, as is the case in the optimisation problems examined in this thesis.

With the use of a super computer, and/or where time is not a factor in the optimisation process, the simulated annealing algorithm will eventually converge to an optimised solution, however a primary objective of this investigation was to focus on the optimisation of a horn antenna design using standard computing power over a reasonable time period. Since the indicated timeframe for SA violates this requirement, the genetic algorithm is clearly preferable. In order to show that the SA will still produce the correct result, certainly within the limits of repeatability that were outlined for the GA, a scaled down optimisation was carried out. The piecewise conical profile horn was optimised with respect to the maximum cross-polar level, however instead of the standard seven optimisation variables, the cost function was made a function of one variable, the aperture radius. The other variables were assigned fixed values, as per table 4.4, with the aperture radius allowed to vary within the limits set out for $y6$ in table 4.2.

After approximately 25 hours, the optimisation process returned a value of 12.02 mm for the aperture radius, compared with a value of 11.59 mm that was obtained using the genetic algorithm. The difference between the two values is of a similar magnitude to those encountered when testing the GA for repeatability, and so it is expected that any differences in performance will not be severe. Figure 4.25 shows the principal planes for the resulting horn at 100 GHz, in addition to cross/co-polar ratio across the band. It is clear from this that once again the impact on the field patterns is minimal. The cross-polar pattern is slightly lower at 100 GHz, which tallies with the cross/co-polar ratio being lower for the SA case at 100 GHz. The ratios take on similar values at the other frequencies, meaning that once again the difference in performance

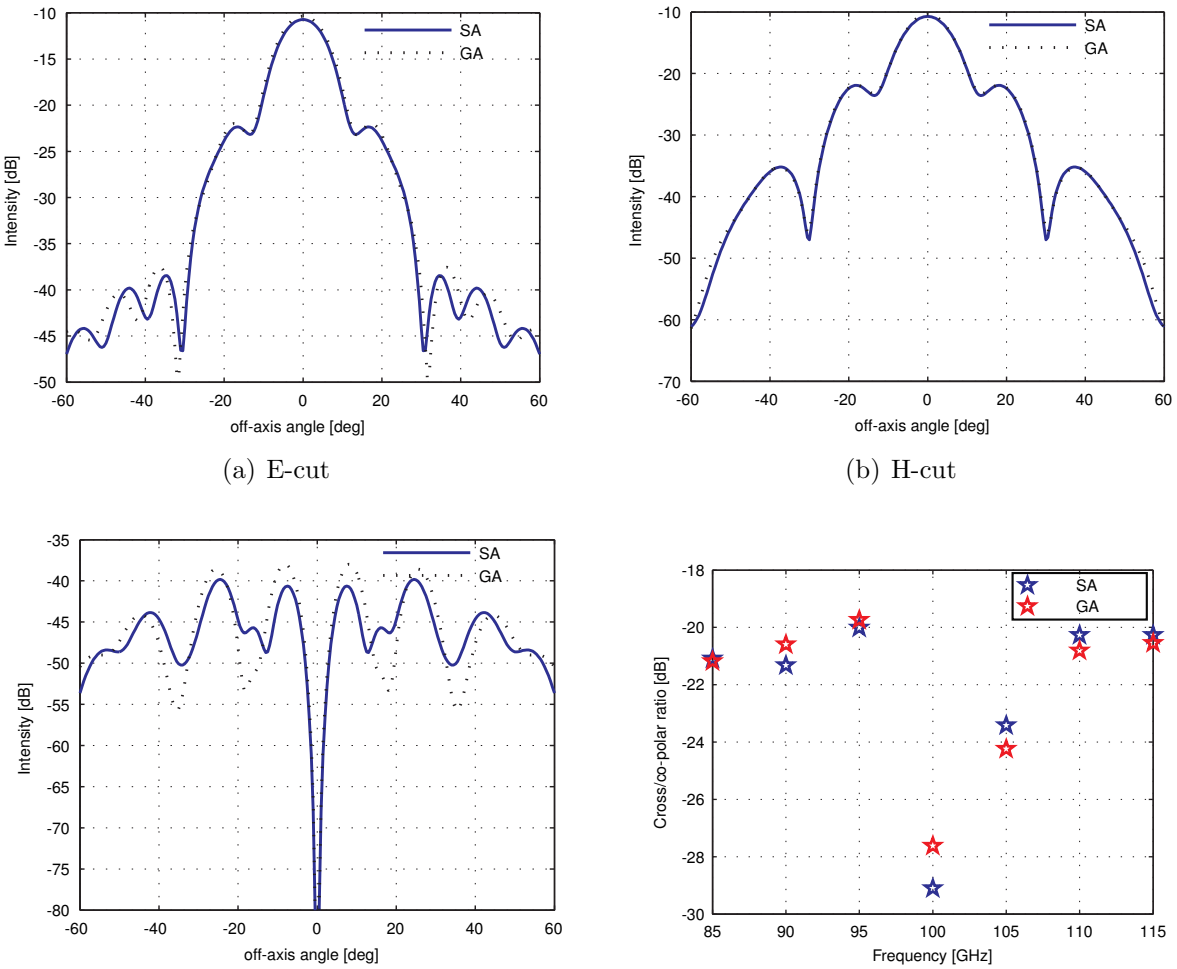

(a) E-cut

(b) H-cut

(c) Cross-polar (45°) cut

(d) Cross/co-polar ratio

Figure 4.25: Principal cuts at 100 GHz and cross/co-polar ratio when the piecewise conical design is optimised for maximum cross-polar power varying only the aperture radius

between the designs that result from the two difference calculations are minimal, certainly when the physical performance of the horn is examined. It can be reasonably concluded from this that given sufficient time, the SA process would iterate towards the same solution as that obtained using the GA.

4.4.1 Summary of Performance of Algorithms/Horns

In this section, a brief summary will be given of the merits of each algorithm along with the performance of each horn design in terms of the overall optimisation process. This is merely to provide a quick reference summary for future work in this area. A comparison of the SA and GA optimisation algorithms in the relevant areas is given in table 4.7. The simulated

	Genetic Algorithm	**Simulated Annealing**
Number of samples per iteration	Many	One
Size of potential parameter space	Large	Large
Need to know approx solution in advance	No	Yes
Execution time for given parameter space	Quick (hours)	Lengthy (days)

Table 4.7: Comparison of the SA and GA optimisation algorithms

annealing algorithm evaluates one sample of the parameters in each iteration, thus searching a small amount of parameter space at a time. The genetic algorithm evaluates many samples per iteration and thus spans a large section of parameters space in each iteration, converging on the approximate minimum rapidly. Thus, although both algorithms can handle large parameters spaces, the SA will take significantly longer to find the approximate minimum of the function. To counter this it is necessary to know the approximate solution in advance when applying the SA, which renders it of little use to the application described in this thesis, where a key factor is not knowing the solution in advance. By using the GA, the entire defined parameter space can be searched relatively quickly and a good approximation to the global minimum obtained, and so this algorithm is used.

Considering the horn types examined, it is clear that the piecewise conical profile horn yields the best solutions in each case. This is likely due to the fact that the algorithm can in effect 'micro manage' this horn more effectively as it has much more flexibility to manipulate the radii of each section directly. In the GPHA case, the optimiser can only manipulate the shape of an overall function and so is limited in what it can achieve by this. It is also possible that whatever mechanisms the optimiser identifies when optimising the GPHA relative to a given metric do not offer as much flexibility or responsivity as in the piecewise conical case.

It is also clear that whilst optimising with respect to each metric gave rise to the best possible performance for that metric at 100 GHz within the limits of what was possible with the control that the optimiser had over the mechanisms operating within the horn by means

of manipulating the geometry,the best overall performance occurred when the horns were optimised with respect to minimising the cross-polar levels. This gave rise to good performance for all relevant metrics rather than just the one being optimised. This is because the other parameters such as symmetry and return loss follow naturally from a low level of cross-polar power, whereas the opposite is not necessarily true. Thus, by minimising the cross-polar power the other parameters are in a sense forced to also improve.

4.5 Optimisation Technique Applied to Absorbing Integrating Cavity

In this section, the method of including a sheet of absorbing material in the mode-matching code SCATTER will be considered. Such absorbers can be used to simulate bolometer detectors which use an absorbing layer that are placed within optical systems. In an integrating detector cavity, such as those proposed in the SAFARI system (this is just one example of such a system and so will be described in context in section 8.4), the detector will ultimately cause the absorption of the radiation present in the cavity, with the cavity itself being fed by a feed horn antenna. As the cavity is closed at one end, radiation that is not absorbed when initially incident upon the detector (as the efficiency will always be less than unity) may be absorbed upon subsequent passes through the detector having been reflected within the cavity, as opposed to being lost to the system. An example of a very simple sealed cavity is given in figure 4.26. Here, a waveguide of radius r is open at one end, allowing radiation to enter it,

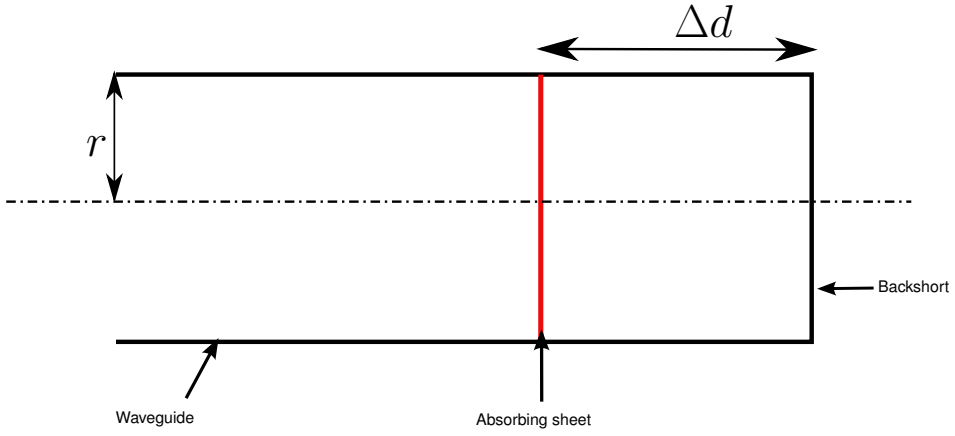

Figure 4.26: Profile view of a simple waveguide based sealed cavity containing an absorbing sheet

however it is sealed at the other end, with this configuration being known as a backshort. A section of absorbing material is located a distance Δd from the backshort. Due to conservation of complex power, any radiation that enters the cavity must either be absorbed by the sheet or re-appear at the entrance aperture of the cavity. The use of an integrating cavity therefore ensures that a significant amount of the radiation that is not absorbed by the detector initially

will be absorbed on subsequent passes, as it remains trapped in the cavity, thus increasing the efficiency, η of the detection system. Naturally, the cavity walls will be subject to ohmic losses, however the detector efficiency should be significantly increased compared to standard systems.

A bolometer is a sensor whose resistance is sensitive to changes in temperature. Typically, it consists of a section of absorptive material (which is what is modelled in the system presented here) that is connected to a body of constant temperature (a thermal reservoir) by means of a thermal link. When radiation is incident on the absorptive material, the material experiences a rise in temperature, and so its temperature will rise above that of the thermal reservoir. The sensor functions by employing a detector (either a resistive thermometer or by measuring directly the resistance of the absorptive material) to monitor the temperature difference between the material and the reservoir. Modern bolometers tend to exploit the properties of semiconductor or superconducting materials, as these can be used to achieve an extremely high level of sensitivity when operated at extremely low temperatures of the order 50-300 mK.

For astronomy that takes place at sub-millimetre wavelengths, bolometers rank highly in terms of being some of the most sensitive detectors available. In order to achieve this sensitivity, and to readily measure radiation whose equivalent temperature is extremely low, they are cooled to temperatures of just above absolute zero, typically in the range 50-300 mK. Examples of projects where bolometers have been used for astronomy at these wavelengths are the Herschel Space Observatory and the James Clerk Maxwell Telescope. The detector will therefore have inherent properties such as, for example, the resistivity of the absorbing material. If these properties can be used to define the absorbing sheet placed in SCATTER, then the behaviour of the sheet, and therefore the overall bolometer, can be modelled, thus giving an indication as to the performance of the system, including the effects of the detector.

4.5.1 Theory

This section describes the implementation of a thin absorbing sheet within the mode matching framework SCATTER, as described by Dr. Stephen Doherty in [57]. SCATTER considers a two port system that consists of several waveguide sections. The material from which the waveguide sections are made is assumed to be perfectly electrically conducting, and this is implemented by the requirement that the total power in the system be equal to the sum of the power distributed across each of the ports when both the forward and backwards propagating modes at each port are considered, in addition to the boundary conditions that result from this assumption, as described in chapter 2. This results in no power being lost from the system. The addition of a sheet of absorbing material amounts to introducing a mechanism for allowing loss of power from the system at a specific plane within the waveguide, with the severity of the loss depending on the properties of the absorbing material.

It is assumed that the sheet is infinitely thin and that it is larger than the spatial extent of the incident modes, whose power is confined to well within the area of the absorber. To consider the case where this is not true would require the addition of two new mode sets; one that describes the propagation of the modes within the dielectric material itself, and another set of co-axial modes that would govern the propagation of the waveguide modes beyond the extent of the absorber, in the region between the material and the walls of the waveguide. Solving the system for such a combination of mode sets would be a complex issue, and so for the purposes of the work that will be described in this thesis, a simplification is employed in which the absorbing material is modelled as a thin, two-dimensional resistive sheet (of dimension greater than the spatial extent of the modes in question), which is found to work very well, yielding against agreement with experiment and other modelling techniques, [57]. One constraint of the approach that will be described is that it does not allow one to directly examine the power lost by the absorbing material. To examine the affect of the absorber, the scattering matrices for the overall system are required, and so if the absorber is the only source of loss in the system then this will of course give the power absorbed. If there is another mechanism for loss in the system, such as a free space gap for example as will be described in chapter 8, then both losses are combined in the overall scattering matrices and it is impossible to distinguish how much loss can be attributed to either mechanism. To do this, a system consisting of several additional ports would be necessary, however this is beyond the scope of what is necessary for this work, as will be outlined in chapter 8.

The boundary conditions that apply to an absorbing sheet that fills a waveguide section mean that the electric field remains continuous across the sheet, with the magnetic field changing value by an amount equal to the current density in the sheet. This means that the co-axial mode set would have to exist over the null space (as the sheet/waveguide separation is zero), and so it is assumed that they need not be considered when the sheet is relatively large, and so is larger than the spatial extent of the incident modes. This is summarised by equations 4.8 to 4.10, where the subscripts indicate the left and right sides of the junction with the resistive sheet. E, H and R_s are respectively the electric field, the magnetic field, and the impedance of the resistive sheet in units of ohms per square, Ω/sq. This value is set according to the detector that it is desired to model.

$$\mathbf{E}_L = \mathbf{E}_R \tag{4.8}$$

$$\mathbf{H}_L = \mathbf{H}_R - \mathbf{K} \times \hat{k} \tag{4.9}$$

$$\mathbf{K} = \frac{\mathbf{E}_L}{R_s} = \frac{\mathbf{E}_R}{R_s} \tag{4.10}$$

It should also be noted that the impedance of the resistive sheet is a function of the radius, r, such that inside the area of the absorber ($r \leq r_{abs}$, if r_{abs} is the radius of the absorber)

$R_s(r) = R_s$ and outside of the absorber ($r > r_{abs}$) $R_s(r) = \infty$.

As detailed in chapter 2, if the n^{th} mode on the left hand side of a junction is represented by the forward and reflected amplitude coefficients A_n and B_n respectively, and the equivalent coefficients on the right had side of the junction are represented by C_n and D_n, then the electric and magnetic fields on each side of the junction can be written as

$$\mathbf{E}_L = \sum_{n=1}^{N} \left[A_n \exp(-\gamma_n z) + B_n \exp(\gamma_n z)\right] \mathbf{e}_n \tag{4.11}$$

$$\mathbf{H}_L = \sum_{n=1}^{N} \left[A_n \exp(-\gamma_n z) - B_n \exp(\gamma_n z)\right] \mathbf{h}_n \tag{4.12}$$

$$\mathbf{E}_R = \sum_{n=1}^{N} \left[D_n \exp(-\gamma_n z) + C_n \exp(\gamma_n z)\right] \mathbf{e}_n \tag{4.13}$$

$$\mathbf{H}_R = \sum_{n=1}^{N} \left[D_n \exp(-\gamma_n z) - C_n \exp(\gamma_n z)\right] \mathbf{h}_n. \tag{4.14}$$

If the boundary conditions outlined in equations 4.8 and 4.9 are applied and an equal number of modes are assumed on each side of the junction, then proceeding by taking the Poynting cross products as in the mode matching technique, yields

$$Q \cdot [A + B] = Q \cdot [D + C] \tag{4.15a}$$

$$Q^* \cdot [A - B] = Q^* \cdot [D - C] R^{-1} \Gamma [A + B], \tag{4.15b}$$

where Q and R are square power matrices with their usual meaning in the context of mode-matching, and $\Gamma = \frac{Z_m P_{mn}}{2R_s}$, with Z representing the mode impedance and P_{mn} controlling the fractional inter-modal scattering. This integral need only be solved for scattering within the same azimuthal order, as no scattering occurs between the different azimuthal orders in the cylindrical case.

The sheet is considered to be infinitely thin, and so for $r > r_{abs}$, we have that $R_s = 0$, therefore the integral required to populate the P matrix is given by equation 4.16, and need only be solved within the *area* of the resistive sheet.

$$P_{absorb} = \int_{S_{absorb}} (\mathbf{e}_m \cdot \mathbf{e}_n)\, dS. \tag{4.16}$$

If equation 4.15 is solved for B in terms of A and C, then expressions for the reflected and

transmitted components can be found, as given in equation 4.17.

$$S_{11} = S_{22} = -\frac{\Gamma}{\mathbf{I} + \Gamma} \tag{4.17a}$$

$$S_{12} = S_{21} = \frac{1}{\mathbf{I} + \Gamma}. \tag{4.17b}$$

These scattering matrices are treated as representing one section in the standard SCATTER framework. They are therefore cascaded with the other scattering matrices in the system in the standard way, allowing the absorber to be efficiently included as an element within the waveguide structure being analysed. The net effect of the absorber is of course for power to be lost from the system. The code required to calculate the scattering matrices for the absorbing section is different to that used for a standard waveguide section. The absorber code is triggered within SCATTER by examining the section of waveguide after which the absorber is to be placed. In the geometry file, the section immediately following this section is defined to be of the required length, but the radius is set equal to that of the section prior to the absorber. When the code observed two consecutive radii of equal size, it triggers the absorber code and moves on to the next section as normal after this. The impedance of the absorber is defined within SCATTER in units of Ω/sq.

4.5.2 Implementation, Testing and Application to an Optimisation Problem

The code was originally implemented in Mathematica, as presented in [57]. The systems presented therein contain an absorber located within a waveguide cavity structure. The details involved in implementing such a structure in SCATTER are given in section 8.2.5. Since the optimiser is in Matlab, it was necessary to implement the code in the Matlab version of SCATTER. This was carried out for a general multi-moded system and tested against the results that were presented in [57], to check that they were reproducible using the Matlab code. Since the optimisation process currently deals with single-moded systems, the remainder of this discussion will be concerned with single-moded systems.

Consider a cylindrical cavity of length 10 mm and radius 1.26 mm that is being illuminated with radiation of frequency 80 GHz (λ = 3.75 mm). At this frequency the cavity is single-moded, supporting only the TE_{11} mode. The absorbed power therefore refers solely to this mode. To evaluate the power absorbed by the absorber, recall that SCATTER illuminates each mode with unity power. In the absence of the absorber, since the back of the cavity is closed and the waveguide is perfectly conducting, all the power would re-appear at port one where the illumination took place, and so the integrated S_{11} power would be unity, i.e. $|S_{11}|^2 = 1$. Since the absorbing sheet provides the only mechanism for loss in such a system, if it is introduced then any power that does not get returned to the excitation port (port 1) from the cavity is

considered to be lost due to the absorber, i.e. it is absorbed. The absorbed power is therefore given by $P_{abs} = 1 - |S_{11}|^2$. This will be examined again in section 8.2.5.

Figure 4.27 shows the absorbed power for the cavity as a function of absorber placement relative to the back (closed end) of the cavity, in terms of wavelength, and as a function of impedance. In the case where the placement is varied, the impedance is fixed at a value of $R_s = 377\ \Omega/\text{sq}$, and in the case where the impedance is varied, the placement is fixed at 1.95 mm from the back of the cavity. The results agree precisely with what was presented

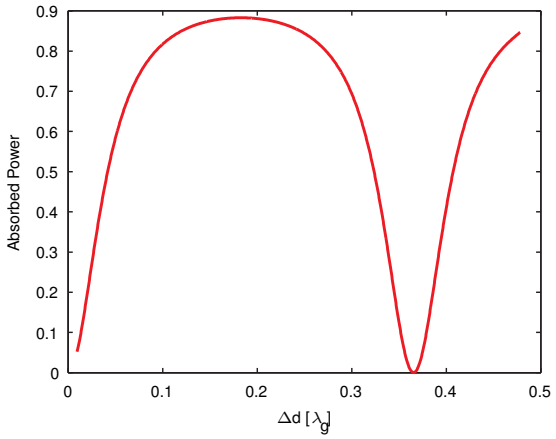

(a) Absorbed power as a function of distance from back of cavity, with $R_s = 377\ \Omega/\text{sq}$

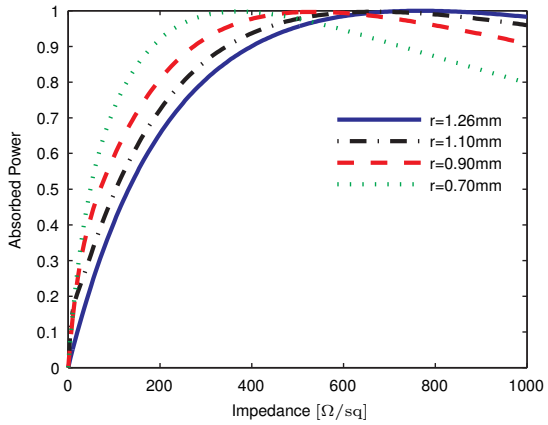

(b) Absorbed power as a function of sheet impedance, 1.95 mm from the back of the cavity

Figure 4.27: Absorbed power in a single-moded cavity as a function of distance to the back of the cavity and sheet impedance, at 80 GHz

in [57], as expected. The response of the absorbed power to the impedance of the absorbing sheet is as expected. The shape of the curve obtained in the case where the absorber placement is varied is as expected, however the value of the separation at which maximum absorption occurs is perhaps not as might be expected. Normally it would perhaps be expected that maximum absorption would occur for a backshort distance of $n\frac{\lambda_g}{4}$, where n is an integer and λ_g is the guide wavelength of the TE_{11} mode. The maximum absorbed power in the range of separations examined occurs at $\approx 0.184\lambda_g$. An examination of the S_{11} matrix shows that the presence of the absorber causes power to be coupled into higher order modes within the $n = 1$ azimuthal order, along the main diagonal of the matrix. This has the effect of altering the guide wavelength of interest, as the field is no longer represented by one mode, so the basis set has effectively changed. Since the guide wavelength of the TE_{11} mode is no longer the 'correct' one in the sense that is being considered, using this value for λ_g results in the indicated result being obtained.

A cost function was implemented for use with the GA, which was a function of one variable, the distance to the back of the cavity, Δd. The overall length of the cavity was fixed at 10 mm, with a radius of 1.26 mm throughout. The cost function receives Δd as its argument and

sets the first section to be of length $10 - \Delta d$, followed by a section of length Δd. Each of these sections is of uniform radius, and so transitioning between the two sections triggers the absorber calculation and then propagates the resulting modes a distance of Δd. The third section is of length 6 mm and radius 0.0000001 mm. At 80 GHz, this section does not support any modes and so effectively simulates closing the back of the cavity. The length of this section is chosen to remove any evanescent modes, even though none should propagate in this section due to the radius. The cost function evaluates the absorbed power and returns the negative of this value, since as outlined earlier the GA locates minima, so minimising the negative value maximises the actual positive value.

Examining figure 4.27a, the absorbed power is seen to vary periodically, and so the maximum absorbed power will always be $\approx$0.9, meaning that the search space for Δd can be restricted to between 0.000 and 0.400 for example, as the local minimum in this range is equal to the global minimum. In a more general case however, the search space for Δd was chosen to be $\Delta d \in [0.000001, 3.000]$, corresponding in terms of the guide wavelength of the TE_{11} mode to $[0.0000002, 0.700]$. The lower bound is expressed with additional decimal places to emphasise that it is close to zero but non-zero, as this situation results in issues when inverting the matrices during the calculation. The optimisation was carried out and returned a result after approximately 8 minutes. The optimised value for Δd, i.e. the distance from the absorber to the back of the cavity resulting in maximum absorption of the TE_{11} mode was found to be 1.911 mm, or 0.183 in terms of the guide wavelength of the TE_{11} mode. This agrees with the data presented in figure 4.27a, showing that the optimiser has obtained the correct solution and so is a reliable method for establishing the optimum position of the absorber. Repeating the process returned a value for Δd of $0.183\lambda_g$ in 8 minutes, showing once again that consistent, reliable results are obtainable with the GA, a critical test owing to the fact that equivalent populations between runs of the optimiser could be different. By making the return value of the cost function directly equal to the absorbed power, the GA was able to minimise the absorbed power. Using the same bounds, the value of Δd giving the minimum absorbed power was found to be of the order of $10^{-8}\lambda_g$, which again agrees with what is observed in figure 4.27, further validating the optimisation process.

Of note here, is that this situation with only one variable is significantly simpler than the case where a feed horn is being optimised. Since there is only one variable, the optimiser can only control one physical process, the one that the parameter is responsible for within the limiting values imposed upon it. There are no interactions with other parameters which could also control other areas of performance. This greatly reduces computational time. Since there are also only three sections to consider, the overall result is an extremely quick optimisation process. Although in this case a simple parameter sweep would achieve the same result, if more free variables were introduced or if the system went multi-moded, then a series of parameter sweeps would become time consuming and difficult to interpret. The optimiser approach would

allow these issues to be dealt with whilst also providing an accurate solution.

This simple optimisation could be carried out by 'brute force' by generating figure 4.1a and reading the appropriate separation value from this directly. It was carried out using the GA in order to verify that the GA would return the correct solution in such a system. Having shown that the optimiser functions correctly, the cavity optimisation process could be combined with the horn optimisation process in order to design a complete pixel, where the cavity is fed directly by the throat of the horn. When the position of maximum absorption is calculated, this means that for this cavity design, the maximum amount possible of any amount of TE_{11} mode that enters it will be absorbed. Using this optimised design for the cavity portion of the pixel, the horn optimisation could be carried out with respect to the metric of interest, with the requirement that the throat radius of the horn match the cavity radius. In this way, the cavity is included in the horn optimisation process as a fixed design that is included in every design that is trialled so the effects of the cavity on field patterns etc. are accounted for in the design process. The horn could for example be optimised with respect to minimising the cross-polar power. As shown, this would result in good field pattern performance and low return loss, giving rise to a high level of absorbed power.

It would of course be possible to optimise the cavity and horn simultaneously in one run, however it was noted during early optimisation work that the relationship between the number of variables and execution time is not linear. Thus, removing as many variables as possible is desirable and so optimising the cavity separately would result in a lower execution time and would guarantee that the cavity design was as efficient as possible as it was designed separately from the horn. Optimising them together could result in a compromise in terms of cavity performance in order to assist in horn performance.

4.6 Conclusions

In this chapter, two global optimisation algorithms that are available for MATLAB as part of the global optimisation package were considered. The standard local optimisers were not considered, due to the fact that to use these it is necessary to know the approximate solution in advance in order to avoid local minima, which defeats the purpose of this work. The two algorithms that were presented were the genetic algorithm and the simulated annealing algorithm. The operation of these algorithms was discussed in detail. Two horn designs were also presented, whose purpose was to act as the basis for the optimisation process, meaning that it would be the parameters which define the profile of these horns which would be adjusted by the optimisation algorithms in order to meet the design goals. The two horns presented were a piecewise conical profile horn and a Gaussian profile horn antenna. These horns were chosen as they, or derivatives thereof, were shown to provide good performance in terms of several metrics which will be important for the horn design presented in the following chapter.

Several other horns exist which could provide the required performance potentially, however a further requirement of this design process was to obtain a design that was relatively easy to manufacture. To this end, these two horn designs were the best option, providing a compromise between ease of manufacture and performance. Since they can also be defined using only a few parameters, they meant that the computational load on the optimiser would be kept as low as possible.

The horn designs and optimisation algorithms were tested subject to three optimisation metrics; co-polar beam symmetry, return loss and maximum cross-polar levels, in order to determine which optimisation algorithm/horn design/optimisation metric gave rise to the best overall horn performance and met all other requirements such as providing results in hours as opposed to days, using standard computing power, and the overall horn design being simpler than a corrugated horn and thus easier to manufacture. The purpose of this investigation was to therefore determine the limitations of the optimisers in terms of horn design and the limitations of the various horn designs/metrics presented in the context of the desired performance requirements.

It was found that optimising the piecewise conical profile horn with respect to the maximum cross-polar level yielded the best results, not just in terms of the metric being optimised but also with respect to the other metrics of interest, whilst also maintaining good bandwidth performance. This is likely due to the fact that low cross-polar levels arise from good performance in the other areas of interest (beam symmetry and return loss), and so forcing the cross-polar levels to be low will necessarily result in high symmetry and low return loss. Optimising these other parameters in isolation will not necessarily give rise to good performance in other areas, as was observed, as they can be forced to be low without any requirement for the other parameters to experience boosted performance. The resulting horn was compared, for context, to an equivalent smooth-walled horn and was found to exhibit superior performance across the band compared to this smooth-walled horn.

This result was obtained using the genetic algorithm. It was found by repeating the optimisation process for the same horn design/optimisation metric/ranges of values for the horn parameters, that results could be obtained that were repeatable on a scale sufficient to give consistent performance between the resulting horn geometries. Using the simulated annealing algorithm in the optimisation process revealed that whilst the genetic algorithm produced results on the order of hours, the simulated annealing approach increased the time to the order of days (5 days minimum), with the process often hanging or crashing, despite using the same horn design as a basis, the same optimisation metrics and the same initial parameter values. This is due to the fact that this algorithm searches much smaller portions of parameter space. When the optimisation process using SA was reduced to one optimisation variable, the results obtained were reasonably consistent with those obtained from the genetic algorithm. Since one aim of this investigation was to obtain results on the order of hours using a standard desktop

computer, the genetic algorithm will be used along with the piecewise conical profile horn for the design work subsequently described in this thesis. This allows a good approximation to the global minimum of the cost function to be found within a reasonable amount of time, whilst still allowing a large search size. The SA would require a reduction in size of the search space which is not desired.

The method for integrating an absorbing sheet within the mode matching technique was also discussed. The optimisation process was applied to an integrating cavity containing an absorber, designed for a single-moded system. The placement of the absorber was optimised for maximum absorption of the TE_{11} mode, and the results obtained agreed with those found by carrying out a parameter sweep of the system for varying absorber position. This shows that if desired, the optimisation process could be applied to a system involving absorbers, as the optimiser returns correct, repeatable results for such systems. It was also outlined how a full horn/cavity pixel could be designed using the optimisation algorithm as the driving process.

Chapter 5

Design and Analysis of a Smooth Walled Horn for Cosmic Microwave Background Measurements

In this chapter, the design and analysis of a profiled smooth walled horn for cosmic microwave background radiation measurements will be presented, using the optimisation process detailed in chapter 4, based on the genetic algorithm approach. The horn that resulted from this design process was constructed by Rutherford Appleton Laboratories (RAL) in the United Kingdom, and was tested experimentally using the vector network analyser (VNA) testbed at NUI Maynooth. The results obtained from the measurement campaign will be presented and compared to the simulated results.

5.1 Requirements

Several performance requirements were specified for this horn based on the performance requirements outlined that were outlined in chapter 1 for a potential future B-mode CMB mission. They are as follows.

- Operate in the 85-115 GHz range (W-Band)
- Single moded across the band
- Cross-polar and return loss levels of below -30 dB
- Symmetric co-polar beam patterns with low sidelobe levels
- Optimisation process must take reasonable amount of time using standard desktop computing power using an arbitrary starting point and no design 'tricks', in order to be robust

- Design must be relatively easy to manufacture and carry lower thermal/mechanical penalties than an equivalent corrugated horn

Clearly, these requirements are highly demanding and difficult to achieve using a smooth walled horn design, particularly when required across a 30% bandwidth. The bandwidth issue is further complicated by the need for the optimisation to be carried out at a spot frequency in order to minimise the computational load, as required by the design goals. These issues arise due to the fact that an increase in performance for one horn characteristic inevitably comes at the expense of others, an issue which is of far less significance when corrugations are used. This necessitates striking a balance between the different metrics at a single frequency, which is achieved by optimising with respect to a metric that results in other metrics improving. The issue arising from the spot frequency optimisation is that smooth walled designs arising from such design processes tend to rely heavily on the interaction of the dimensions of the horn with a specific wavelength, and so bandwidth performance is difficult to achieve. Any deviation from the design wavelength would be expected to result in a severe degradation in horn performance. Corrugated horns avoid this by utilising the balanced hybrid mode condition, as outlined in [37, 45].

With the above in mind, the design process is now considered.

5.2 Horn Design

In this section, the design of the horn outlined above will be considered. Following significant investigation, as reported in chapter 4, it was observed that in order to meet the general performance requirements for a CMB mission, the best compliance was achieved when the genetic algorithm was applied to a smooth-walled piecewise conical profile horn, with the aim of the optimisation process being to minimise the cross-polar levels. This chapter therefore applies the genetic algorithm to a piecewise conical profile horn optimised with respect to minimum cross-polar levels, using different constraints on the overall length and aperture radius compared to the constraints used in chapter 4 (see table 5.1), in order to obtain a design for a horn that complies with the performance requirements outlined above. The results of the optimisation process when applied to these specific requirements and an analysis of the resulting design will now be presented.

5.2.1 Optimisation Results

As before, a few sectioned spline profile horn was used, employing six sections. Following extensive testing of various performance metrics and initial bounds, the metric that gave rise to the best result in terms of all desired parameters was the maximum level obtained in the cross-polar radiation pattern of the horn. Using other metrics resulted in good performance for that

parameter at the expense of other parameters, but using this characteristic gave rise to a good compromise between all parameters of interest and relatively good bandwidth performance, as will be shown. Optimisation with respect to several simultaneous horn characteristics was also attempted, however it either took too long to justify the small increase in overall performance (keeping in mind the computational time requirement), or ultimately still optimised one parameter at the expense of all others.

When minimising the cross-polar radiation pattern, the $\phi = \frac{\pi}{4}$ cut was considered as opposed to a full 2D plane. Due to the mode sets used in calculating this field pattern (in this case the y-components of the TE and TM modes without including the orthogonal polarisations), the maximum of the radiation pattern occurs along the $\frac{\pi}{4}$ cut and so there is no need to test a full 2D plane. Several standard horn geometries were tested in order to validate this, by calculating the complete 2D cross-polar radiation pattern in the farfield, and it was noted that the maximum level always occurred along the $\frac{\pi}{4}$ cut. Therefore, if the $\frac{\pi}{4}$ cut is calculated and the maximum level along this direction is found, then the maximum cross-polar level achieved by the horn is known. This has the advantage of reducing the time required to calculate the optimisation metric, as only one cut is calculated as opposed to having to finely sample all azimuthal angles from $\phi = 0$ to 2π in a 2D plane.

Using GA terminology, each individual (horn geometry under test) that is evaluated by the algorithm is defined with respect to the length of each conic section (each section being this length other than the throat which is fixed at 5 mm to further reduce the computational load. 5 mm was chosen as this is long enough to suppress evanescent modes) in addition to the radius of each section, where the radius of each section is an independent, bounded variable. Table 5.1 shows the initial upper and lower bounds of each of the optimisation parameters, as well as their optimised values which are the output values of the design process. The horn was designed for

Parameter	**Lower bound /mm**	**Upper bound /mm**	**Optimised value /mm**
Section length	8.000	10.000	8.781
r1	1.200	2.000	1.604
r2	1.600	4.000	3.417
r3	3.400	6.000	3.528
r4	5.500	8.000	5.694
r5	7.500	10.000	7.837
r6	9.500	13.000	12.969

Table 5.1: Lower and upper bounds for the generation of the initial population, with the final optimised values.

use over the 85-115 GHz band, corresponding to a 30% bandwidth centred on 100 GHz. This band was chosen as it is a common frequency range to operate in when probing the cosmic origins, however as stated previously, despite the specified bandwidth the optimisation was

carried out at a spot frequency of 100 GHz in order to reduce the computational load and to see how well a horn designed in such a manner would perform across the band. A 30% band is ambitious for a horn designed at a spot frequency. If the horn proves to operate well over a smaller bandwidth of 15-20%, then it will prove extremely useful and relevant in future CMB and Far-IR missions. Such bandwidth are common in these frequency ranges in instruments designed to, amongst other goals, observe the CMB such as the **C**osmic **O**rigins **E**xplorer, COrE, as described in section 1.4. Many of the bands used in COrE are 15 GHz wide, which in the 100 GHz range is approximately a 15% bandwidth. The indicated 85-115 GHz band for the horn being designed in this thesis is typical of a channel that could be used in such missions, hence its selection.

Since the horn was being designed with a view to being fed by a section of WR10 waveguide (dimensions 1.27 $\times$ 2.54 mm), consideration was given to the modes that would be present at the input. At this frequency, the allowed values for the radius of the throat mean that modes of azimuthal orders 0,1 and 2 would potentially be supported. The horn was designed as single-moded in the $n = 1$ azimuthal order, but based on previous investigative simulations, it was deemed necessary to allow the horn to be this wide in order to maximise power throughput. Previous simulations had displayed low power at the output when smaller throat radii were used. It was possible to exploit this fact in this situation, as although modes of three azimuthal orders are supported by these potential throat radii, once the horn is attached to the rectangular WR10 waveguide that feeds the system, it is valid to consider only the $n = 1$ azimuthal order. Due to the relative symmetries between the rectangular and circular mode sets, no power is coupled into the $n = 0$ azimuthal order in the horn as modes of this order in the rectangular WR10 waveguide will either couple into the circular TE_{11} and TM_{11} modes in the horn, or not couple into the horn at all. Thus, $n = 0$ need not be considered in the horn. Similarly, rectangular modes with $n = 2$ will either couple into circular modes with $n = 1$, or not couple into the horn at all, meaning that at all times only $n = 1$ need be considered, as these modes are neither excited by the rectangular waveguide source nor excited by scattering from $n = 1$ modes within the horn, as modes in different azimuthal are incoherent. The assumption based on the transition section does not hold for $n > 2$, however modes of such azimuthal orders are not supported by the horn, and so this is not an issue. Of the modes present in the horn, the TE_{11} mode is dominant in the throat, and so the horn is essentially single-moded as required, and so only the $n = 1$ azimuthal order need be considered in the optimisation process. It is stressed that this assumption is only valid when the horn is fed by a WR10 waveguide, however this is assumed to always be the case.

In reality, the transition between mode sets will occur via a rectangular to circular transition, as will be considered in section 5.3.2. The primary function of the transition section is to match the two different mode sets. Without this transition, the abrupt change at the junction from rectangular to circular would result in a large mismatch between the mode sets and so give rise

to a significant amount of power being reflected at the interface between the two geometries due to coupling into higher order modes. Adding a transition section reduces such reflections and reduces the detrimental effect of the geometry transition on the performance of the horn. As SCATTER cannot currently combine rectangular and circular geometries, this transition was not included in the design process.

The bounding values for the other parameters were determined by setting reasonable limits on the physical dimensions of the horn. The ultimate application of the designed horn was as a feed for one pixel as part of a focal plane array unit. Although not designed for a specific project and so there were no dimensions available for the array, it made sense to restrict the dimensions of the horn in order for it to remain a 'traditional' size. This is important, as if no restrictions were placed on the length for example, then the horn could become extremely long in order to phase up the propagating modes to minimise the cross-polar power. If the aperture radius were not limited then it would grow extremely large in order to minimise the impedance mismatch between the modes at the aperture and free space, as this would also reduce cross-polar power. To this end, the horn was limited to be of length $\approx 19\lambda$ (a reasonable length) and of aperture radius $\approx 4\lambda$. The bounds of the individual sections were defined such that the lower bound of one section lay within the bounding values of the previous section. This gave the algorithm the ability to explore geometries where the horn profile was not monotonically increasing in radius, and so a section could have a smaller radius than the previous section.

The design process giving rise to the optimised values quoted in table 5.1 took 6.5 hours, a very reasonable amount of time considering the number of variables present and the complexity of the cost function. The overall length of the resulting horn was $\approx$ 49mm, with an aperture radius of 12.969 mm. These dimensions correspond to 16.3λ and 4.3λ respectively for $\lambda = 3$ mm, which is reasonable. A schematic of the horn is given in figure 5.1, (a). Considering a total

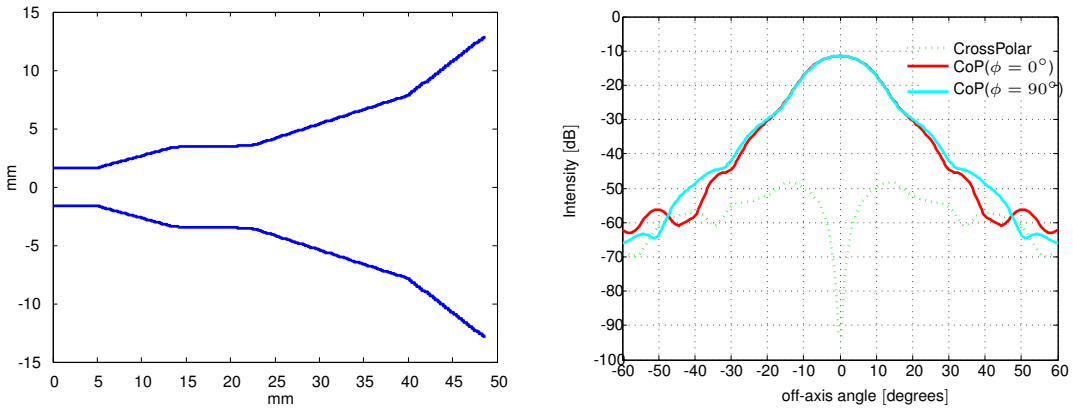

(a) Schematic of optimised horn

(b) Co and cross polar radiation patterns of the optimised horn

Figure 5.1: Horn schematic and farfield radiation patterns

power of unity in the $n = 1$ azimuthal order, 0.994 of the power was transmitted, with the

majority of the power being contained in the TE_{11} and TM_{11} modes, with some power also found in the TE_{12} mode. The remaining modes contained a rapidly decaying amount of power. When a singular value decomposition was carried out on the S_{21} matrix, it was found that only one mode was supported and so the horn is single-moded, as required. The balance of the power was found to be in the S_{11} matrix, contained solely in the TE_{11} mode (in reflection), giving rise to a total reflected power of -32 dB, as required. This is summarised in table 5.2, which shows the power (with respect to unity excitation) in each mode at the exit aperture of the horn that contains non-negligible power. This corresponds to the first column of the S_{21} matrix, implying that these modes are excited by the TE_{11} mode scattering into higher order modes within the horn. The total reflected power across the band (integrated S_{11}) is

$\mathbf{TE}_{nl}$	**Power [arb. units]**	$\mathbf{TM}_{nl}$	**Power [arb. units]**
1 1	0.3890	1 1	0.3160
1 2	0.1659	1 2	0.0672
1 3	0.0292	1 3	0.0185
1 4	0.0091	1 4	0.0019
1 5	0.0013	1 5	0.0005
1 6	0.0004	1 6	0.0001
1 7	0.0000	1 7	0.0001
1 8	0.0000	1 8	0.0000
1 9	0.0000	1 9	0.0000
1 10	0.0000	1 10	0.0000

Table 5.2: Transmitted power in each of the modes at the aperture of the optimised horn geometry

shown in figure 5.2. The requirement is to have the reflected power below -30 dB (a figure originating from the difference in magnitudes of the E and the B-modes) across the band, and this requirement is clearly met. The reflections are highest at the optimisation frequency of 100 GHz. This illustrates the trade off in performance for parameters other than the one being optimised. The cross-polar power is very low at this frequency for example, but this comes at the expense of a higher amount of reflected power within the horn. Despite this, the levels still meets the requirement and so this is not an issue. The trade off still allows the requirement to be met.

The ratio of the maximum cross-polar level to the maximum co-polar level, the former quantity being the parameter that is being minimised, is -36.86 dB, meeting the design requirements. This means that the separation between the co and cross-polar fields is large, as required by the polarisation sensitive nature of the measurements. A visual examination of the $\phi = 0^\circ$ and $\phi = 90^\circ$ co-polar fields that are shown in figure 5.1, shows that the co-polar field is highly symmetric, again as required. A significant mismatch of these fields would result in the horn being highly sensitive to azimuthal orientation, which is not a desirable feature for measurement sys-

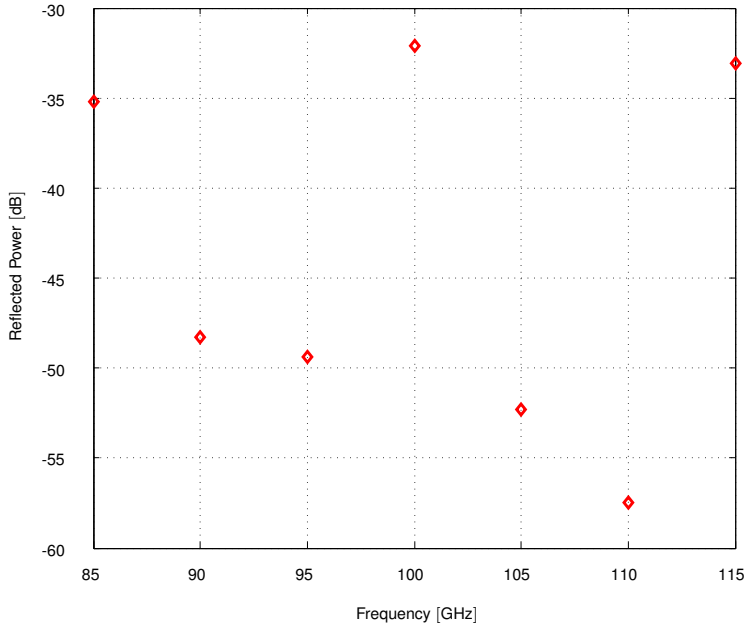

Figure 5.2: S_{11} (total reflected power) of the optimised horn across the 85-115 GHz band

tems such as those in question. It is clear that minimising the cross-polar power also enhances other parameters, as low sidelobe levels are noted, and good performance is noted across the band, despite the horn design only being optimised at one frequency; 100 GHz. These points will be clarified in section 5.3, by means of a comparison with a comprehensive measurement campaign. These results are not surprising, as low cross-polar levels follow from symmetric beams with low sidelobe levels, low reflection etc. This performance was not noted when other optimisation metrics were used, as detailed in chapter 4. In these cases, the other parameters not being optimised would suffer massively to accommodate the optimised parameter, or the bandwidth performance would be significantly compromised.

Figure 5.3 shows the E and H planes of the horn across the 85-115 GHz band. Considering the E-cut (azimuthal angle of 0°), it can be seen that as the frequency increases, the off-axis extent of the field increases meaning that the field is no longer as spatially confined or as focussed, once this additional mode appears. The field distribution remains sufficiently constant across the band, showing only slight variation at lower power levels. Just below 115 GHz there is a significant change in the radiation pattern, corresponding to an additional modes cutting in at around this frequency, namely the TM_{11} mode. This additional mode is expected, as again, the horn was optimised for 100 GHz only and so there is no consideration given to modal content at other frequencies. Optimising for additional frequencies would perhaps suppress this additional mode, however since it only cuts in at the edge of the band its effect is negligible and is not overly detrimental. The H-cut (azimuthal angle of 90°) exhibits similar features, despite being slightly narrower in spatial extent, and so this indicates a large degree of symmetry across

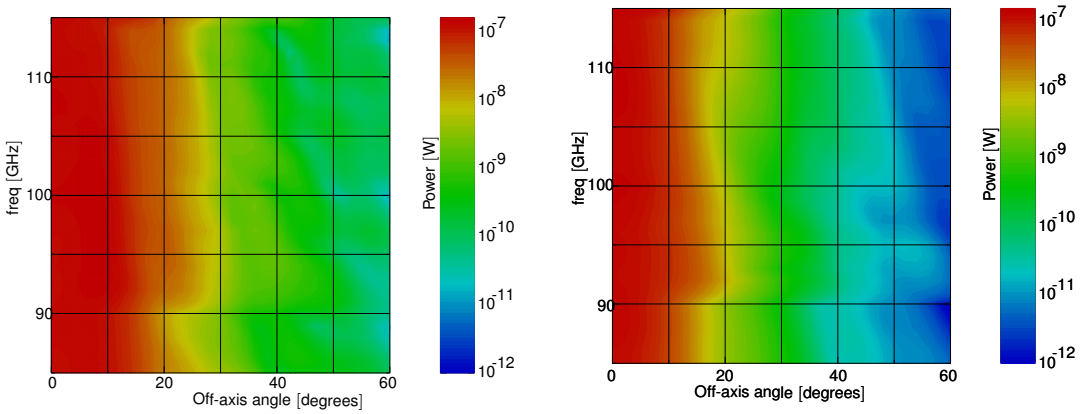

(a) E-cut (b) H-cut

Figure 5.3: Bandwidth analysis of E and H field cuts for optimised horn

the band, as required (the precise degree of symmetry will be inferred from the experimental measurements later in the chapter). This deteriorates towards the edges of the band, again due to the spot frequency nature of the optimisation, and a bandwidth optimisation (subject to the availability of sufficient computational power) would extend the symmetry across the band, although it achieves a very acceptable level as is. The principal planes at either end of the band are shown in figure 5.4.

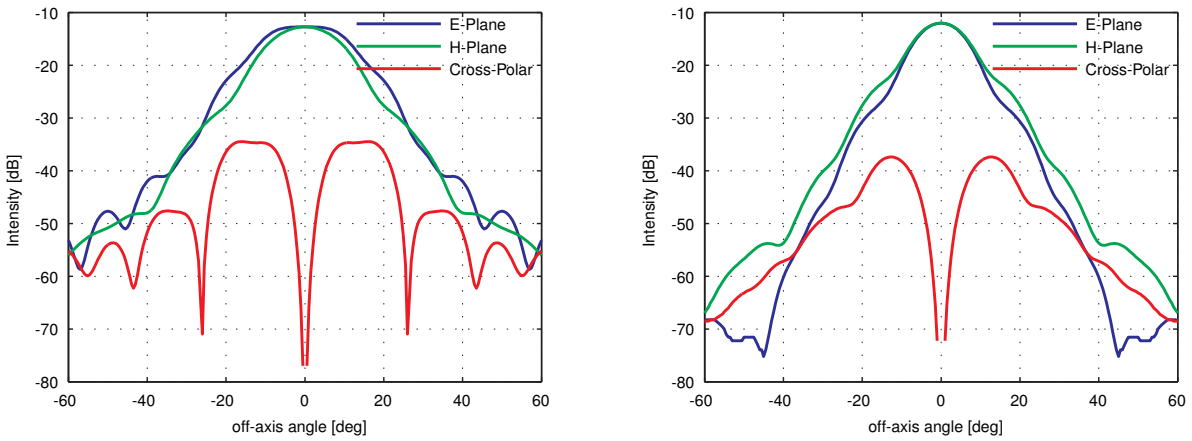

(a) 85 GHz (b) 115 GHz

Figure 5.4: Principal planes of the optimised horn at 85 and 115 GHz. The cross-polar plane refers to the $\frac{\pi}{4}$ cut

Considering the excellent performance of the horn across the band despite the spot frequency optimisation, a bandwidth optimisation was not undertaken, as the performance goals were well achieved using the minimum amount of computational effort that is associated with a spot frequency optimisation, which was another goal of the project. Bandwidth optimisations have been run, and unusable results have been recorded despite execution times of approximately

a week. In each case, regardless of the bounds and the range involved, the final values would simply converge to the lower bound value in the case of each parameter.

The likely reason that the bandwidth optimisation process tends to the lower bound of each parameter is the fact that the multiple frequencies results in an extremely complicated mathematical form of the cost function, and so it is difficult for the algorithm in its current form to minimise the cross-polar power for all values of frequency. The algorithm functions by calculating the maximum cross-polar value at each of the frequencies and minimising the average. It appears that if a new geometry improved the performance of the horn at one frequency, it would be at the expense of another frequency and so on, resulting in an average that stalled at a similar value regardless of the geometries that were tested. This means that the only way the algorithm could iterate towards reducing the average in a meaningful manner was to reduce the geometry of the horn to its smallest value, choking the number of modes supported at the throat at each frequency and so reducing the power entering and propagating through the horn. This would of course reduce the cross-polar level (by minimising the power arriving at the aperture) and so achieve the goal of the optimisation algorithm, but at the expense of any useful power arriving at the exit aperture of the horn. This behaviour was also observed for other complicated cost functions, for examples cost functions that are evaluated with respect to several metrics.

Figure 5.5 shows the bandwidth performance of the cross/co-polar ratio parameter of the optimised horn compared to that of a standard corrugated horn. This comparison is instructive, as the aim is to achieve comparable performance to a corrugated horn in this area. Considering

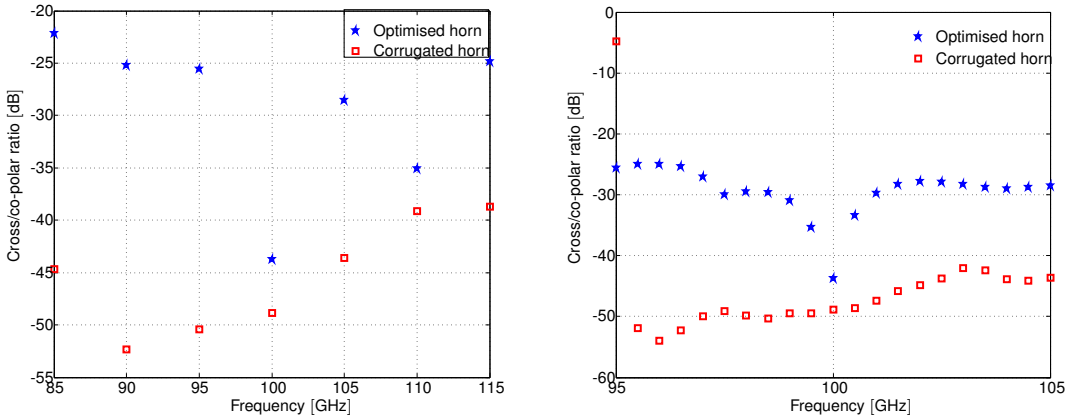

(a) 30 % bandwidth

(b) 10 % bandwidth

Figure 5.5: Bandwidth analysis of cross/co-polar ratio for optimised and standard corrugated horn, with only azimuthal order $n = 1$ considered

the 30% bandwidth analysis, it is clear that the performance of a standard corrugated horn exceeds that of the optimised horn across the band. This is to be expected. The design and manufacture of a corrugated horn is highly complex and costly. The optimised design

is significantly simpler, yet at the frequency at which it was designed, the performance of the optimisation metric (cross/co-polar ratio) is comparable to that of the corrugated horn. The fields are highly symmetric as previously discussed (and as will be shown in section 5.3), with a low return loss. Thus, the optimised horn meets the criteria described at the design frequency, rivalling the performance of the corrugated horn. The difference between designs becomes apparent however, when the performance of the horns across the band is considered. The corrugated horn maintains low ratio levels, however the optimised horn does not. Despite this, between 100 and 110 GHz the low cross/co-polar levels are maintained by the optimised horn. At worst, the ratio strays 5 dB above the required maximum, however this occurs at the edge of the band. The bandwidths encountered in missions such as COrE are, as stated earlier, narrower. Applying the 15 GHz bandwidth utilised by COrE in the 100 GHz region, it is clear that the performance of the horn in this band is very good considering the spot frequency optimisation.

Examining the 10 % bandwidth, a likely channel for future missions (such as proposed for the COrE mission), it is clear that the optimised horn is comparable to the corrugated horn and certainly achieves the desired cross-polar performance, with extremely low levels being achieved in addition to other parameters meeting acceptable standards, however at times the requirements are breached. It is not feasible that such a simple geometry, designed at only one frequency, could compete across the band with a corrugated horn with respect to all performance characteristics, however given the contrast in complexities, the results are encouraging. It is very difficult to replicate the performance of a corrugated horn with such a simple design, even using horns that were designed using a highly mathematical design process, however at the design frequency this is achieved.

In the requirements set out for the horn antenna design, there were no specific constraints on either the overall length or aperture radius of, as there were no specifications available for the focal plane unit or the optical system (hence Gaussicity and associated parameters would be pointless to calculate as there are no matching optics currently detailed). To this end, as outlined earlier, the horn is of a reasonable length and aperture radius, and if necessary these parameters could be constrained in order to meet the physical requirements of any system. The effect of each conic section on the overall horn performance was therefore investigated, to see if sections could be removed with the result of shortening the horn and making it narrower. Removing the final section resulted in the maximum levels achieved by the co-polar field reaching -6.71 dB, with the cross/co-polar ratio rising slightly to -34.29 dB, however the field patterns exhibited more pronounced sidelobes with increased levels. Some symmetry was also lost between the principal planes of the co-polar beam. Removing subsequent sections further increased the maximum co-polar levels to as high as -5.72 dB with one section, however the cross/co-polar ratio was seen to increase rapidly, eventually reaching -19.44 dB with only one section beyond the throat. The field patterns were also observed to deteriorate in

terms of sidelobe structure and symmetry. The performance of the optimised horn is therefore dependant on the entire horn. The first few sections excite and phase the waveguide modes so as to reduce the total cross-polar power and associated parameters, with the final section allowing the modes to phase up so and a highly symmetric co-polar field with low sidelobe levels results. Using differing numbers of spline points was also trialled, however changing this number resulted in a similar deterioration in horn performance. This was due to increased inter-modal scattering due to more radius changes being present in the increased case, and more dramatic radius changes occurring in the case where less spline points were used. If a shorter and/or narrower horn was required, it would therefore be necessary to repeat the entire optimisation process again, subject to the new constraints.

It is necessary to provide context for the performance of the designed horn. As stated in chapter 4, the algorithm does not have a fixed starting point that it uses as an initial point, the design of which it then optimises. Rather, it takes a pool of geometries that are derived from a range of values for each variable that span the entire parameter space. This corresponds to an entire set of horn designs. It then uses an iterative process to combine these horns in various ways to narrow the search space to find the optimal solution. It does not optimise a particular starting horn. As this horn was designed to provide an alternative to corrugated horns in future many pixel arrays, it therefore makes sense to provide context not by comparing it to an initial design but rather to the performance of an actual corrugated horn operating at 100 GHz. The horn that will be used is one of the corrugated horns at NUI Maynooth, with entrance and exit aperture radii of 1.535 mm and 7.1345 mm respectively, with an overall length of 69.9 mm. To make a valid comparison, it was also assumed that the corrugated horn was illuminated via a WR10 waveguide, and so only modes of azimuthal order $n = 1$ were considered.

Figure 5.6 shows the principal planes of the corrugated horn at 85, 100 and 115 GHz as well as the co-polar symmetry across the band. With reference to figures 5.1b and 5.4, the relative performance of the horns can be compared. In terms of beam quality, the optimised horn required a high level of co-polar symmetry and low sidelobe levels. For the corrugated horn, shouldering effects/sidelobes begin to appear at approximately -30 dB at the various points sampled across the band. Examining the plots for the optimised horn reveal that it is required to drop to a similar power level before any significant amount of these effects are noticed. This is a feature across the band. A high degree of co-polar symmetry is also evident across the band in either case, fulfilling both requirements for the quality of the main (co-polar) beam of the optimised horn, when compared to the performance of the corrugated horn which was the goal. Figure 5.6d shows the return loss of the corrugated horn across the band, which should be compared to figure 5.2, showing the same parameter for the optimised horn. The maximum level of return loss experienced by the optimised horn is less than that of the corrugated horn, (≈-30 dB and -26 dB respectively). The corrugated and optimised horns both achieve extremely low levels of return loss across the band, meeting the mission requirement of

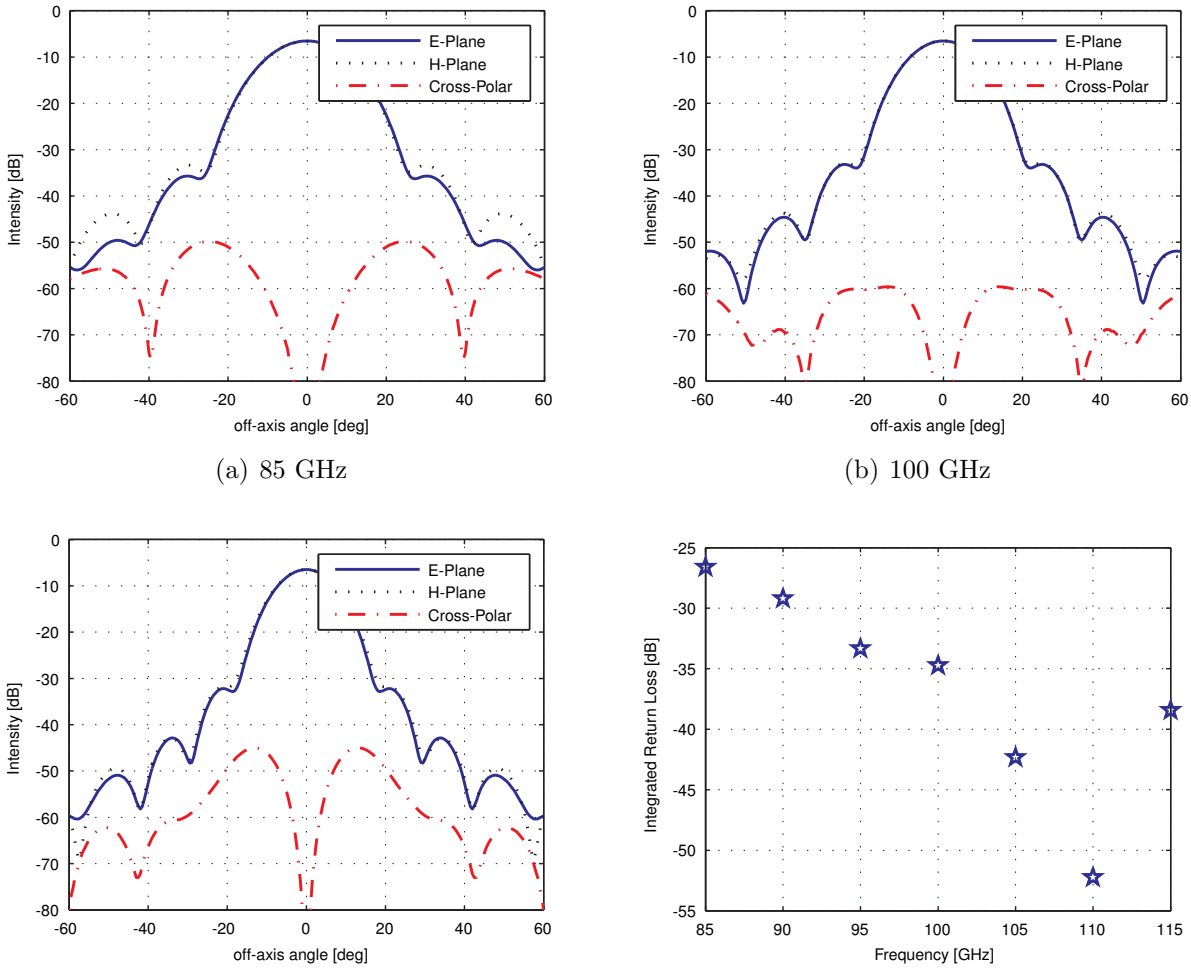

(a) 85 GHz

(b) 100 GHz

(c) 115 GHz

(d) Return loss

Figure 5.6: Cuts and return loss for the NUI Maynooth corrugated horn

maximum levels of -30 dB, with the optimised horn performing slightly better in this regard. The performance of the horns in terms of return loss is thus comparable. Figure 5.5 already gave a comparison between the two horns with respect to the ratio of the maximum cross and co-polar levels. This parameter illustrates the superior performance of the corrugated horn, particularly away from the design frequency. As expected, the corrugated horn realises better performance, however the optimised horn still closely approaches the -30 dB requirement, so although the performance levels are not as close as was observed for the other parameters, the optimised horn still realises the required performance. The two horns provide competitive levels of performance in terms of the parameters of interest. This is realised across the band, which is an impressive result as the optimised horn is smooth-walled and was designed at a spot frequency. The optimised horn design is therefore a success when compared with the corrugated horn, as it delivers comparable performance in the required areas.

As previously alluded to, the horn was designed for experimental verification in systems that use a rectangular waveguide to feed components on a Vector Network Analyser (VNA) setup. A rectangular to circular transition must therefore be incorporated in order to match the two geometries, minimising reflections and the associated return losses at the interface. This could not be carried out using the mode matching technique, as although both circular and rectangular waveguide structures are supported in this method, it is not yet possible to combine both geometries without further verification. It was necessary to design this transition in order for the horn to be manufactured, and so the design process and the transition section will be discussed in more detail in section 5.3.2, where the measurement campaign is considered.

Since this horn geometry satisfies all the design requirements and performs relatively well across the band, this was taken to be the final design. Prior to selecting this as the final design, the impact of altering the variable ranges and the optimisation metrics was examined, to illustrate the impact of this on the performance of the horn relative to what became the accepted design. This analysis will now ne presented.

5.2.2 Impact of Altering Optimisation Parameters

In this section, the impact of different ranges for the variables and different optimisation metrics will be briefly examined purely for illustrative purposes. The optimisation was repeated with the same cost function as used above, returning a minimum value for the maximum level realised in the cross-polar field, however this time although the mid-point of the search space for each section remained the same as in the previous optimisation, the overlap between the search spaces of consecutive variables was increased. This gave the algorithm more scope to search non-standard designs. The resulting optimisation took 27,465 seconds, and yielded a horn geometry (shown in figure 5.9) with a maximum co-polar level of -14.6 dB, a maximum cross-polar level of -49 dB (giving a ratio of -35.3 dB), a Gaussicity of 0.9996 and an E/H coupling of 0.9996. The cross-polar beam, as shown in figure 5.7, displayed increased levels

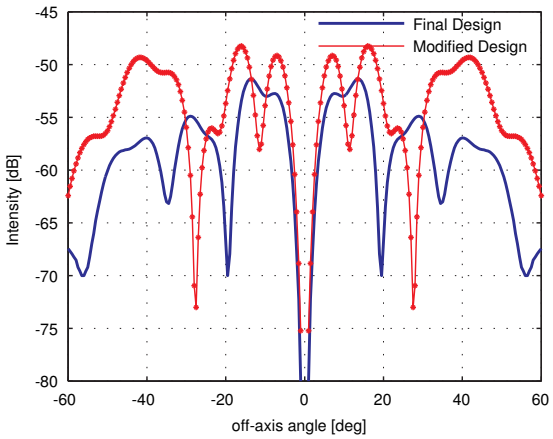

Figure 5.7: Cross-polar radiation pattern for the final and modified designs at 100 GHz

overall as evidenced by the maximum level increasing, however the beam pattern remained sensible. The co-polar cuts, as shown in figure 5.8, were less directional than in the final design and showed a reduction in on-axis power as a consequence of this. The reduction in co-polar

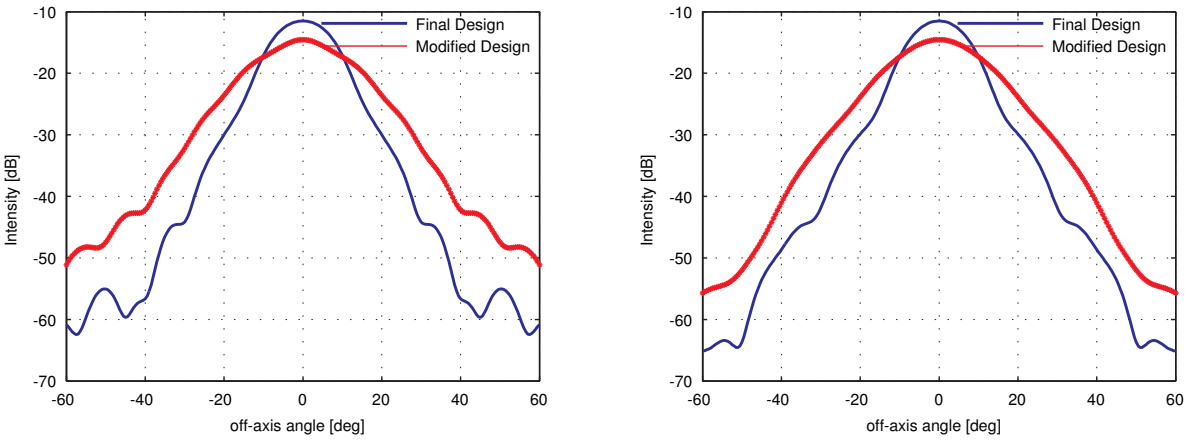

(a) E-cut

(b) H-cut

Figure 5.8: Co-polar radiation patterns for the final and modified designs at 100 GHz

power combined with the increase in cross-polar power results in a reduction of the cross/co-polar ratio, and so this is a poorer horn design in terms of one of the main requirements of the design, namely to retain good discrimination between the co-polar and cross-polar fields. The modified horn geometry would also be much more difficult to manufacture due to the mixture of sections involving increasing and decreasing radii, as shown in figure 5.9. The less directional main beams coupled with the loss of cross/co-polar ratio performance and a much more awkward horn design render this horn inferior to the previous design.

An optimisation in which the coupling of the average co-polar field to a Gaussian of fixed width was the metric of interest was carried out in order to see if it was possible to boost performance. This corresponds physically to matching the co-polar radiation pattern to some

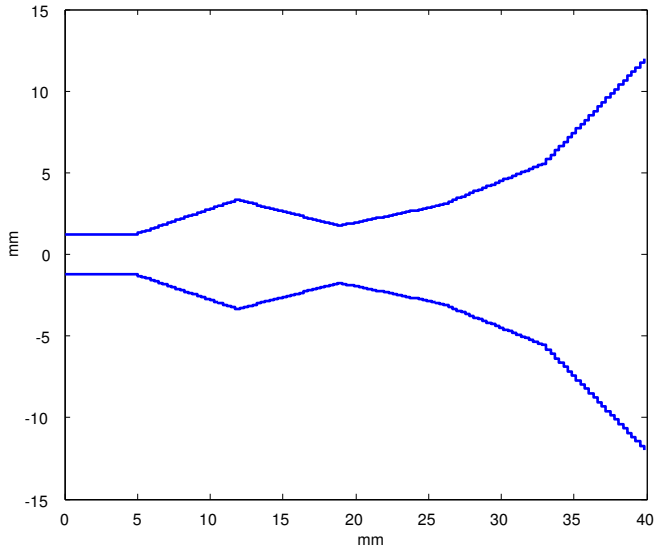

Figure 5.9: Profile of the modified horn geometry

fixed radiation distribution, perhaps generated by some focussing optics where the output of this system is known. In this case the optics have not been decided, however for completion this parameter was briefly investigated. This design yielded a cross/co-polar ratio of -33.24 dB which is worse than the final design, with a maximum co-polar level of -7.39 dB. The E/H coupling remained approximately the same at 0.999, with a Gaussicity of 0.98785. This latter characteristic was the characteristic that was optimised. Of note is that the final section of the horn also did not tend toward its maximum permitted value at the aperture, but rather tapered to a radius of approximately 10 mm. Other than this, it was found to have effectively the same profile as the horn presented in figure 5.1. The difference due to this final section is large however. There is an introduction of significant amounts of higher order modes, and this additional modal content manifests itself clearly in the co-polar fields with significant shouldering effects and higher sidelobe levels. This renders this horn inferior in terms of having a clean beam, despite the higher transmitted power, and so the best horn is still that presented in figure 5.1.

To test the effects of simultaneously utilising several different optimisation metrics, the same initial conditions that gave rise to the accepted horn design were used. The metrics selected were the Gaussicity as above, and also the maximum level in the cross-polar field as both of these parameters previously yielded good results. The resulting horn design showed no marked improvement relative to the accepted design in terms of the horn characteristics, or indeed the field patterns, as shown in figure 5.10. There is no deterioration in performance, and so this horn is a viable design. It is clear that in fact the cross-polar levels associated with this design are slightly lower, however this design was disregarded as the optimisation process took 610119 seconds ($\approx$ 169 hours), which is of the order of a week, as opposed to the six hours that the cross-polar level only optimisation took. Thus a significantly longer time was required

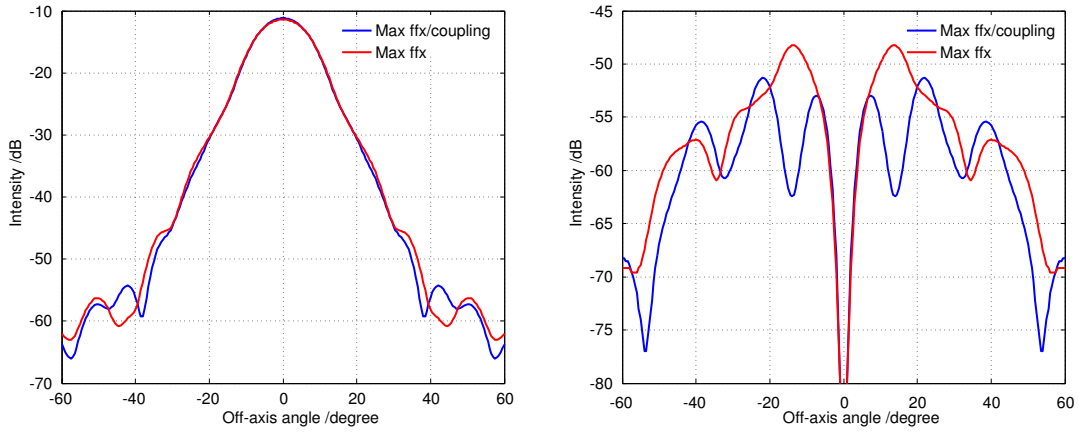

(a) Co-Polar E-plane

(b) Cross-Polar

Figure 5.10: Co-Polar (E-plane) and cross-polar cuts for the final design where only the maximum cross-polar (ffx) level was optimised and where both this and the Gaussian coupling were optimised

for very little improvement, and seeing as one aim of this design process is to produce a design using a standard computer on the order of hours, this means that the original optimised design is still superior in terms of all of the goals of the design process (computational time etc.).

Finally, it was desired to see if similar performance could be achieved using fewer or more spline sections than the six that were used to obtain the accepted design. For the case where fewer sections were used, four sections were chosen as this would give the system a sufficient number of degrees of freedom to achieve a reasonable design. Nine sections were chosen for the case where more sections were used, as this balances an increased number of sections with maintaining a reasonable computational load. The ranges allowed for each of the variables in either case were defined so as the throat and aperture radii were the same as in the six section case, and the section length variable was adjusted so as to give equivalent lengths to the six section horn. The same parameter was optimised as in the six section case; the maximum level in the cross-polar field. A comparison of the profiles of these two horn designs, along with the six section horn, is shown in figure . The results for the optimised parameter, along with several others, for each of the horn designs are shown in table 5.3, with the co-polar and cross-polar fields of the three horns compared in figure 5.12. (Ratio refers to the ratio of the maximum value of the cross-polar field to the maximum value of the co-polar field.) It can be seen clearly that the six section horn gives the smoothest co-polar beam with higher directionality, less shouldering and lower sidelobe levels, while the cross-polar beam pattern is significantly lower for this horn, and more regular. When the various characteristics shown in table 5.3 are examined, it is clear that all three horns perform similarly, aside from the maximum cross-polar levels achieved, and therefore the cross/co-polar ratio. This is a natural consequence of the cross-polar beams indicated in figure 5.12. Since, due to the beam patterns, the six section horn performs better on balance, it can be concluded that the six section spline profile horn

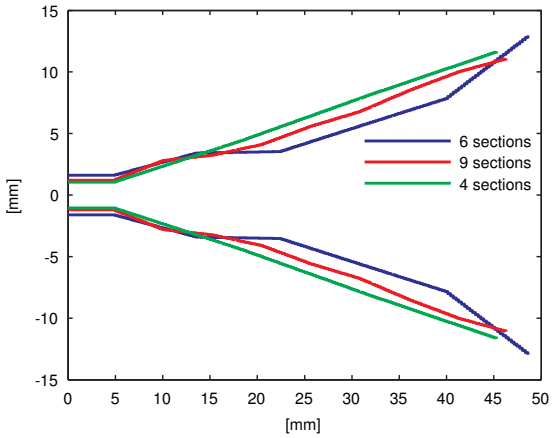

Figure 5.11: Profiles of the spline horns defined using a different number of sections

Parameter	Original (6)	9 sections	4 sections
max co-polar [dB]	-9.38	-9.2	-8.58
max cross-polar [dB]	-48	-38.08	-27.68
ratio [dB]	-36.8576	-28.7	-19.1
Gaussicity	0.9986	0.998	0.992
E/H coupling	0.999	0.999	0.994

Table 5.3: Values for various horn metrics for the 6,9 and 4 spline section designs

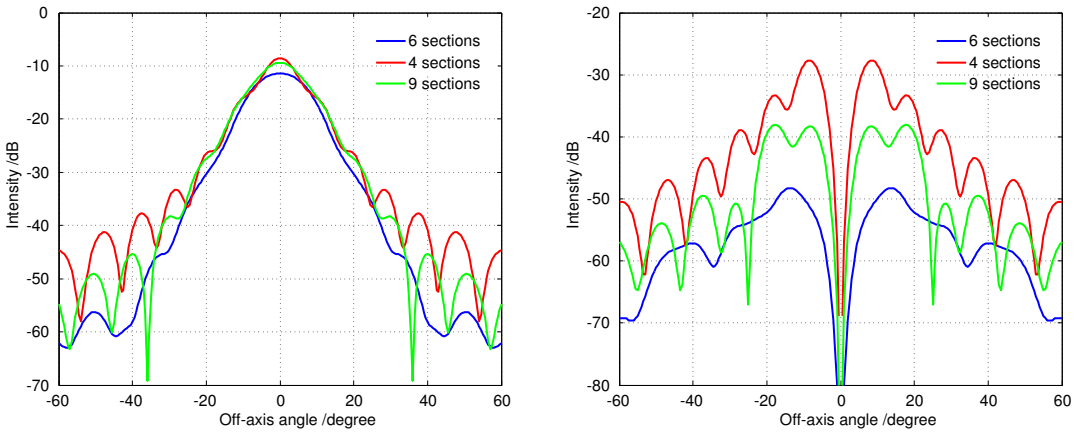

(a) Co-Polar (E)

(b) Cross-Polar

Figure 5.12: Co-Polar (E plane) and cross-polar cuts for the 6, 4 and 9 spline section horns where only the maximum cross-polar (ffx) level was optimised

whose design was optimised with respect to the maximum cross-polar levels only, with some overlap in the search space, gives the best overall performance. This is likely due to an increase in scattering between modes for the other designs. In each case there is more scope for such scattering due to a more 'jagged' design in the 4 section case, and simply more sections in the 9 section case.

It would perhaps be expected that the nine section horn would iterate towards a solution that was similar to the six section horn, particularly given that the physical limits imposed on the parameters at the throat and exit aperture were the same, and that the intermediate limits were adjusted to reflect the differing number of sections. The optimisation process was run several times in the case of the nine section horn, and a similar solution, within the limits described previously, was found each time. There are several possible reasons for this. One possibility is that the bounds placed on the intermediate variables may not easily permit convergence between the two models. More generous overlaps were allowed between the sections, however this resulted in the horn attempting to minimise its dimensions, as previously noted, in order to choke all power from the horn and consequently reduce the cross-polar power. Therefore, the bounds that resulted in the illustrated horn were necessary in order to maintain horn performance and so the overlap between sections could not be increased so convergence between the models was not possible. Another potential source of this apparent discrepancy is that the cost function may be mathematically different when a different number of sections are used, and so the solution space may adopt a different shape. This would result in the possibility that results may converge for a different number of sections, however it would not be a necessary outcome.

5.3 Horn Measurements

The piecewise conical profile horn presented in section was manufactured by Rutherford Appleton Laboratories (RAL) located in Oxford in the United Kingdom. Two versions of the horn were manufactured, with the difference between the two being the material and the manufacturing techniques. This was carried out in order for RAL to evaluate how well they could manufacture such a design using two different materials and manufacturing techniques. One horn was drilled from aluminium, and the other was etched from copper. Figure 5.13 shows both of these horns, in addition to a profile of the horn that was also drilled from aluminium. In this section, the results of the measurement campaign carried out using the millimetre wave testbed in NUI Maynooth will be presented. In addition to this, the rectangular to circular transition region that was necessary to match the feed horn to the system source will be discussed.

The millimetre wave measurement testbed uses a Rohde and Schwarz ZVA-24 (4 port) vector network analyser (VNA) in a two port configuration to measure the complex transmission and

Figure 5.13: Horns and horn profile that were manufactured by RAL

reflection coefficients of each port in the W-Band (75-110 GHz), by means of two ZV-110 frequency converter heads, [58]. Since the system is used in a two port configuration, one of the converter heads acts as the system source and the other acts as the detector. Since both heads simultaneously act as both receivers and transmitters, either head can be used in either capacity. This allows the phase and magnitude/intensity of any system under test to be measured accurately in either forward or backward configurations. In this case, the system under test will simply be the feed horn, which will be attached to the head that will be deemed the source, and the beam patterns will be measured by the other head, acting as a detector, by means of planar scans in some detection plane. A waveguide probe will be attached to this head in order to give good measurement resolution by increasing the coupling of the detector to the signal of interest.

The system, at the time of these measurements, offered an effective dynamic range of ≈80 dB over the 75-110 GHz band, based on a measurement bandwidth of 100 Hz. The scanning probe used on the detector head to improve the coupling of the detector to the signal of interest, is a truncated section of WR10 waveguide of wall thickness 1 mm and length 12.5 mm. When this is combined with the 40 mm WR10 section present in the converter head, this is long enough to ensure that only the TE_{10} mode contributes to the signal detected. This probe is also sensitive to orientation, in that it does not detect both orthogonal polarisation of the TE_{10} mode at once. This allows polarisation sensitive measurements to be carried out, enabling the cross-polarisation pattern of the horn to be measured as is desired in this measurement campaign.

The system alignment procedure will first be described, followed by setup modifications that were made during the final measurement campaign in order to improve the alignment of the system and thus the dynamic range. The results of the measurement campaign will then be

Figure 5.14: Testbed setup used in the measurement campaign to evaluate the performance of the optimised horn

presented and compared to simulations that were carried out in MODAL (Maynooth Optical Design and Analysis Laboratory, see chapter 2), using a full physical optics propagator that allow exact comparisons to be made with the measurement system. The results presented in this thesis pertain to the aluminium horn.

5.3.1 Alignment of Measurement System

The scanning system consists of a z-axis translation stage, and an xy raster scanner. One VNA head is mounted on the raster scanner and serves as the detector, while the other head acts as the system source, although due to reciprocity, either head can be used in either mode. For the purposes of these measurements, and with reference to figure 5.14 which shows the experimental setup, the optimised horn, the horn under test, was mounted on the VNA head on the translation stage which was taken to be the system source. This allowed the horn/scanning plane separation to be adjusted to a high degree of accuracy. A waveguide probe was mounted on the second head to act as the detector, and a waveguide probe was mounted to this head as described above. Using Eccosorb, the majority of the detector head was covered, leaving

only the protruding waveguide belonging to the probe visible, to the source. This minimises both reflections in the system and the number of standing waves that form between the source and detector heads of the VNA. In this case, head two was the source head, and head one the detector head. In this configuration, by measuring the S_{12} parameter at all points by means of a planar scan, it is possible to characterise the performance of the horn. By examining the complex value of this parameter at each point in the detection plane, it is possible to measure the complex field of the horn, and thereby recover both amplitude and phase. The other S-parameters can also yield useful information, particularly when determining reflections within the horn etc.

The initial alignment of the system was carried out mechanically. This involves ensuring that in the centre position of the system (the origin of the scanning grid coordinate system) the heads are pointing at each other and at the same height. This is verified by running several one-dimensional test scans using small off axis distances in order to adjust the positions until they are aligned in the x and y directions, resulting in maximal power coupling between the two heads. By analysing the power output from one head relative to the input from the other head, the alignment can be further enhanced. Following the completion of this step, it is possible to perform an adjustment of the alignment on a finer scale by carrying out full two dimensional scans and by performing a Gaussian fit on the measured data. This simulated fit is accomplished by means of a program created in LabVIEW, [59]. Using an initial estimate of the 2-D Gaussian parameters of the measured 2-D field (achieved by a superficial analysis of the data), the program returns values such as the equivalent beam waist at the scanning plane, W, the initial waist radius, W_0, the phase radius of curvature, R, and the distance propagated by the beam from the waist to the scanning plane, z. Of particular interest for the alignment process are the values X_c, Y_c, α and β, all of which are also returned by the program. The first two of these values represent the displacement of the beam maximum in each of the indicated planes from the scanning plane centre, as such these two values taken together give the central point of the measured field. α and β give the angular tilts (about the x and y axes respectively) of the source and detector relative to each other. An additional angle considered is an azimuthal rotation of the horn relative to the waveguide probe, i.e. a relative rotation of these components about the optical axis of the system. This angle is known as γ.

These four values allow the alignment of the system to be improved significantly, relative to the initial alignment. By carrying out relatively quick 2-D scans for a variety of propagation distances (varied by means of the translation stage) and tracking the changes in these four values as a function of this distance, one can determine in what way the system needs to be altered. For example, to reduce the value of α, trigonometry can be used to calculate the distance by which the front or back of the translation stage needs to be raised or lowered in order to affect the correct angular change about the x-axis required to minimise the angular tilt. This results in a well aligned system.

x_c [mm]	y_c [mm]	α [°]	β [°]
-2.28	-1.51	0.21	-0.20
-1.24	2.54	0.38	0.02

Table 5.4: Alignment parameters for identical runs showing inconsistency

x_c [mm]	y_c [mm]	α [°]	β [°]
-2.65	1.66	-0.35	-0.18
-3.36	1.42	0.29	-0.38

Table 5.5: Alignment parameters for identical runs (with waveguide probe having been rotated through 180°) showing inconsistency

While it is desirable to have the system as well aligned as is possible, it is not necessary to perfectly align it or to spend a disproportionate amount of time on this task. MODAL, used for the simulations with which these measurements will largely be compared, allows one to add the small lateral displacements and angular tilts indicated by the fitting program, to the simulation model. MODAL also allows the horn to be mounted at some azimuthal angle relative to the source. Since the horn will likely not be mounted precisely end-on with the source, it is important to be able to simulate this, as will be seen later. In addition to modelling various misalignments in the system, MODAL can also include the affect of the scanning probe, and can carry out the scan of the detector plane by means of an xy raster scanning regime. All of these features mean that it is possible to simulate the actual measurement system very accurately and therefore make accurate comparisons between the simulated and measured datasets, as opposed to relying on a perfect model and attempting to recreate this experimentally.

Following the adjustment of the experimental setup in an attempt to minimise the angles and to keep the offsets constant as a function of propagation distance (proof that the angles have been correctly accounted for), a further higher resolution scan was carried out to accurately determine the magnitude of the offsets in the system. It was clear that despite the careful alignment, there were still issues present in the system. Although the offsets themselves were of an acceptable magnitude, the consistency between identical measurements of these parameters was extremely low, considering the alignment procedure. A sample of the figures obtained are shown in tables 5.4 and 5.5. The offsets also lacked consistency relative to the system configuration. For example, if the waveguide probe was rotated by 180°, then despite the waveguide probe itself being carefully modified to remove asymmetries, a different set of values would be obtained for the offsets, as shown in table 5.5. It was clear from this investigation that other systematic effects were affecting beam quality and alignment. The placement of the Eccosorb in the vicinity of the detector head front (waveguide probe) is the one element of the setup that is reset regularly in a manner that is not entirely reproducible between measurement runs. The placement of this material gives rise to shadowing due to the absorber protruding

in a non-uniform manner from the receiver probe. Since the placement is not entirely regular, the shadowing effect regularly changes, leading to discrepancies between measurements that are dependent on the Eccosorb potentially blocking the beam signal.

In order to address this issue, a piece of Thomas Keating absorber material was modified to be mounted over the waveguide probe in a fixed, and therefore repeatable manner, meaning that the discrepancy between runs should be reduced. The remainder of the detector (behind the waveguide probe) was covered with standard Eccosorb, however this area is not critical for measurements. This arrangement is shown in figure 5.15. It is possible to place this absorber

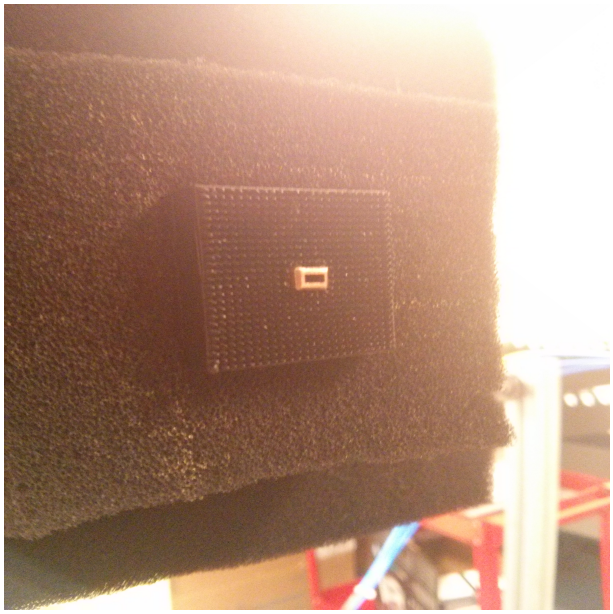

Figure 5.15: Thomas Keating absorber and Eccosorb mounted on the detector head and waveguide probe

consistently, as it is designed to fit over the front of the detector head according to the location of the screws and the waveguide probe itself, thus consistency and good repeatability are achievable. With the probe in place, a new set of scans were carried out in confidence with good alignment and repeatability. It was found that the offsets were consistent with distance, as expected, but critically the values were of an acceptable level and at a fixed distance they were self consistent once the absorber was attached, even when the scan was repeated after rotating the waveguide probe by 180°. These results are summarised in tables 5.6 and 5.7. Run 1 and run 2 are exact repeats of each other, with the setup reset in between scans to test the consistency achievable when the absorber is removed and then replaced.

These results clearly show that the original inconsistent placement of the Eccosorb was a large source of the issues with the beam quality, leading to large misalignments. Examining the measurements with the Thomas Keating absorber material present, it is clear that in isolation both orientations of the waveguide probe give excellent results as the values are individually well within what is acceptable for each parameter. There is also relatively little variation

Absorber	Orientation	x_c [mm]	y_c [mm]	α [°]	β [°]
YES	UP	5.39	1.41	-0.12	-0.26
YES	DOWN	5.54	1.34	-0.3	-0.26
NO	UP	5.57	2.64	-0.47	-0.24
NO	DOWN	4.7	1.95	-0.41	-0.39

Table 5.6: Alignment parameters calculated in various probe/absorber configurations at 100 GHz for run 1

Absorber	Orientation	x_c [mm]	y_c [mm]	α [°]	β [°]
YES	UP	5.77	0.85	-0.2	-0.21
YES	DOWN	5.51	0.93	-0.24	-0.27
NO	UP	5.14	3.49	-0.61	-0.3
NO	DOWN	6.25	1.92	-0.4	-0.15

Table 5.7: Alignment parameters calculated in various probe/absorber configurations at 100 GHz for run 2

between the two orientations of the waveguide probe when the absorber is in place. This is important as it shows clearly that the measurements are extremely sensitive and that there is little asymmetry within the probe and so it is not critical to ensure that the same orientation is used when scans are being carried out for the purpose of comparison. Since the probe must be rotated in order to measure the different polarisations of the horn, it is important that these asymmetries be reduced.

In the case where no absorber is used, the alignment parameters have higher values and there is a greater variation. This is as expected, as the absorber is not present and so the system is more susceptible to multiple reflections in the system or standing waves (on-axis) forming between the probe and the horn. If the surface presented at the detector head is not uniform then this will result in inconsistent results which would otherwise not occur if the absorber were present. This shows clearly that the absorber plays a critical part in accurate and sensitive measurements and allows the measurements to be carried out to a high standard.

5.3.2 Rectangular to Circular Transition

The horn is of circular cross section, however as assumed during the simulations, it is fed by a rectangular WR10 waveguide, which is the case in the testbed in NUI Maynooth. Simply mounting the circular horn onto the rectangular waveguide would result in losses and reflections (due to the mismatch of the geometries) which could easily destroy the performance of the horn. To address this issue, it was necessary to include a transition section that would match the two geometries. This transition was manufactured as part of the horns, and included prior to the throat as an additional section added to the design presented above. The transition is visible in figure 5.13 at the throat of the cross-section model. The transition was designed by

Dr. Stephen Doherty and will now be briefly considered.

The transition could not be included in either the optimisation process or when simulating the measurement setup in MODAL, as although SCATTER exists in both rectangular and cylindrical coordinate systems, it cannot currently handle combinations of the two geometries. MODAL is subject to the same limitations, as it is also based on the mode matching technique when modelling waveguide structures. Due to this, the transition had to be designed separately using a full electromagnetic solver. CST (Computer Simulations Technology) Microwave Studio was used for this purpose.

CST Microwave Studio

CST is a computer aided design (CAD) based program that utilises finite integration techniques in order to calculate the performance of the structure under consideration. The structure is constructed within the program using standard CAD tools. Once the model has been constructed, it is divided up into individual cells by means of imposing a mesh structure of variable density on it. Maxwell's equations are then solved in each of the cells (using discrete versions of the integral form of the equations) in order to determine the electromagnetic performance of the structure, [57]. For this to occur, then mathematically two orthogonal mesh grids are required, and the equations are solved in each cell of each grid.

CST implements the above approach by means of three separate solvers, the choice of which solver being determined by the efficiency and accuracy required as well as by the type of system under consideration, for example, resonant versus dynamically stable systems. The three solvers are:

- Transient solver
- Frequency domain solver
- Eigenmode solver

The eigenmode solver is not typically used in the applications considered in this thesis (usually used in the analysis of *closed* resonant structures), and as such only the first two solvers shall be examined.

Transient solver: This solver allows the calculation of the performance of the structure under test over a finite bandwidth and time period, where all frequencies to be considered are calculated in one simulation. It proceeds by using a hexagonal mesh grid structure and considers the electric voltage and magnetic flux in each cell, where the variables are located alternately in time. This allows full bandwidth simulations to be carried out in a highly accurate manner, accounting for losses. The full set of scattering matrices, along with the beam patterns associated with the structure can be calculated across the defined range of frequencies in one iteration of this solver. It is possible in CST to measure many different parameters,

and monitors to record these parameters can be easily integrated into any given run of the solver. One limitation of this solver can often be encountered when resonant structures are considered, for example cavity structures as will be encountered later in this thesis. Depending on how resonant the structure is, the initial excitation may result in eddy currents and various instabilities in the system which initially return erratic and potentially non-physical results. If examined over a sufficiently long time interval, the natural damping of the systems will suppress these effects (with the system entering a steady state) and return sensible results, however the default time interval for this solver will often be insufficient to allow this to happen, and a suitable time interval will often result in unacceptable run times. One method to solve this issue is to use the frequency domain solver.

Frequency domain solver: This is best used for simulations that require the simulation of a narrow band. The band should be narrow as each frequency requires a separate iteration of the solver which could be time consuming. It utilises hexahedral or tetrahedral grids, whichever best suits the structure, and is capable of performing efficient scattering parameter calculations on structures that are loss free and highly resonant. As stated above, it is particularly appropriate to analyse resonant structures whose steady state solution is required, in the frequency domain, as they require large settling times in the time domain, and so using the transient solver would yield potentially inaccurate results, or accurate results with a severe time penalty. It utilises the time harmonic case for Maxwell's equations, as outlined in [57]. This solver does not calculate field patterns directly. For CST to calculate the fields associated with a structure that was solved using this solver, an add-on is required that carries a significant time penalty. Rather, using the calculated scattering parameters, the fields can be calculated outside of CST if required.

In creating the structures to be analysed, CST allows the consideration of various materials. The two basic materials are perfectly electrically conducting (PEC) and vacuum. There is also a library available that contains many additional 'real' materials such as lossy metals and conductive and dispersive materials. These realistic materials have real electrical and thermal properties. It is also possible to customise materials by defining required values for parameters such as, for example, the electrical and thermal properties, permittivity and magnetic permeability, as well as considering whether the material is isotropic or anisotropic.

In designing the transition, the transient solver was used, along with PEC. This choice of material ensures that currents can flow unimpeded along the bounding walls of the structure, resulting in a lossless structure. The design of the transition will now be presented.

Transition

In designing the transition, several designs were considered. Rectangular to circular transitions of the geometries required by the frequencies in question are difficult to manufacture within the tolerances required (of the order of μm). Their performance is highly dependant on the

structure and so any deviations from the design will result in a degradation in performance. To this end, a design had to be found that made a good compromise between performance and manufacturability.

One design that was considered but had to be rejected ultimately, was the tapered transition, as shown in figure 5.16. In this case, the guide wavelength [57] of the TE_{10} mode in

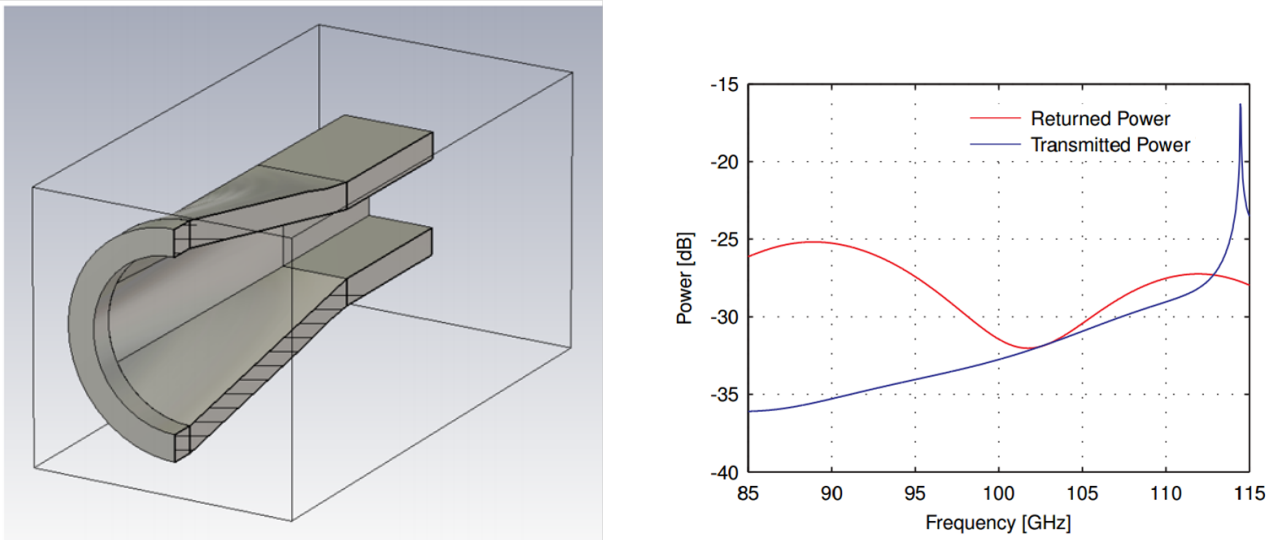

Figure 5.16: Tapered rectangular to circular transition, with total returned and transmitted power across the band

the rectangular waveguide was matched to that of the TE_{11} mode in the circular section, with the radius of the flared portion of the circular section gradually increasing to the radius of the throat of the horn. The matching of guide wavelengths was carried out in order to minimise reflections (returned power) at the interface between the geometries, as this would negatively impact on the power entering the horn. Matching the two modes in this manner reduces reflections, as the modes can couple strongly to each other, allowing much of the power contained in the TE_{10} mode to couple into the TE_{11} mode. The flared section was of sufficient length to allow a gradual change in radius, so as to not excite additional modes via the inter-modal scattering that can occur when the radius of a waveguide changes suddenly. These features of the transition section ensured that the expected mode mixture arrived at the assumed 'input' of the horn, namely the TE_{11} mode. The design works well, as is clear from the reflected and transmitted power shown in the accompanying plot. Ultimately the design was rejected due to manufacturing difficulties.

The design that was settled upon is shown in figure 5.17. This is a single step rectangular to circular transition. The single step makes the transition between the geometries more extreme than before, and so will naturally result in higher reflections due to the resulting mismatch between the modes. This is seen clearly by noting the reflected power across the band being significantly higher than in the tapered transition case, as shown in figure 5.18. Matching the

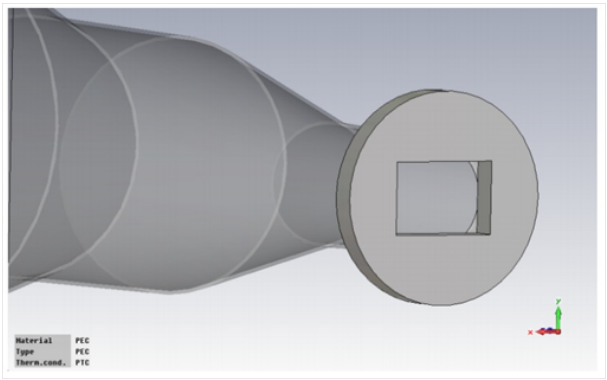

Figure 5.17: Single step rectangular to circular transition, mounted on horn throat

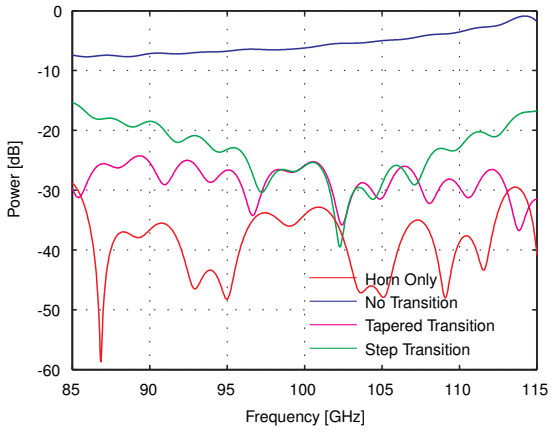

Figure 5.18: Returned power for different transitions to the WR10 waveguide

guide wavelengths during the design process allowed this issue to be minimised to a degree. This acceptably low level of reflected power results in reduced levels of inter-modal scattering at the throat of the horn. The transition was designed for operation at 100 GHz, which is why the returned power is particularly low at this frequency.

Figure 5.18 shows the returned power (S_{11}) for four different cases; the horn only, the horn mounted directly onto the WR10 waveguide, and the horn mounted to the WR10 waveguide via both of the transition sections described above. The isolated horn displays the lowest return loss, as expected. Since there are no transitions between geometries, there are no additional reflections relative to those that already occur in the horn. When the horn is mounted directly to the WR10 waveguide, the curve corresponding to 'no transition' in figure 5.18, there is a sudden change from rectangular to circular waveguides and this increases reflections relative to the case where the horn is isolated. This case gives the highest level of return loss, as there is no attempt to match the geometries in such a way as to reduce this. When the tapered transition is included, the return loss is lower than the no transition case, giving good performance across the band, remaining within approximately ±5 dB of an average level of -30 dB. The step transition gives a higher level of return loss across the band due to the fact that the geometries

are not as well matched as a result of the change occurring via a step rather than gradually. The return loss in this case varies between -15 dB and -39 dB.

Although the tapered transition gives lower levels of return loss, manufacturing difficulties prevented this transition from being included with the horn, and so the step transition was selected, as it was possible to manufacture this, and although not as high performance as the tapered transition, it still delivers satisfactory levels of performance, including at the design frequency.

5.3.3 Scan Results

Several scans were completed using this aligned setup, using the aluminium horn, in order to evaluate its performance. The scans were compared with the results obtained from simulating the experimental setup using the in-house developed software package MODAL (Maynooth Optical Design and Analysis Laboratory), as was described previously, which in this instance uses the mode matching technique to simulate the horn and a full physical optics treatment to propagate the beam to the desired scanning plane.

For each scan, a frequency range of 75-110 GHz was swept, using steps of 5 GHz, and a measurement bandwidth of 100 Hz. Table 5.8 shows the off-axis distances used as the limits for the xy raster scan either side of the aligned centre of the system, and the distance set using the z-axis translation stage, relative to the zero point. The zero position of the translation stage

Scan Type	**X/Y [mm]**	**Z [mm]**
Co/Cross-polar Scans	$\pm$129(3)	$\pm$150(50)

Table 5.8: Dimensions used in scans. The bracketed figures indicate the stepsize used for the respective parameters in the indicated units.

corresponds to a horn aperture/waveguide probe separation of 190 mm. Using the translation stage to vary this distance results in the propagation distance varying between 40 and 340 mm, in steps of 50 mm.

Figures 5.19 to 5.21 show the measured co-polar beam pattern at distances of 40, 190 and 340 mm at 100 GHz. The indicated off-axis distances are in units of mm, with intensities in units of dB and phase in units of radian. These beam patterns clearly demonstrate that the measured beam shows a high degree of symmetry and a well defined main beam with low sidelobe levels, as required. The beam retains these properties across all distances, despite the fact that at its furthest it has propagated 300 mm (approximately 100 wavelengths). The regular and symmetric beam pattern is retained with no nulls appearing and no appreciable deterioration in sidelobe performance. The beam simply increases its spatial extent whilst reducing in on-axis power, as expected due to diffraction. This is indicative of largely single-

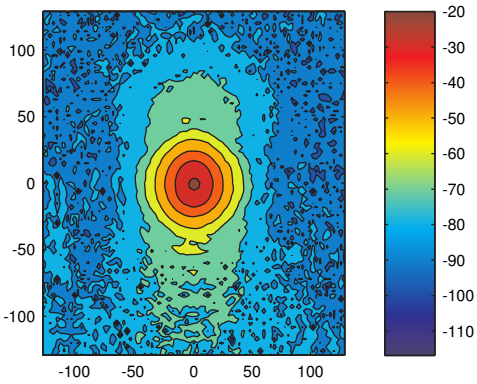

(a) 40 mm Intensity

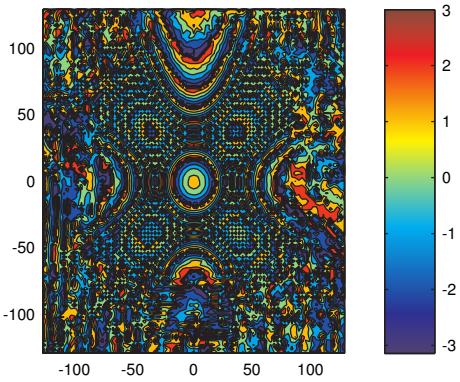

(b) 40 mm Phase

Figure 5.19: Measured intensity and phase patterns for a propagation distance of 40 mm at 100 GHz

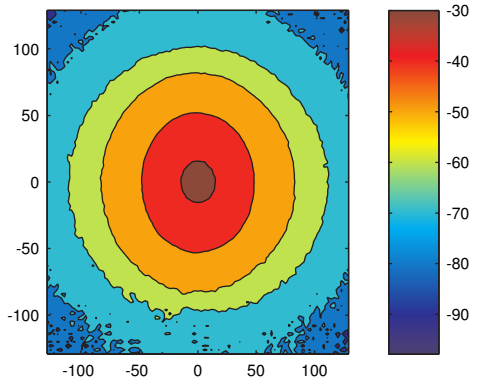

(a) 190 mm Intensity

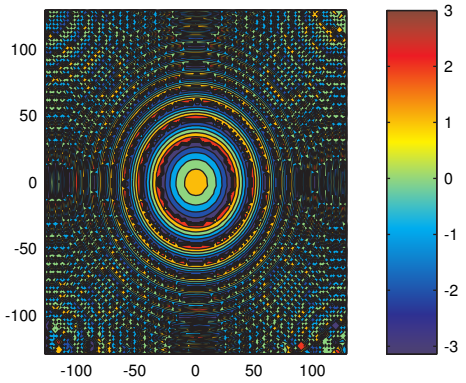

(b) 190 mm Phase

Figure 5.20: Measured intensity and phase patterns for a propagation distance of 190 mm at 100 GHz

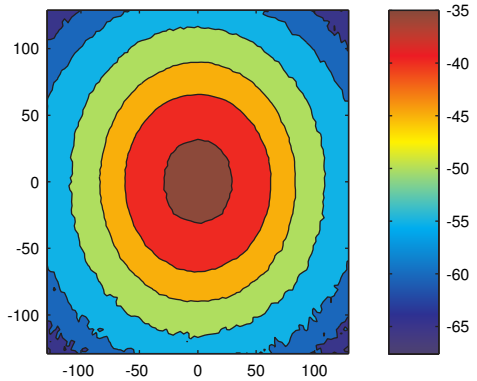

(a) 340 mm Intensity

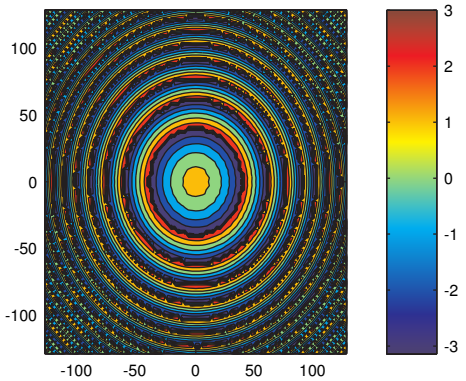

(b) 340 mm Phase

Figure 5.21: Measured intensity and phase patterns for a propagation distance of 340 mm at 100 GHz

moded performance, as required, as if additional modes were present then it would be possible for unwanted contributions to appear over such large propagation distances.

Figures 5.22 to 5.24 show the measured cross-polar beam at distances of 40, 190 and 340 mm at 100 GHz. The indicated off-axis distances are again in units of mm. These measurements were obtained by adding a length of waveguide between the detector head and the waveguide probe that features a 90° twist along its length. Since the waveguide probe is polarisation sensitive, this rotates the probe such that it is detecting the polarisation that is orthogonal to the co-polar beam, giving the cross-polar beam. At the low levels expected for the cross-polar

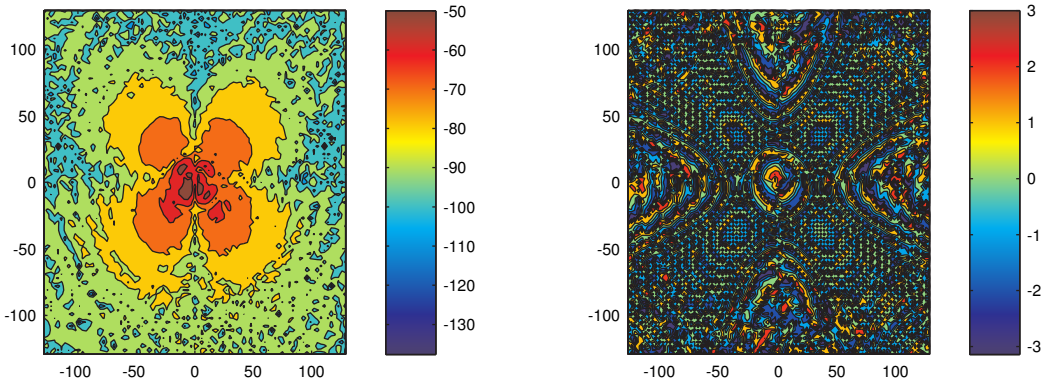

(a) 40 mm Intensity (b) 40 mm Phase

Figure 5.22: Measured intensity and phase patterns for a propagation distance of 40 mm at 100 GHz

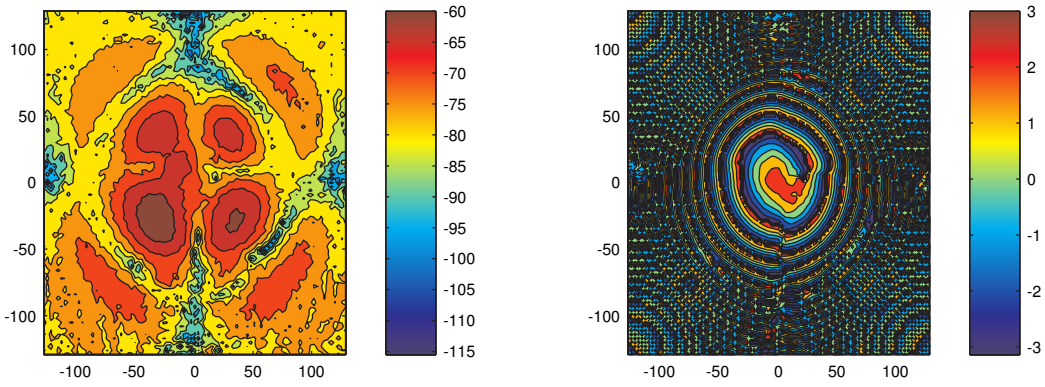

(a) 190 mm Intensity (b) 190 mm Phase

Figure 5.23: Measured intensity and phase patterns for a propagation distance of 190 mm at 100 GHz

power, the field patterns are difficult to measure, however at both extrema of the distances measured, clear and well defined patterns were obtained. This is due to the fact that the system was very well aligned and carefully calibrated, giving a high dynamic range. As in the co-polar case, they diffract with distance in a well defined manner, as expected. The quad lobe

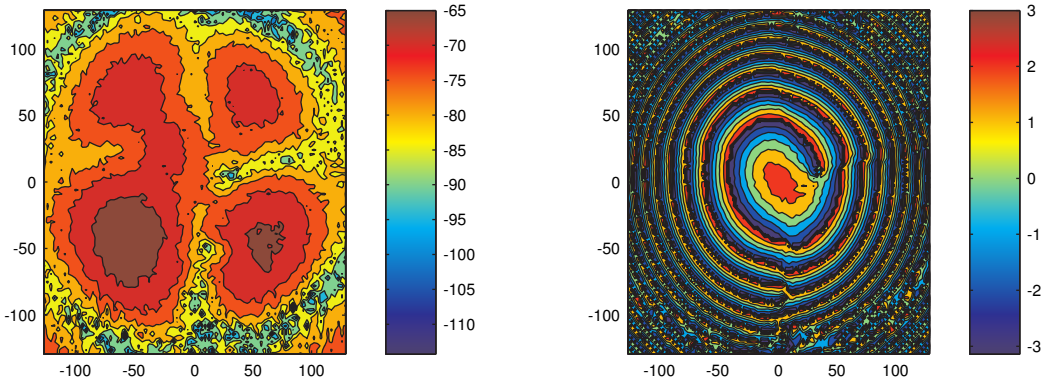

(a) 340 mm Intensity

(b) 340 mm Phase

Figure 5.24: Measured intensity and phase patterns for a propagation distance of 340 mm at 100 GHz

structure normally associated with the cross-polar field is clearly visible in each case. The fact that these features are visible clearly illustrates the value of careful alignment and calibration of the system, as it is not possible to measure these features so clearly otherwise, apart from away from the design frequency where the cross-polar levels rise to higher levels, for example at 75 and 110 GHz.

Of particular interest in these field patterns is the top left quadrant. For each of the propagation distances, there is an irregularity here that results in the expected null along the x axis being completely or largely obscured, along with half of the expected y null. This is likely due to the co-polar field 'leaking' into the measurement. Several possible reasons exist for this. It is possible that the horn and/or twisted waveguide section was mounted at a slight azimuthal angle relative to the waveguide from the VNA head. Since to measure the cross-polar beam one must rotate the detector by precisely 90° relative to the 0° azimuthal plane of the horn, if either of these components are not mounted on the VNA head waveguide so as to be aligned with this plane, then a rotation of 90° will not result in a perfect measurement of the cross-polar pattern, but rather some convolution of the cross-polar and co-polar patterns. Since the co-polar beam is of a much higher level than the cross-polar beam, the effect when attempting to measure the cross-polar beam will be tangible. Of course, if this mounting issue is the source of the discrepancy then the co-polar patterns presented above will also have been affected. This will have resulted in some cross-polar power leaking into the co-polar measurement, but since the cross-polar levels are so low relative to the co-polar levels, the effect will not have been noticeable. Another possible explanation is that the two polarisations are not completely separable due to the modal composition of the field, and so the observed leakage is inherent and unavoidable.

Vertical and horizontal (E and H plane) cuts of the co-polar plane, as well as the 45° of the cross-polar plane, are illustrated for propagation distances of 40 and 340 mm in figure 5.25. It

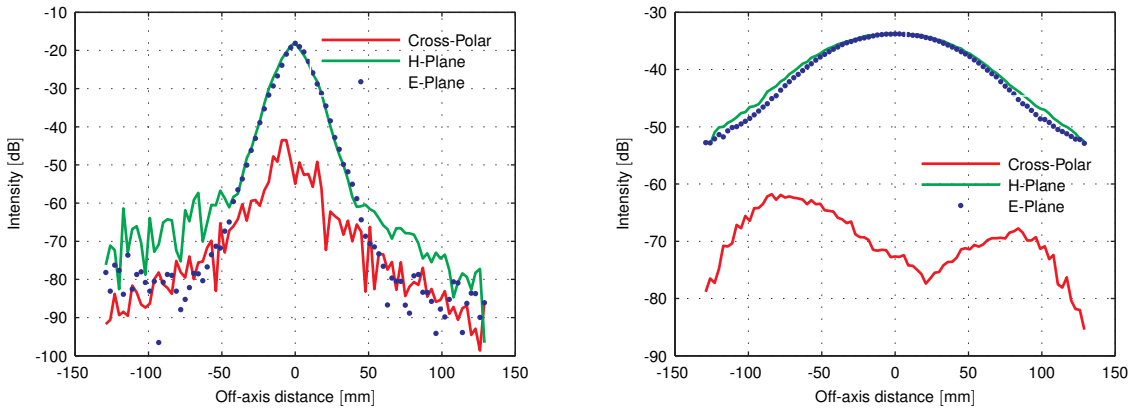

(a) 40 mm

(b) 340 mm

Figure 5.25: Cuts from the measured beam patterns of the horn at 40 and 340 mm at 100 GHz

is clear from these beam patterns that the horn has the predicted polarisation properties (i.e. a large separation in terms of power levels between the co and cross-polar fields), as required. The features such as co-polar leakage into the cross-polar field are clear from figures 5.25(a) and (b), where the on-axis null is not reproduced and there is some asymmetry in the left side of the beam pattern, just off-axis. In figure 5.25(a), the co-polar field has not sufficiently diffracted to remove the sidelobes from the measurement plane. The cuts show low sidelobe levels, a good Gaussian main beam with very little beam shouldering and a high degree of symmetry. Figure 5.25(b) shows the effects of diffraction on the beam, it having propagated one hundred wavelengths. The on-axis power has reduced significantly, however the beam is still highly Gaussian with little shouldering and the high degree of symmetry is retained, indicating excellent performance.

A critical design parameter of this feedhorn is the symmetry of its main beam. This ensures that the measurements that the horn makes will be independent of azimuthal orientation, which is a highly desirable characteristic. In order to quantify the degree of symmetry of a beam is to calculate the coupling between the 0° and 90° cuts of the field. The degree of the coupling is indicative of the degree of circular symmetry present in the (in this instance) co-polar beam. If the coupling of the two cuts is low, then the beam is likely to be elliptical, whereas a high coupling of unity would mean that the beam would be perfectly circular. The coupling between the two cuts, c_{eh}, is given by equation 5.1.

$$c_{eh} = \frac{\int_0^b E_E E_H^* 2\pi x \mathrm{d}x}{\sqrt{\int_0^b E_E E_E^* 2\pi x \mathrm{d}x}\sqrt{\int_0^b E_H E_H^* 2\pi x \mathrm{d}x}}, \tag{5.1}$$

where E_E and E_H represent the 0° and 90° cuts of the complex co-polar beam respectively, and b is the off-axis distance corresponding to the maximum angle considered in the farfield representation of the field. Since the complex forms of the fields are used, the calculated

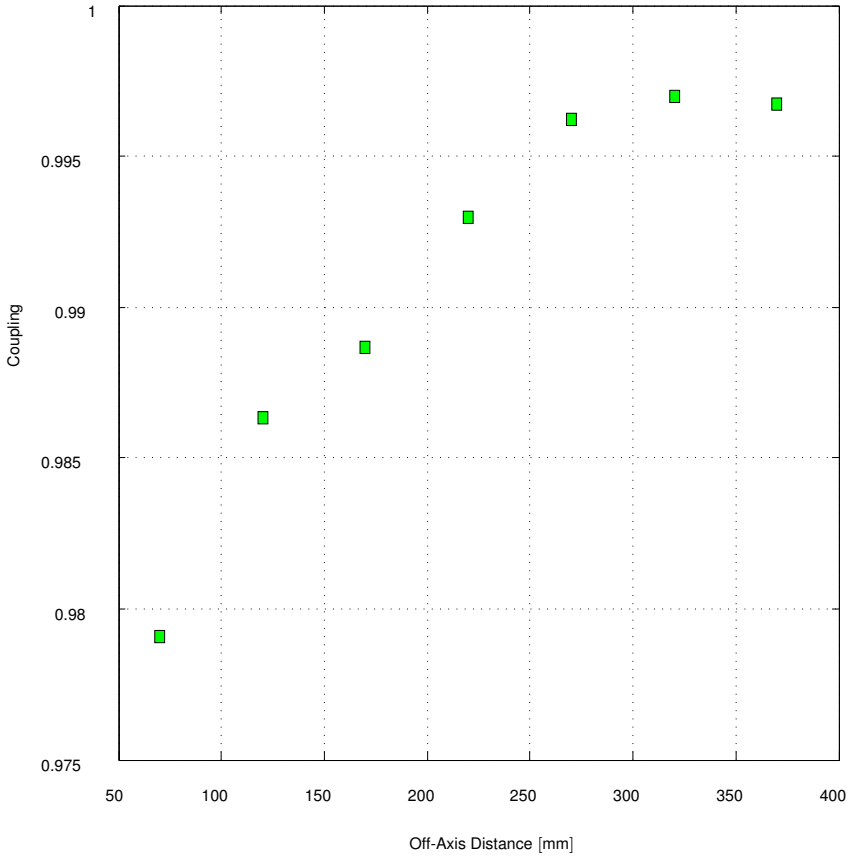

Figure 5.26: Coupling between 0° and 90° co-polar cuts (beam symmetry) as a function of propagation distance at 100 GHz

coupling considers potential phase mismatches and so gives a true indication of beam symmetry.

The symmetry of the co-polar beam as a function of propagation distance at the central frequency of 100 GHz is displayed in figure 5.26. It is clear that in general the coupling increases as propagation distance increases. This is due to the fact that as the propagation distance increases, one enters the farfield and so the field patterns become more stable. It should be noted that for all distances considered, the coupling remains above 97.5%, and so the horn performs excellently at the design frequency for a large range of propagation distances. For most distances, the coupling is above 98%, rising steadily to over 99.5%.

An examination of the beam symmetry as a function of frequency for a fixed propagation distance of 40 mm (the shortest distance to the horn) was carried out. Figure 5.27 shows the coupling between the 0° and 90° cuts of the co-polar plane as a function of frequency at this distance. The symmetry of the co-polar field across the band is excellent, remaining high. At the lower end of the band, 75 GHz, the symmetry reaches its lowest value of approximately 94%, however this is still indicative of very good beam symmetry. The coupling shows an initial steep increase in symmetry to 96.5% at 80 GHz. A further increase is observed at 85 GHz with a slight decrease at 90 GHz. Beyond this, the coupling increases even further, tapering

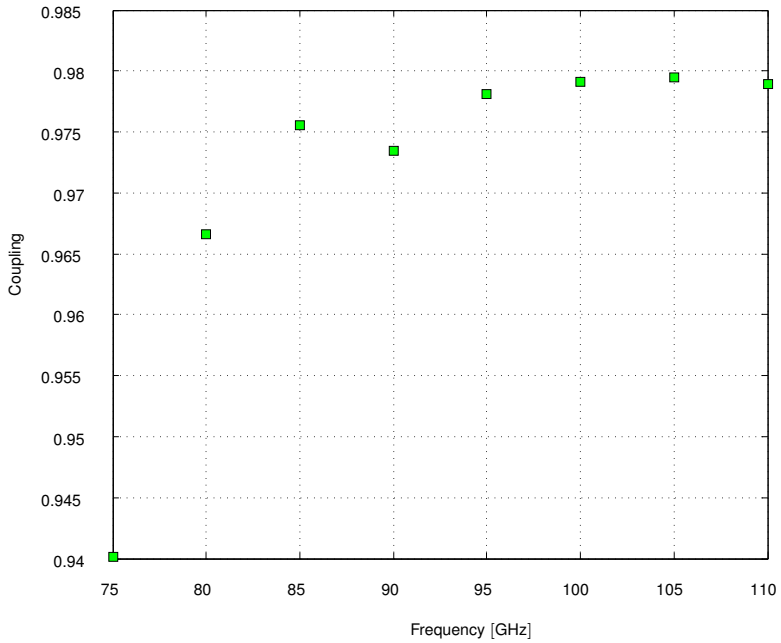

Figure 5.27: Coupling between 0° and 90° co-polar cuts as a function of frequency, z=70 mm

asymptotically to a value of just below 98%. Other than at 75 GHz, the coupling between the two planes does not go below 96.5%, and so the horn performs well in terms of producing a highly symmetric co-polar beam across the band. Considering that the optimisation algorithm tailored the design for operation at exclusively 100 GHz, the fact that such good performance is noted in terms of beam pattern and beam symmetry across the band is an excellent result, although the performance is markedly in the portion of the band above the design frequency compared to the portion below it.

The horn was designed to minimise the power contained within the cross-polar field at the design frequency of 100 GHz. As shown previously in figure 5.5, the predicted performance of the horn indicates good performance from this point of view at 100 GHz and above, but poorer performance below 100 GHz. The simulated and measured values of the cross to co-polar ratio are shown in figure 5.28. The simulated value (using SCATTER) for this ratio at 100 GHz is just above -45 dB. This exceptionally low was obtained as this is the frequency at which the horn was designed, whilst being optimised for this particular parameter. This was carried out, in theory, at the expense of performance at other frequencies and metrics, however despite this, as shown by the excellent beam symmetry and low sidelobe levels, the overall performance of the horn is still relatively good across the band for the metrics of interest. Towards the higher end of the band the simulated performance of this metric is still excellent, and although below the design frequency the cross-polar power increases, the performance is still excellent down to 80 GHz when the fact that the optimisation process was carried out at a spot frequency is kept in mind.

Comparing the measured values for the cross to co-polar ratio to the simulated values

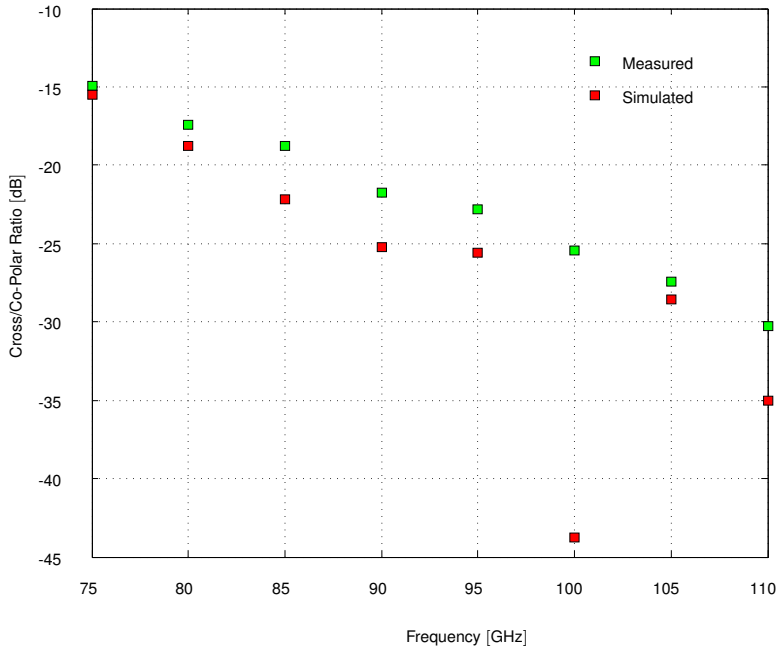

Figure 5.28: Peak cross-polar to co-polar ratio (simulated and measured) as a function of frequency, z=70 mm

reveals close agreement, apart from at 100 GHz, where the predicted value is significantly lower than that measured. The origin of this discrepancy is possibly in the rectangular to circular waveguide transition that exists in the physical system but is not currently accounted for in the simulations illustrated. In the simulations, the source feed is assumed to consist of circular waveguide modes of unity power and that these are present at the opening of the horn throat, exciting the modes that then propagate through the horn. In the physical system however, the source is a rectangular WR10 waveguide, feeding the circular horn via a rectangular to circular transition. This results in potentially significant reflections occurring due to the mismatch between the modes of differing geometry. These reflections will result in less power than expected entering the system but will also result in an increase in cross-polar power. These two factors combine to give the resulting increased value of the cross-polar field relative to the co-polar field. This will be accentuated if the horn is not mounted perfectly on the source head, i.e. there is a misalignment between the transition and the source waveguide. The transition was based on a perfectly aligned system so any deviation from this will result in reduced performance from further increased reflections, and so the cross-polar levels will be increased owing to the fact that such misalignments could result in the fractional excitation of the orthogonal polarisation in the horn. Since the transition was designed at one frequency, it suffers from deteriorated performance away from the design frequency and so this will clearly impact upon the performance of the horn.

Another possible source of this discrepancy is that the solution obtained from the optimisa-

tion algorithm is mathematically unstable. This is likely, as the design is so heavily customised for operation at 100 GHz that although the performance at this frequency is mathematically real, it depends on an absolutely perfect system, and so any deviation from this will result in degraded performance. The measured values shown in figure 5.28 show an almost linear decrease in cross-polar power as a function of frequency. The simulated values also show a relatively steady decrease, apart from at 100 GHz where the value simulated in SCATTER shows a sudden, large, increase. The fact that this data point is such an outlier relative to the other values implies that it is based on a mathematically unstable solution, likely due to the fact that the horn was designed solely for 100 GHz operation and the design was optimised to give the best performance at this frequency, assuming a perfect horn and source. Any deviations from this would clearly result in a degradation in performance, which would be particularly true if the solution was not stable. In the real measurement system such imperfections will exist, including the presence of the transition which was not accounted for in the design process, and so the cross-polar field will be different to what is predicted, particularly at 100 GHz where the performance depends on a perfect system. Since the trends and values of both datasets agree so closely, this is likely to be the case. The measured performance at 100 GHz is still excellent and so this deviation from the predicted value is not an issue. It is also to be noted that the measured values are higher than the simulated values at each point. This is due to the transition section not being included in the simulations and so the additional reflections are not accounted for. Overall, the measured values for this metric across the band are satisfactory, considering the simple design process.

The location of the phase centre of the horn was also examined as a function of frequency. This was calculated by using the Gaussian fitting program that was used in the alignment stage in order to fit a Gaussian to the full two dimensional scan of the co-polar field of the horn at each frequency at the furthest, nearest, and central values for the experimental propagation distance. In the scans that were undertaken to determine the phase centre, this corresponds to horn aperture/detector plane separations of 246, 346 and 446 mm. Large distances were chosen so as the scan would be undertaken in distances that entered progressively deeper into the farfield. The beam patterns in the farfield are relatively stable and are free from nearfield effects, thus the average phase centre position at a given frequency across the three distances will be a good indication of the true phase centre position.

The location of the phase centre is found as follows. The fitting program returns a value for the propagation distance z, which gives the distance propagated by the fitted Gaussian beam from the source. Comparing this to the experimental propagation distance (z_{exp}) allows the location of the phase centre, Δ_{pc}, to be determined according to equation 5.2.

$$\Delta_{pc} = z - z_{exp}. \tag{5.2}$$

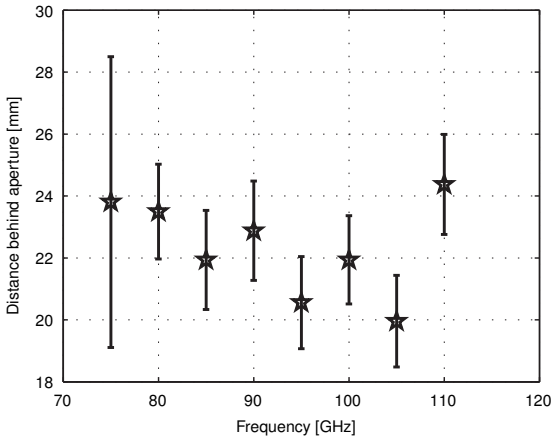

Figure 5.29: Phase centre location behind the horn aperture as a function of frequency

From this equation, if Δ_{pc} is positive then this implies that the fitted Gaussian travelled farther than the experimentally measured distance, meaning that the source of the beam, the phase centre of the horn, is located behind the aperture of the horn, a distance Δ_{pc}. A negative value for this quantity implies that the phase centre is located in front of the aperture of the horn by that amount. It was found that for all frequencies the phase centre was located behind the horn aperture, as shown in figure 5.29 which shows the location of the phase centre as a function of frequency. This quantity was calculated by finding the average location of the phase centre for each propagation distance at each frequency. The phase centre was found to vary between approximately 20 and 24 mm across the band, remaining behind the aperture of the horn at all times. The errors were calculated by calculating the standard deviation of the location of the phase centre across the three distances for each frequency and adding 1 mm to the standard deviation in each case, as this is the error on the metre stick which was used to determine the horn aperture/waveguide probe aperture separation. The error on the 75 GHz is large as the location of the phase centre at this frequency was calculated to be 28 mm behind the aperture for a 346 mm separation. For the other separations, the location was approximately 21/22 mm and so this gives rise to the large standard deviation as the average value is also elevated.

One further parameter to be considered was the return loss of the horn, which physically corresponds to the S_{11} parameter. Figure 5.30 shows a comparison between simulated and measured values across the band, where the simulated values represent the integrated S_{11}, and the measured values represent the S_{11} averaged across all distances at a given frequency. Three datasets are given. The measured results are obtained by direct recording of the S_{11} parameter from the VNA. The simulated values were obtained independently using SCATTER and CST. SCATTER does not consider the rectangular to circular waveguide transition section, whereas CST does. The results are as expected. The presence of the transition should increase the return loss of the horn due to the increased reflections, and so not including it should result in maximum throughput across the band, which is exactly what is observed for the results

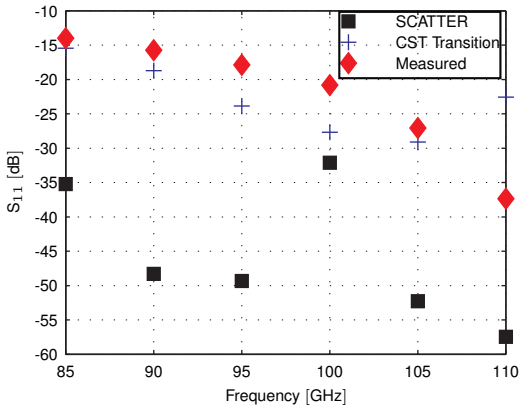

Figure 5.30: S_{11} as a function of frequency, as measured and in two simulated cases

obtained using SCATTER which does not include the transition. Here, levels of <-30 dB are realised across the band, with a minimum of <-55 dB at 110 GHz. The CST and the measured results agree more closely, as the transition is included in CST. They both follow the same decreasing trend across the band, broadly agreeing, apart from at 110 GHz where CST predicts an increase to ≈-22.5 dB and -37 dB is measured. The measured values are averaged across the scan-plane for all distances, and so discrepancies can occur due to phase issues when the cables bend etc. however these discrepancies should average out. The likely cause of the disagreement lies with the mounting of the horn on the VNA head and the mounting of the head itself, as the performance of the transition is based on the angle gamma (relative angle of horn/waveguide probe, corresponding to a rotation about the axis joining the horn and detector in figure 5.14) being zero, and the simulated figures also depend on no other misalignments in the system. These misalignments are not currently considered in the CST model, however their impact will be examined later in this section using MODAL. Overall, close agreement is observed between simulation and measurement, and the measured values are satisfactory, particularly when it is again remembered that the horn was optimised at one frequency with respect to a different metric.

Figures 5.31 to 5.33 show a comparison between the simulated and measured intensity and phase distributions of the horn for the lower and upper limits of the propagation distance, 40 and 340 mm. In each case, the plots are normalised to the maximum value contained within the respective datasets. To gain an insight into the difference between co and cross-polar levels, figure 5.28 can be referenced for both the measured and simulated cases.

It is clear in each case that there is excellent agreement between the measured and simulated beam patterns in both the co-polar vertical and horizontal cuts, and as such the model implemented in MODAL predicts the behaviour of the entire measurement system accurately. The presence of the rectangular to circular transition in the physical system appears then to not have an impact on the beam patterns, implying that the transition does not excite any

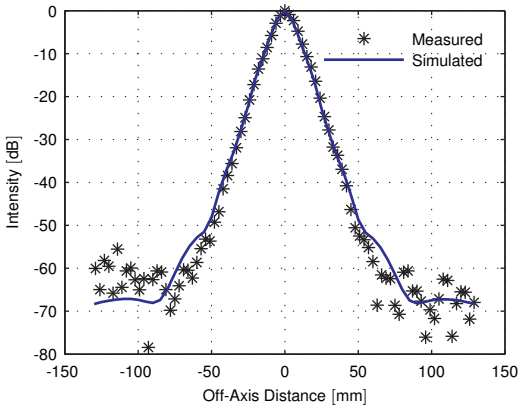

(a) 40 mm Intensity

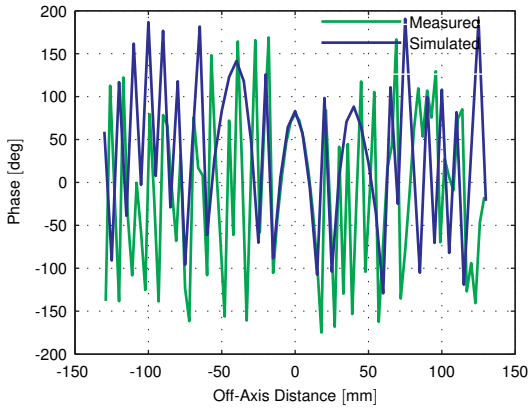

(b) 40 mm Phase

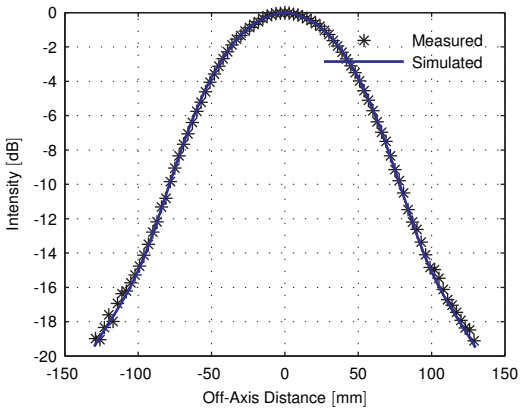

(c) 340 mm Intensity

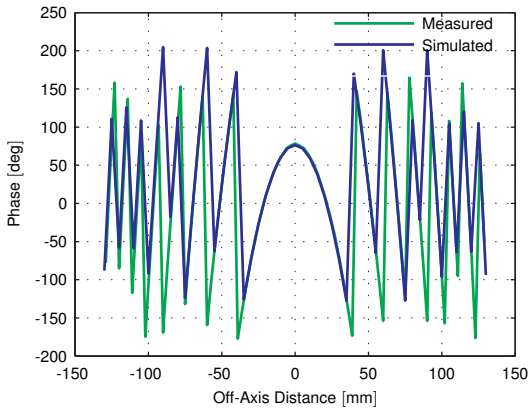

(d) 340 mm Phase

Figure 5.31: Measured and simulated intensity/phase comparisons for horizontal cuts through the co-polar plane at 100 GHz

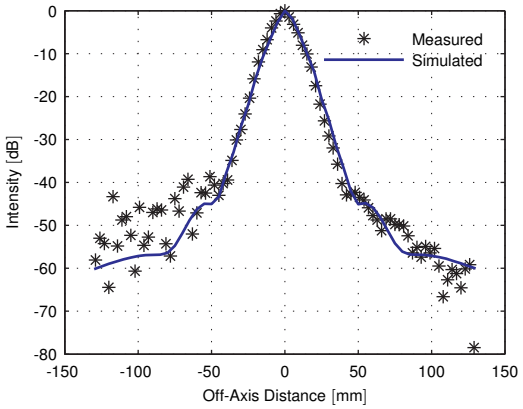

(a) 40 mm Intensity

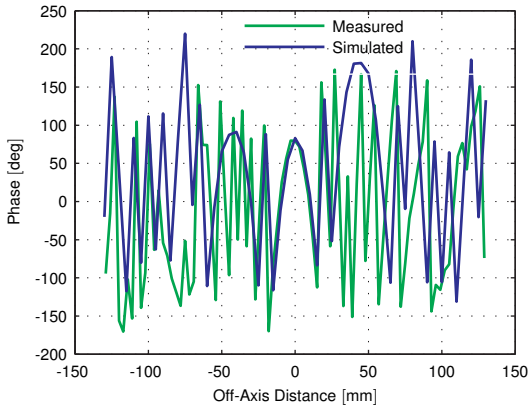

(b) 40 mm Phase

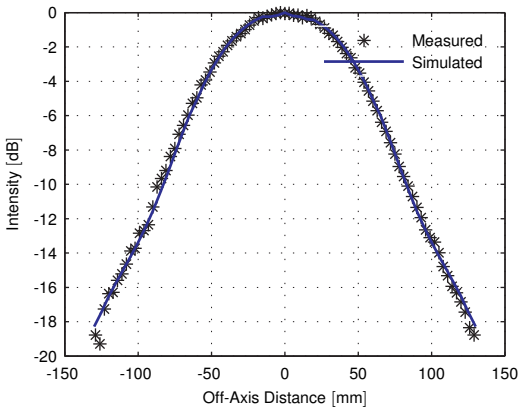

(c) 340 mm Intensity

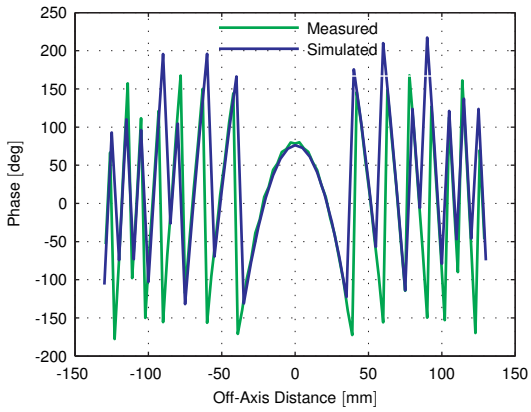

(d) 340 mm Phase

Figure 5.32: Measured and simulated intensity/phase comparisons for vertical cuts through the co-polar plane at 100 GHz

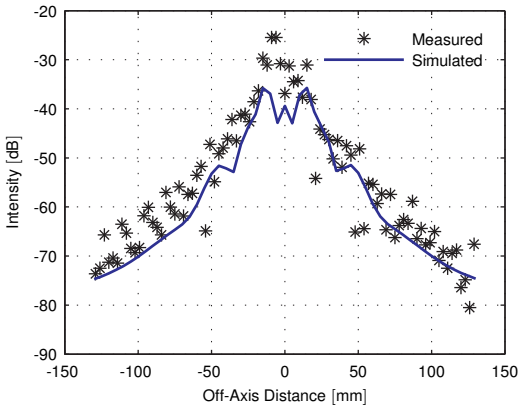

(a) 40 mm Intensity

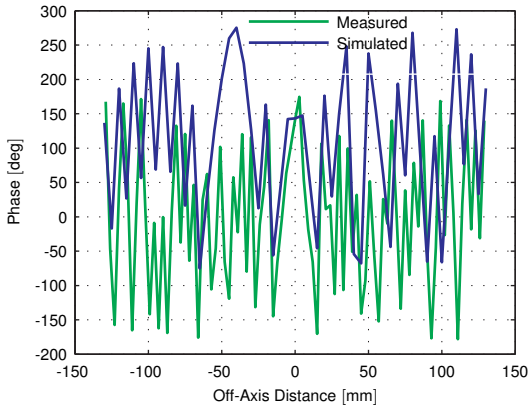

(b) 40 mm Phase

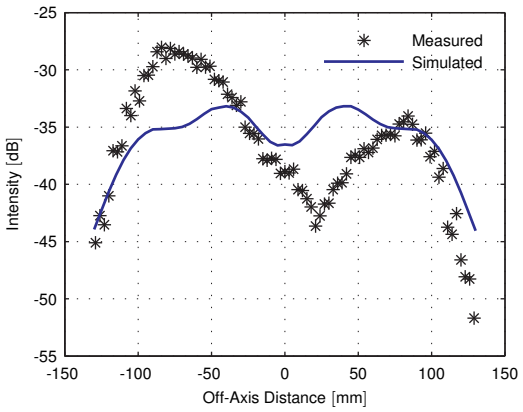

(c) 340 mm Intensity

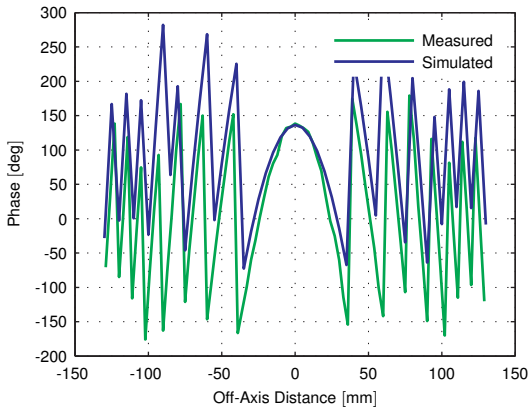

(d) 340 mm Phase

Figure 5.33: Measured and simulated intensity/phase comparisons for 45° cuts through the cross-polar plane at 100 GHz

modes which propagate through the horn that are not accounted for by MODAL. These are normalised comparisons however (as is necessary when comparing simulations and measurements), and so the transition may reduce co-polar power which would appear as the increased levels of reflection which would in turn increase cross-polar levels. In order to quantify the agreement between the simulated and measured datasets, the coupling between a full two dimensional scan plane and the equivalent measurement plane in MODAL was calculated according to equation 5.3, where A is the area of the scanning plane in xy coordinates, E_{meas} is the measured complex field and E_{sim} is the simulated complex field.

$$c = \frac{\int_A E_{meas} E^*_{sim} \mathrm{d}A}{\sqrt{\int_A E_{meas} E^*_{meas} \mathrm{d}A} \sqrt{\int_A E_{sim} E^*_{sim} \mathrm{d}A}}. \tag{5.3}$$

The coupling was evaluated both at a fixed frequency for varying propagation distance (40 to 340 mm in 50 mm steps) to examine the accuracy of the model as a function of distance, and also at a fixed distance of 40 mm for frequencies varying in the range 75-110 GHz in 5 GHz steps to test the model and physical horn as a function of frequency. The results are given in figure 5.34. As a function of distance, at 100 GHz, the coupling increases steadily as distance

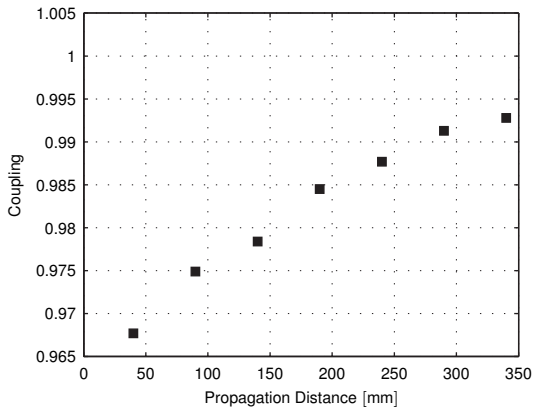

(a) Simulated/measured data coupling as function of distance at 100 GHz

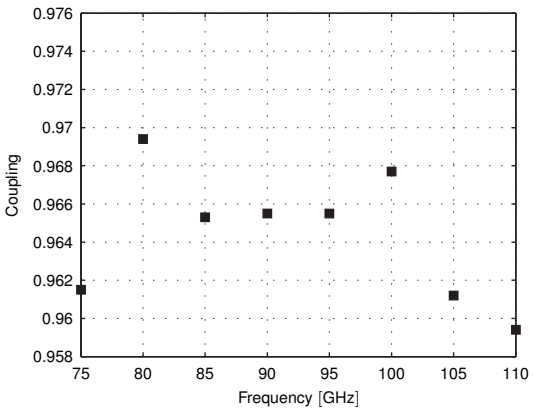

(b) Simulated/measured data coupling as function of frequency at 40 mm

Figure 5.34: Two dimensional model/measured coupling

increases. From 40 mm to 340 mm the coupling increases from ≈96.55% to ≈99.35%, with the rate of increase of the coupling beginning to taper off for larger propagation distances. This is sensible, as when the farfield is approached the field pattern begins to stabilise as the induced currents on the measurement probe which are not accounted for in the simulation become negligible. Thus, the agreement with the model should begin to become more consistent which is as observed. For the case where the frequency was varied for a fixed distance of 40 mm, the coupling oscillates between ≈95.9% and ≈96.9%. This still gives excellent agreement, however it is slightly less than the case where the distance is varied at a constant frequency and is more erratic. This is explained by the waveguide transition step. This is not modelled in MODAL

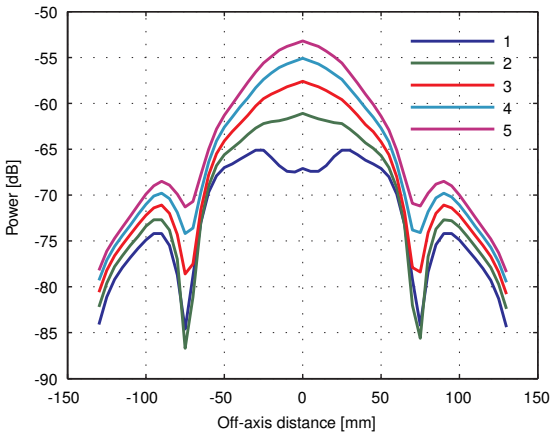

Figure 5.35: Simulated variation of cross-polar field with gamma, for gamma $\in [1°, 5°]$

and is highly wavelength dependent, thus a notable discrepancy will be introduced when the measured results (which do include the transition) are compared to the simulated results (which do not include the transition), away from the design frequency. Despite this, the coupling is still excellent and will improve if the transition is integrated with the model. Figures 5.27 and 5.26 show a high degree of symmetry in the co-polar beam. Since 5.34 shows a high degree of coupling between the two dimensional simulated and measured fields, this implies that the simulations carried out on the horn would also predict a high degree of symmetry in the main beam.

Examining the cross-polar fields, the agreement between the measured and simulated fields is not as good, as was expected. The cross-polar field is naturally difficult to measure due to the inherently low levels involved (in this case for example $<$ 50dB) and the fact that it is difficult to be certain that the measured pattern is indeed purely the cross-polar field due to potential misalignments within the measurement system. In contrast, it is straightforward to extract only the cross-polar field from the simulation model by virtue of considering only x or y components of the modes, depending on the polarisation of the simulated source. In the practical setup however, this is not easily achievable. Co-polar components can 'leak' into the cross-polar field and can also be an inherent inseparable part of the cross-polar field, as previously noted. Despite this, the agreement between the simulated and measured fields is certainly satisfactory in terms of the levels being predicted and the general beam patterns. The disagreements lie mainly with the depth of the on-axis null and the increased power just off-axis on the left of each cut.

In an attempt to reconcile the measured and simulated cross-polar patterns, namely to approximate the measured central null and the area of enhanced power, the gamma parameter in MODAL was adjusted. This parameter controls the relative angle between the probe waveguide and the horn, in other words it represents an azimuthal rotation between the horn and probe. Figure 5.35 shows how changing gamma alters the cross-polar field. Altering gamma

means that the measurement (both real and simulated) is no longer exclusively looking at the cross-polarisation but is rather measuring a combination of co and cross-polar components. Since the co-polar is dominant and has a high on-axis power level, this results in the on-axis null being eroded with on-axis power appearing. Increasing gamma allows a greater co-polar component to enter the measurement, and so the observed increase in on-axis power is consistent with this. It is clear that altering gamma in the simulation does not account for the bump observed in the experimental measurements, as was previously thought. To this end, the impacts of α and β were examined, which correspond to rotations of the horn about the x and y axes respectively.

Figure 5.36 shows how the cross-polar field responds to changes in α and β that are of the order of the misalignments that may occur in the real system. Physically, it would be

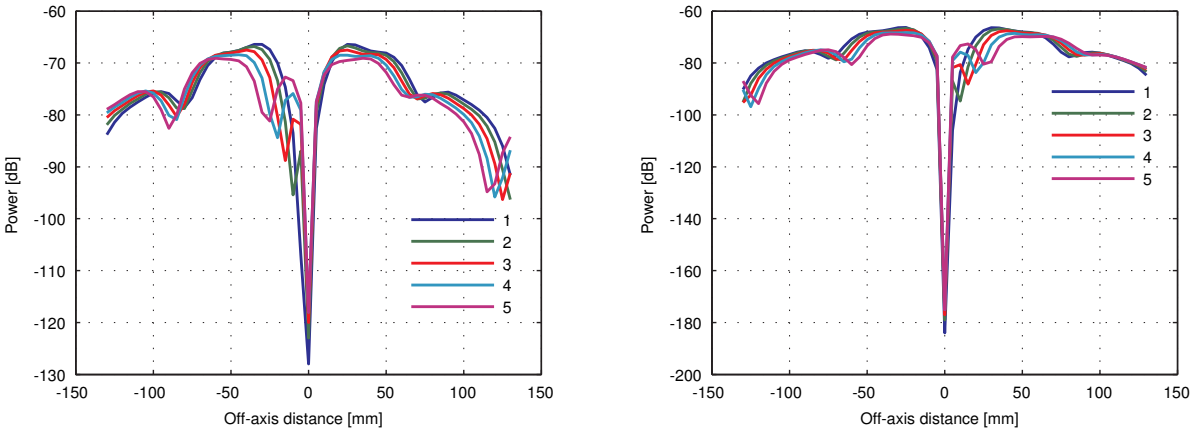

(a) Variation of α

(b) Variation of β

Figure 5.36: Simulated variation of cross-polar field with α and β, for α and $\beta \in [1°, 5°]$

expected that such tilts about the x and y axes would result in a projection effect appearing in the measured data which would manifest itself as an asymmetry in the measured data, with the lobes on one side of the axis showing altered levels and spatial extent, when compared to the corresponding lobes on the opposite side of the optical axis. This is what is observed in figure 5.36. Since this has the effect of making a bump appear on one side of the pattern, these parameters were investigated in order to reconcile the measured and simulated data.

Various combinations of α and β were considered. Setting $\alpha = 10°$ and $\beta = 5°$ yielded the best agreement. The simulated cross-polar field for these parameters is compared with the measured field in figure 5.37. Although changing gamma increases on-axis power levels, it does not reconcile the difference noted and as such it was set to zero. The difference between the measured lobes either side of the optical axis is approximately 7 dB, and a very similar difference was noted in the simulated case, using the indicated parameter values. This was obtained largely by variation of α. The off-axis shift of the beam pattern was obtained by adjusting β. The value of 10° used for α that gave rise to this agreement is quite large,

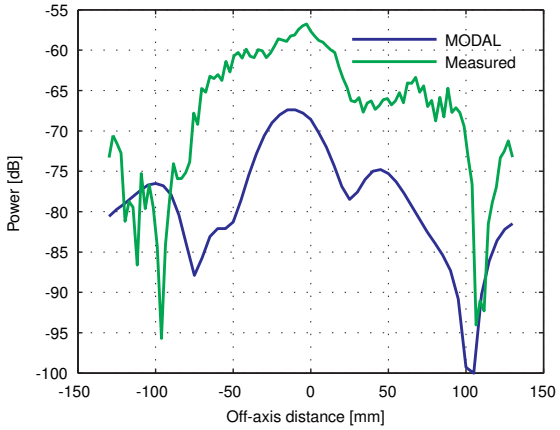

Figure 5.37: Comparison of simulated ($\alpha = 10°$ and $\beta = 5°$) and measured cross-polar fields (45° cuts)

certainly larger than any tilt that was present in the measurement system. This implies that although misalignments relative to these axes contribute to the observed beam pattern, the cause of the bump is still likely to largely be due to a component of the co-polar field leaking into the measurement, as opposed to being purely based on projection effects. It is possible that tilts or rotations of the waveguide probe , particularly when the waveguide twist section is used, is likely to contribute to the observed pattern and it is not yet possible to investigate these effects.

As stated previously, two horns were manufactured by RAL as part of this project. Each followed the same design, the difference being the material from which the horn was constructed. The measurements outlined above relate to the aluminium horn. The second horn was etched from copper. Due to laboratory availability, a full measurement campaign using this horn on the optimised testbed was not possible, however the copper horn was measured using an earlier, less well aligned iteration of the testbed. Figure 5.3.3 gives a comparison of the E and H planes of the co-polar field of the two horns at propagation distances of 53 and 213 mm. Excellent agreement is observed at these distances and at other propagation distances that were also considered. The performance is almost identical, as expected. Although no large difference was expected due to the different materials, it was possible that slightly different results may be obtained due to the potential for the geometries to be different owing to the different manufacturing methods and the resulting differences in tolerances. Since no appreciable differences were observed, this shows consistency between the manufacturing methods. It is planned to perform a comprehensive analysis of the second horn in the future in order to evaluate the cross-polar performance of the horn in addition to the performance of the other metrics of interest, however on the basis of these initial field pattern comparisons it is expected that the two horns will perform very similarly.

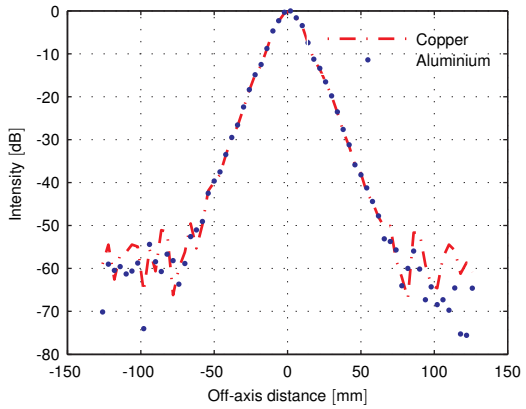

(a) E-Planes @ 53 mm

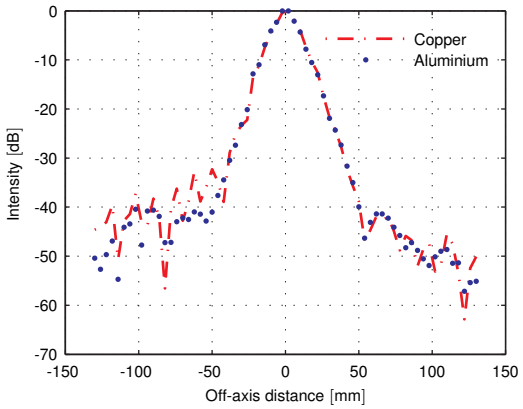

(b) H-Planes @ 53 mm

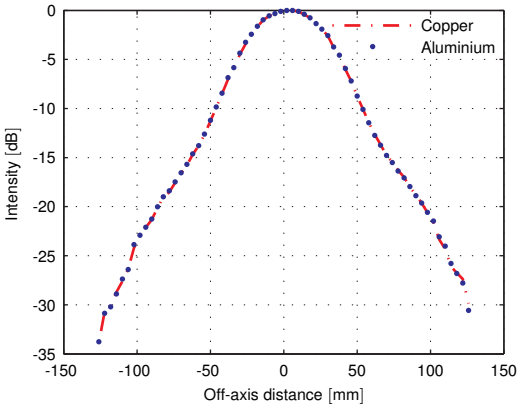

(c) E-Planes @ 213 mm

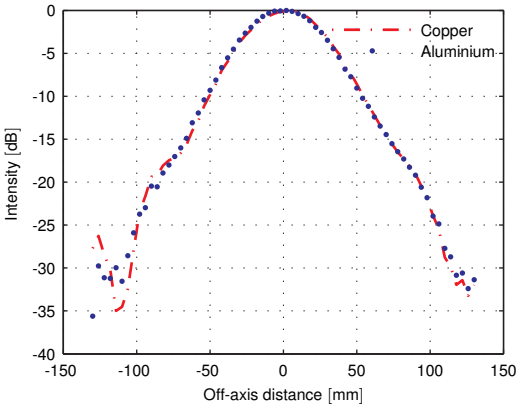

(d) H-Planes @ 213 mm

Figure 5.38: Comparison of E and H planes for the two horns at different distances at 100 GHz

5.4 Conclusions

In this chapter, the genetic algorithm based optimisation method that was presented in the previous chapter was used to design a horn antenna for a future CMB mission (the same process could also be readily applied to far-IR missions), subject to demanding performance requirements. The performance of the horn was found to be comparable to that of a standard corrugated horn with respect to the parameters of interest which is an excellent result considering the fact that the optimised horn is smooth-walled and customised for spot frequency operation. The resulting optimised horn geometry was manufactured by RAL and subjected to a rigorous measurement campaign. The measurements were compared to predictions made using CST, MODAL and SCATTER. Excellent agreement was observed between the simulated and measured data, indicating that the horn is operating as intended. The co-polar field was found to remain highly symmetric across the 75-110 GHz band, approximately the same band as required by the design brief. In this range the sidelobe levels also remained low and the beam retained a Gaussian distribution. With propagation distance the horn also behaved as expected, implying no unexpected modal contributions and largely single-moded performance. The cross-polar field retained a low level relative to the co-polar throughout, particularly at the design frequency of 100 GHz, as required. Disagreements were observed when the S_{11} and cross-polar ratio were measured across the band, however the trends still agree with predictions and the actual performance still yielded excellent results, even if they did not precisely match those predicted. The main factors in the disagreements are likely to be the fact that rectangular to circular transition is not included in the MODAL and SCATTER simulations, and also potential misalignments within the physical measurement system.

The design process was required to use a standard desktop computer and to take a reasonable amount of time, of the order of hours, without using any special design tricks. This goal was achieved, with the design process taking only 6 hours. One way the computational time was kept down was by carrying out the design at the central frequency of the band, 100 GHz. This meant that the horn was tailored for operation at precisely 100 GHz and with respect to one metric, thus achieving this performance at the expense of other metrics and frequencies. The fact that all of the metrics of interest perform excellently across the band despite this is an extremely positive result and validates the approach, showing that such challenging design briefs can be met once efficient modelling techniques are employed.

Chapter 6

Simulation of Off-Axis Illumination

In this chapter, the modification of SCATTER to allow for horn illumination by an off-axis plane wave will be discussed. In the standard implementation of SCATTER it is assumed that all modes being considered in the calculation (all radial modes in each azimuthal order) are excited equally with unity power. This represents the illumination of the horn with a source that effectively acts like a blackbody in a cavity, i.e. an incoherent source. All modes are thus excited, with the radius of the smallest section of the horn (typically the throat section) determining the actual modes that can propagate within the horn, which will form a subset of the modes initially excited. The modification presented in this chapter replaces this incoherent source with a plane wave that can be incident on the input at some user defined angle. This allows for the off-axis excitation (illumination) of the horn in question. Rather than exciting all modes, this will result in the excitation of only the modes supported by the horn. More correctly, it will result in the excitation of a subset of the modes supported by the horn, as even on-axis illumination will not necessarily excite all of the supported modes. Additionally, the excited modes are not excited with unity power but rather with a varying amount of power which depends on the coupling of the plane wave to each mode for a given angle of illumination.

The immediate motivation for this work is related to the measurement and analysis of multi-moded horns. It is highly desirable to be able to measure the various modes present in a horn individually so as to be able to fully characterise and understand the performance of multi-moded horns. One possible way of accomplishing this goal, which will be the approach adopted in this thesis, is to illuminate a back to back horn at an angle, using a plane wave initially. The back to back horn is simply two versions of the horn that is to be measured, mounted together by the throat filter sections. Off-axis illumination of this system with a plane wave could, in theory, result in the excitation of a subset of the modes supported by the horn, with the modal content of the horn defined by the angle of illumination. In this way, individual modes could be measured which would improve the understanding and testing of multi-moded systems, which has clear benefits going forward for the design of multi-moded systems for future astronomical missions. Other applications for this technology could include encoding

different information within different modes and so illuminating the horn at different angles and measuring the resultant beam patterns accurately could lead to an effective 'multiplexing' system, allowing for the transfer of larger amounts of information in a given amount of time.

6.1 Background Theory

In this section, the background theory relating to the implementation of off-axis plane wave illumination in SCATTER will be discussed. The overall principle is that the system is processed as normal in SCATTER and a singular value decomposition (SVD) analysis is carried out on the resulting transmission scattering matrix, determining the modes or power channels that are present in the system (back to back horns). The orthogonal polarisations of the waveguide modes are included in this. Since back to back horns are being considered in this case, it is these singular modes that the off-axis plane wave will couple to when the horn is being excited at the aperture at one end, and so they will determine the power ultimately entering the system. Therefore, each of the singular modes are considered separately, and their coupling to an angular plane wave is found. This is considered to be the 'amount' of that singular mode that is excited at the aperture of the back to back horn, and these scaled modes are then used as the initial input matrix to the horn in SCATTER, allowing the new output fields to be found. A description of the plane wave will now be given, followed by its integration with the process described above.

6.1.1 Implementation of the Off-Axis Plane Wave

SCATTER assumes that all modes that are to be considered in each azimuthal order (the number of modes and the number of azimuthal orders are controlled by the user) are excited in equal amounts (with unity power) at the horn throat, (or more generally the portion of the horn that is being excited). The modes that do not propagate are then filtered out by the initial section of waveguide, forming the excitation port of the horn. The model that will be used to excite the horn with off-axis plane waves initially assumes that a plane wave is present at the horn aperture The implementation of the plane wave in SCATTER, as a function of incidence angle, is now given.

Figure 6.1 displays the phase front off an off-axis plane wave as it propagates. The phase front of a plane wave is flat across the span of the wave and impinges on a flat plane that represents the aperture of the horn, the part that is being excited by the plane wave. For clarity, the diagram shows the plane wave propagating beyond this, as in the off-axis case it is the waves corresponding to these subsequent phase fronts that excite the horn at greater off-axis distances. The plane wave is defined such that the wave which crosses the throat at the centre is spread equally on the input and output sides of the throat plane. Figure 6.2 displays the situation in more detail. The plane wave is assumed to be an angle θ off-axis,

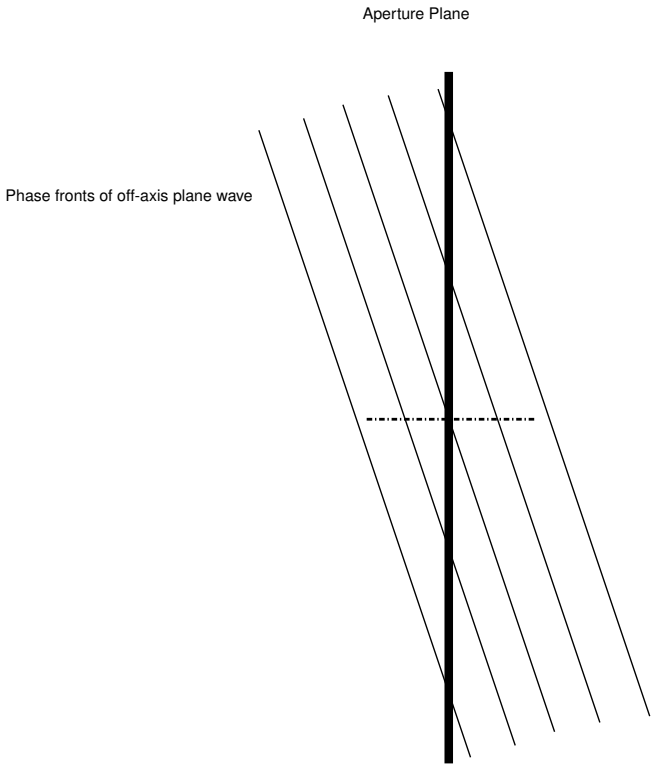

Figure 6.1: An off-axis plane wave propagating to and beyond the throat plane

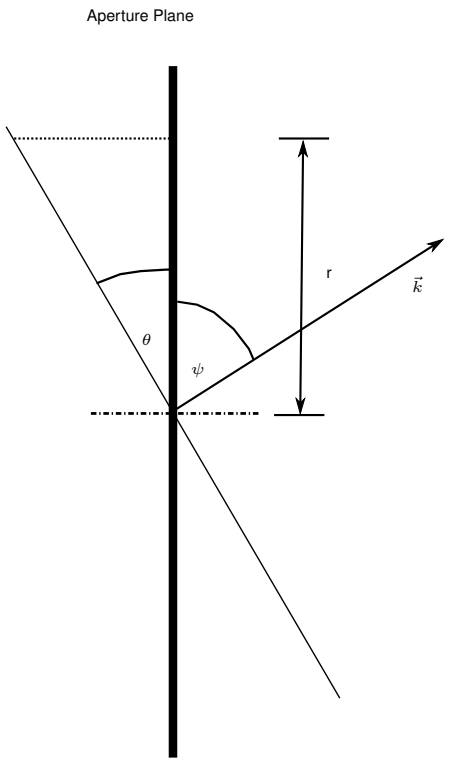

Figure 6.2: Off-axis plane wave schematic

where setting this angle to zero results in on-axis illumination of the aperture by the plane wave. This angle is known as the off-axis angle. This is the standard normal illumination, whereby the propagation vector of the plane wave, $\vec{k}$, is normal to the plane of the aperture, or parallel to the normal of the plane. The off-axis plane wave is implemented by assuming that an on-axis plane wave would have an arbitrary constant phase across the aperture, in this case zero, and that an off-axis plane wave would have this phase difference at the centre of the throat, but the phase would lag or lead that of the on-axis wave at all other points on the aperture plane, depending on the off-axis angle. Thus, for each point sampled in the aperture plane, the phase difference between the on-axis and off-axis plane waves must be found. This can be achieved by two methods:

1. Geometric optical path difference
2. Plane overlap phase calculation

In either case, the phase difference between the two waves is found, which must be added to the initial phase of the wave. Since this is completely determined by the location of the source of the wave, which is completely arbitrary, this initial value is also arbitrary, and so can be assumed to be zero. Thus, the phase of the off-axis wave in the aperture plane is simply equal to the phase difference between it and the on-axis wave. The geometric optical path difference calculation is now examined.

Optical path difference

By examining figure 6.2, the optical path difference (OPD) can be readily calculated. In the aperture plane, each of the sampled points is a distance r off-axis. For each of these points, a line is traced back to intercept the illustrated phase front normally. This line represents the propagation of the phase front (at that point) to the horn aperture plane. This line represents the additional distance that the off-axis wave must travel to reach that point on the aperture, relative to an on-axis wave. The length of this line is thus the OPD, δ, of the two waves, and so the phase difference between the two waves, $\Delta\phi$, can be represented as

$$\Delta\phi = \exp\left(jk\delta\right), \tag{6.1}$$

where k is the wavenumber of the plane wave (wavelength λ), given by $k = \frac{2\pi}{\lambda}$. Again, this figure is the phase of the off-axis plane wave at the aperture, as a zero initial phase is assumed, such that the phase would be zero in the aperture plane for an on-axis wave. Using the sine relation, it is clear that $\sin\theta = \frac{\delta}{r}$, and therefore that

$$\delta = r\sin\theta. \tag{6.2}$$

This gives the phase difference (i.e. phase of the off-axis plane wave) at a point in the throat plane a distance r off-axis as

$$\Delta\phi\left(r\right) = \exp\left(jkr\sin\theta\right). \tag{6.3}$$

The phase difference is now calculated using an analytical analysis of the system, as opposed to the geometric argument presented above.

Plane overlap phase calculation

In this section, the phase difference is calculated using an analytical approach. The phase difference is calculated by finding the phase along the aperture of the horn as before, however this time it is achieved by calculating the dot product of the propagation vector of the plane wave and the position vector of the point on the aperture plane under consideration. This is again consistent with the phase difference at the centre of the aperture being zero. The position vector has a magnitude equal to the off-axis distance of the point under consideration, r, and a direction along the plane. It is represented by $\vec{r}$. The phase difference is thus given by

$$\Delta\phi = \exp\left(j\left(\vec{k}\cdot\vec{r}\right)\right). \tag{6.4}$$

Examining figure 6.2, it can be seen that $\vec{r}$ makes an angle ψ with $\vec{k}$. Thus, the dot product is given by $\vec{k}\cdot\vec{r} = kr\cos\psi$. The angle between the propagation vector and the direction parallel to the phase front is, by definition, 90°, and it is given by the sum of θ and ψ. Thus, $\theta + \psi = 90°$, and so $\psi = 90° - \theta$. This allows equation 6.4 to be rewritten as

$$\Delta\phi = \exp\left(j\left(kr\cos\left(90° - \theta\right)\right)\right). \tag{6.5}$$

However, it is known that $\cos\left(90° - \theta\right)$ is equal to $\sin\left(\theta\right)$, and so the phase of the plane wave at a point in the aperture plane that is a distance r off-axis, is given by $\vec{r}$, and can be written as

$$\Delta\phi = \exp\left(jkr\sin\theta\right), \tag{6.6}$$

which is equal to equation 6.3, which is the corresponding expression obtained by using a geometric argument.

It is clear from this expression that for an off-axis angle of zero, the phase difference factor is one for all points on the throat. This corresponds to a phase difference of zero across the throat, as required by the assumed conditions, and so the model is consistent.

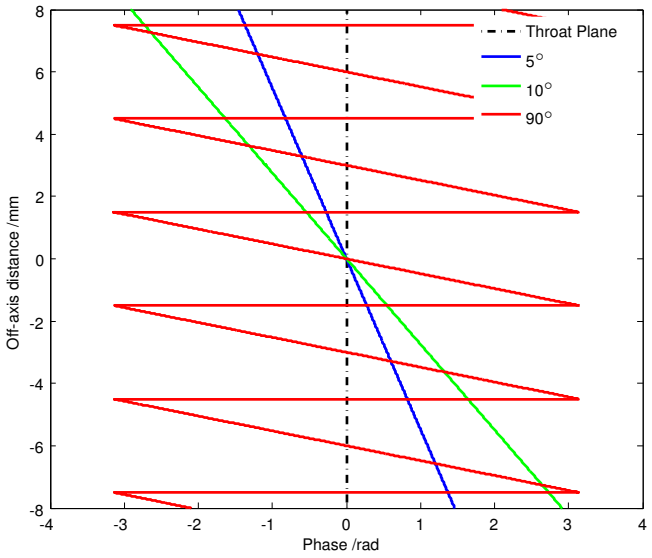

Figure 6.3: Phase fronts of generated off-axis plane waves

6.1.2 Overall Plane Wave

Assuming a zero phase for the wave front shown in figure 6.2, the off-axis wave can be represented by the expression

$$E_{PW}(\theta, r) = A \exp(jkr \sin\theta), \tag{6.7}$$

where A is the amplitude of the wave at each point. For a plane wave this is a constant. The power is normalised by choosing a value of A such that

$$\int_{\text{Aperture}} E_{PW}(\theta, r)\, E_{PW}(\theta, r)^* \, dA = 1. \tag{6.8}$$

By taking the argument of equation 6.6, the plane wave can be visualised. Figure 6.3 illustrates off-axis plane waves relative to the aperture plane for a variety of off-axis angles. When zero degrees is set, the plane wave is coincident with the aperture plane, as required, validating the approach. For other angles, the plane wave is seen to be slanting relative to the aperture plane, as expected. Beyond ten degrees the phase begins to make 2π jumps due to its sinusoidal nature. This is illustrated by the 90° wave. A phase smoothing algorithm could be used to smooth this jump out in order to illustrate the plane wave in a clearer fashion, however this is not carried out in order to demonstrate the emergence of this effect, and this is the form which the phase must be left in for further calculations.

The plane wave will be used in overlap integral calculations with the aperture fields calculated in SCATTER, and so to match the formats it is necessary to convert from the polar form of equation 6.7 to complex plane representation. This is accomplished by invoking Euler's

formula ($e^{jx} = \cos x + j \sin x$), such that the field can be represented as

$$E_{PW}(\theta, r) = A\left(\cos\left(kr\sin\theta\right) + j\sin\left(kr\sin\theta\right)\right). \tag{6.9}$$

Using this form of the off-axis plane wave, off-axis illumination of a horn can be simulated in SCATTER.

6.1.3 Implementation of Off-Axis Illumination

The off-axis illumination of the horn depends greatly on the singular value decomposition (SVD) method of reconstructing fields that was discussed previously in section 3.2. Using the SVD method, the overall field patterns and modal content of a horn can be represented as a sum of independent (singular) modes (numbering significantly less than the number of modes required if the raw output scattering matrices are used). Each of these modes is an independent power carrying channel in the horn that does not mix with the other singular modes. They are added together in the same way that modes of different azimuthal orders are added in SCATTER (incoherently) to recover the overall field. Since a back to back horn is being considered, the throat of the horn is not being excited, but rather the aperture of the horn, as assumed in section 6.1.1. Since an actual wave is now being used to illuminate the horn (rather than simply illuminating each mode at the aperture with unity power, similar to the situation that would occur if the source were a blackbody cavity), the coupling of the illuminating wave, the off-axis plane wave, to the aperture field must be considered. It would be possible to simply consider the aperture field, however since individual modes are to be measured, it makes sense to consider the singular modes as these are the actual modes that propagate through the horn.

It is therefore necessary to perform an overlap integral between the off-axis plane wave and each of the modes supported in the horn that were obtained using an SVD analysis of the horn. This overlap integral will return a coupling coefficient which represents how well the plane wave and the mode couple. Thus, if the coupling is low, then the plane wave will excite little of that mode at the aperture of the back to back horn, and so it will not propagate within the horn in any appreciable amount. The coupling will vary as a function of excitation angle, and so as the angle changes, different amounts of each mode will be excited at the aperture, resulting in different modes propagating through the horn, resulting in different output fields, depending on the excitation angle. The coupling of the plane wave to the i^{th} singular field, $U_i(r)$ is given by

$$C_i = \int E_{PW}(\theta, r)\, U_i(r)^*\, dr. \tag{6.10}$$

SCATTER is run using a back to back version of the horn that is being analysed, and the aperture field calculated using the SVD method. The coupling between each of the modes and the plane wave is then calculated, and the original modes multiplied by their corresponding

coupling factor, effectively exciting the mode by the correct amount at the aperture. Since the amount of each singular field present is varied as a function of angle, the overall field at the aperture varies with excitation angle, resulting in the variation of the input excitation field to SCATTER.

As previously described, SCATTER proceeds by splitting the horn into many sections and considering each junction. When moving to a subsequent junction, it takes the input to this section to be the output of the previous section, in other words the S_{21} matrix for the previous section. Since for the first section there is no previous section, the input scattering matrix is assumed to contain each of the waveguide modes being considered in equal amounts. The radius of the first section filters out the modes that cannot be supported in a waveguide of that radius at the frequency in question. The remaining modes then propagate through the horn. Using the scaled singular modes, the input scattering matrix for the first section of the back to back horn must be found, as will now be described.

For a given section, the transmission scattering matrix, S_{21}, can be represented as the product of the three matrices U, Σ and V such that $S_{21} = U\Sigma V^{\dagger}$. This is the SVD representation of the scattering matrix. Columns of U which have a non-zero associated singular value form a basis set for the field on the output side of the section in question, while columns of V form a basis set for the field on the input side, with the diagonal entries of the Σ matrix control the scattering between the modes, as described in detail in section 3.2. This ensures that the modes remain independent of each other and do not mix, i.e. the singular modes are incoherent. This representation is used in the off-axis illumination step.

The overall S_{21} matrix of the (back to back) horn is calculated in SCATTER and an SVD is carried out on this matrix. This allows the V matrix for the horn to be found, which gives the field at the input of the horn, i.e. the initial conditions that SCATTER places on the horn. These initial conditions are used as the overall source in the off-axis mechanism. This initial (on-axis) field is assumed to be present prior to the horn, however between the field and the horn aperture an additional section is inserted of zero length, and this section facilitates the simulation of the off-axis illumination due to the plane wave.

The overall affect of the off-axis plane wave is simply that depending on the angle, it couples with a varying amount to the various modes present at the aperture of the horn, due to phase differences. Thus, after illumination by the off-axis plane wave, the modes present at the aperture are still present with the same phase and shape (assuming that the coupling of a given mode to the plane wave is non-zero), it is simply the amplitude of these fields that changes with varying illumination angle. This change in amplitude is modelled by inserting the section described above. The fields are scattered into the required smaller amplitudes inside this equivalent section to facilitate calculation. Critically, the section is of zero length so as it does not alter the phase of the modes. In SCATTER, the scattering between modes

is calculated, and then a phase term is also added to each mode in order to account for the phase shift in the modes due to propagation along the length of the section. If the length of the section is assumed to be zero, then the required conversion occurs, but without any phase shift or additional inter modal scattering as there is no change of radius and the modes are incoherent so no inter modal scattering is permitted in any case. Since the section in question here is simply for rescaling the various singular fields present and is not associated with any physical section or propagation, the section is assumed to be of length zero. The amplitude conversion of the modes is governed by the scattering matrices for that section, which are calculated as follows.

The input to this zero length section is taken to be the V matrix of the original horn once an SVD is performed on its S_{21} matrix. This keeps the initial conditions the same as in the plane wave case, allowing comparisons to be made between results. The output field of this section is simply given by the U matrix of the original horn, where each of the columns corresponding to propagating modes is scaled by the coefficient that gives its coupling to the off-axis plane wave, as outlined previously. All other columns are set to zero, thereby not exciting modes which are not supported. This new matrix is known as U_{new}. The purpose of this section therefore, is to scatter the initial modes present, given by V, into the modes given by the propagating columns in U_{new}. Therefore the singular value matrix, Σ_{new}, must be found that relates these two columns, as this will then allow the scattering matrix for this section to be found and used to excite the system. Equation 6.11 gives this new scattering matrix in terms of its SVD components.

$$S_{21}^{\text{new}} = U_{new}.\Sigma_{new}.V_{new}^{\dagger}. \tag{6.11}$$

U_{new} and V_{new} represent the output and input mode distributions respectively. They represent the distribution of these modes, with the amount of each mode at the output related to the amount excited at the input by the scattering matrix Σ_{new}. In off-axis illumination, the relative amount of each mode is all that changes, not the modal field distribution. Thus $U_{new} = U$ and $V_{new} = V$. Of the modes that are supported in the system, the coupling of each of those modes to the plane wave is given by the value determined by equation 6.10. Since for this section, the singular value of each mode is unity as no scattering occurs due to the fact that it is assumed to be of uniform radius (and so will transmit through the section with the input amplitude equal to the output amplitude), all non-propagating modes may retain this singular value and will be cut-off when they enter the rest of the system. The β modes that do propagate however must be scaled relative to unity in order that they exit this section in the proportions excited by the plane wave. Thus $\Sigma_{new} = C_i\Sigma_i, \forall i \leq \beta$. This matrix can then be combined with U_{new} and V_{new} to evaluate the new scattering matrix given by equation 6.11. In the special case where the singular modes are not a composite of several waveguide modes but rather represent one mode. For example in the case of a horn which supports two waveguide modes, each in different azimuthal orders, then the above need not be applied. In this case, the scattering

matrix will be a diagonal matrix with the diagonal entries being unity. The relevant modes along this diagonal can be directly scaled by the corresponding C_i and this matrix then used as the excitation source.

The new scattering matrix is passed to the standard SCATTER code with one modification, that is to accept this matrix as the initial excitation source, i.e. the initial S_{21} matrix. Where multiple azimuthal orders are to be considered, the new scattering matrix is calculated for each order, and the overall multi-dimensional matrix is passed to SCATTER, which then proceeds with the calculation as normal. In this way, off-axis excitation is simulated in SCATTER, starting at the aperture of the horn.

6.2 Examples of Off-Axis Illumination

In this section, the code that was developed based on the approach outlined above will be applied to the following two examples.

- Back to back horn, considering one azimuthal order (single-moded)
- Back to back horn, considering six azimuthal orders (few-moded)

The horn in question, paired back to back with an identical horn, is the corrugated horn described in section 3.2.3. At 100 GHz, assuming that all modes are excited, this horn is few-moded, containing power in azimuthal orders $n = 0, 1, 2$. Corrugated horns actually support hybrid waveguide modes rather than the standard TE and TM waveguide modes. The mode supported by this corrugated horn in the $n = 1$ azimuthal order is known as the HE_{11} mode, [46]. Such modes can be expressed as linear combinations of the standard waveguide modes. The HE_{11} mode can be approximated to consist of $0.85TE_{11}$ and $0.15TM_{11}$. Thus, when the scattering matrices for this horn are examined, power will be scattered into the standard waveguide modes in such a way that this mode is formed, and if an SVD were calculated and the resulting single mode plotted, it would be a combination of the TE_{11} and TM_{11} modes, i.e. it would be the HE_{11} mode. In the $n = 1$ azimuthal order in this horn, only the TE_{11} mode is supported in the throat section. As the radius increases along the flared section of the horn, higher order modes are supported and this effectively allows the HE_{11} to be formed. The HE_{11} mode is the only mode supported by the horn in this azimuthal order due to the fact that the TE_{11} mode is the only mode supported in the throat in the $n = 1$ azimuthal order. The two modes are therefore associated. Although the hybrid and standard waveguide modes are not interchangeable, since the TE_{11} mode gives rise to the HE_{11} mode, in this thesis the standard waveguide modes will be used when discussing the modes supported by the horn in the throat (as these give rise to the associated hybrid modes), however it should be remembered that physically the hybrid modes are propagating. Thus when the TE_{11} mode is referred to, the resulting excited hybrid mode, the HE_{11} mode is what is physically in question. This is done

to establish consistency in the examples presented, as in chapter seven a smooth-walled horn is considered in detail.

The modes supported in each of these azimuthal orders in the throat are, respectively, TM_{01}, TE_{11} and TE_{21}. These modes can scatter into other modes along the length of the horn, forming the hybrid modes. The TE_{11} mode is the fundamental mode of the system and so the associated hybrid mode is the first one to be activated, followed in sequence by the associated hybrid modes excited by the TM_{01} and TE_{21} modes. The horn has a throat radius of 1.535 mm, which is what defines this modal content. Although at the excitation and exit ports of the back to system (radius 7.1345 mm) several more modes are supported, due to the joined throat filter section that occurs when the horns are joined together, only the indicated modes are supported in the horn, and so only these modes (or more correctly their associated hybrid modes) need be coupled to the plane wave. The modes obtained from the SVD analysis represent these hybrid modes, and so coupling these to the plane wave allows them to be excited in the correct proportions.

6.2.1 Back to Back Horn, Run Single-Moded

In this section the back to back corrugated horn system will be analysed at 100 GHz, however in order to clearly observed the effects of off-axis illumination, it will be considered to be single-moded, i.e. only the $n = 1$ azimuthal order will be considered. In this azimuthal order, only the HE_{11} hybrid mode associated with the TE_{11} mode is supported in the throat, and so there will only be one mode retrieved from the SVD analysis, along with its orthogonal polarisation. Physically, the modes from the other azimuthal orders still exist, however they are independent from $n = 1$, as no mixing can occur between azimuthal orders. To this end, they can be ignored at this point as their presence or absence has no bearing on the performance of the TE_{11} mode that is supported in the throat.

The standard geometry file for this horn is passed to SCATTER, with the back to back configuration realised within the modified code. Once the length and radius arrays are read in from the standard horn geometry file, the arrays themselves are modified so as a mirror image of the horn is produced and attached to the original horn, throat to throat. The excitation aperture of this new structure is now the exit aperture of the horn defined in the geometry file. Following this is the flared portion of the horn, tapering towards the throat. The two throats are then joined together, and modelled as one continuous section. Following this, the flared portion of the other horn forming the back to back arrangement is added, culminating in the exit aperture of the second horn. A profile view of such an arrangement for an arbitrary smooth-walled horn is shown in figure 6.4. The number of sections in the back to back horn is thus $2\,(\text{sections} - 1)$. These new arrays for the horn lengths and section radii are then passed to the mode matching portion of the code and the calculation proceeds completely as normal, as though they were read in directly from a geometry file.

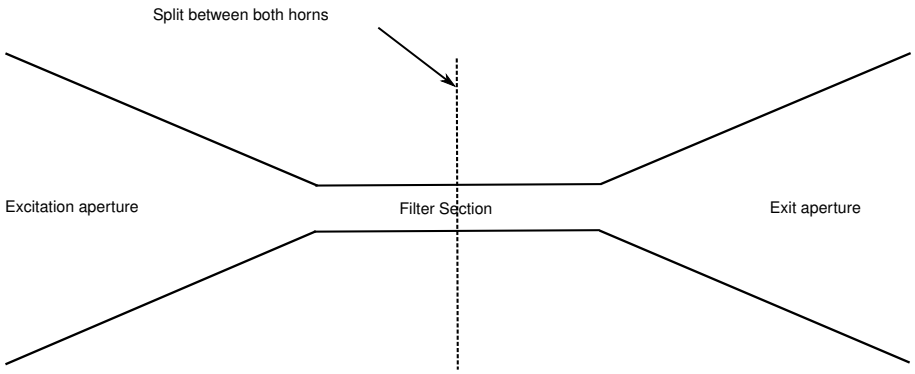

Figure 6.4: Schematic of two identical smooth-walled horns arranged in the back to back configuration outlined in the text

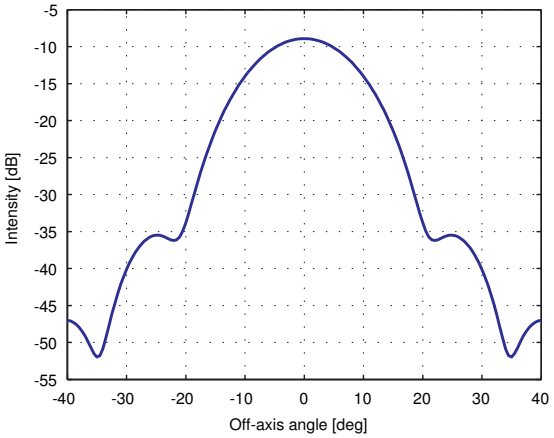

Figure 6.5: Farfield (E-plane) of the mode of azimuthal order $n = 1$, supported by the back to back system, the HE_{11} mode which is a linear combination of the TE_{11} and TM_{11} modes.

This geometry was processed using a frequency of 100 GHz and considering only $n = 1$, which again, in this case, means only the TE_{11} mode was supported in the throat. As outlined previously, the standard SCATTER code is used to calculate the transmission scattering matrix for the back to back horn, and an SVD is carried out on the overall resulting matrix. It was found that there was only one singular value with a value of greater than 10^{-6}, the threshold for which a singular value is considered non-zero in this code. Plotting the farfield (E-field, but as it is a hybrid mode, this is symmetric with the H-plane) of this mode, as in figure 6.5, reveals that, as expected, it is the HE_{11} mode, which consists of the modes that are expected to be supported by the system at 100 GHz. Using the first column of the U matrix associated with the SVD of the S_{21} matrix, the (singular) mode was reconstructed in the complex form $a + ib$. Typically, these modes may not correspond to a given waveguide mode but rather a linear combination of several. In this case, since only the TE_{11} mode is supported, then the mode does correspond directly to this waveguide mode. For a given off-axis illumination angle, the appropriate off-axis plane wave was constructed, normalised so as to contain unity power. Due to the nature of the phase calculation, this was initially calculated in complex

polar form, however after this it was transformed into the complex planar form in common with the singular mode.

The singular mode was normalised to unity power and a complex coupling integral was carried out with the off-axis plane wave. The mode was then scaled by its corresponding plane wave coupling factor and re-inserted into the overall U matrix. The new U matrix was then constructed by zeroing all columns that had a zero singular value previously. This was then used, along with the appropriate input field matrix, V, to find the new scattering matrix that was used as the initial excitation to the SCATTER calculation which then proceeded as normal, as outlined in the previous section. This effectively simulates the off-axis illumination of the back to back horn.

When the angle of illumination, the off-axis angle, was varied between 0° and 60° in steps of 10°, the output composite farfields were calculated and are shown in figure 6.6. These were calculated by using the single output mode as a basis set for expanding the farfield form of the waveguide modes, including the orthogonal polarisations, as outlined previously when the application of SVD to field calculations was discussed in section 3.2. It can easily be seen

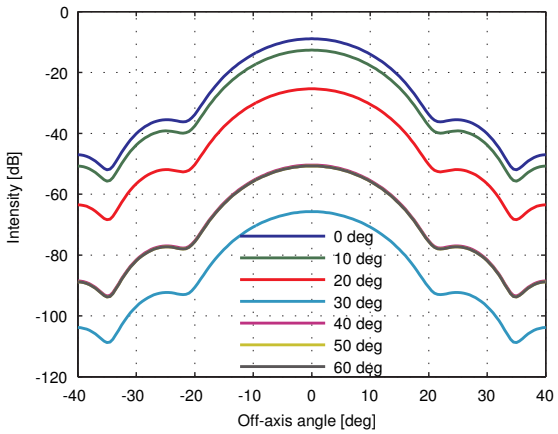

Figure 6.6: Response of the farfield of the single-moded back to back corrugated horn system for various off-axis angles of illumination

that as the off-axis illumination angle varies, the amplitude of the farfield pattern changes, however the field distributions remain very similar in profile. This follows directly from the fact that there is only one mode supported within the horn at this frequency when only the $n = 1$ azimuthal order is considered, and so although it will be excited by different amounts as a function of illumination and its magnitude will resultantly vary, the distribution will remain the same, as it is the same mode being excited in each case. The amount by which this mode is excited is determined by the degree of coupling between the excitation plane wave and the distribution of this mode at the horn aperture. Since the back to back system is symmetric, the output field will retain the same modal content, merely having its amplitude scaled by a proportionate amount.

Of note is that the pattern for an illumination angle of 30 ° is of a lower level than that for higher angles of illumination. This can be explained by considering figure 6.7, which shows the coupling of the off-axis plane wave to the singular mode as a function of illumination angle. The trend in the coupling between the plane wave and the mode supported by the system is

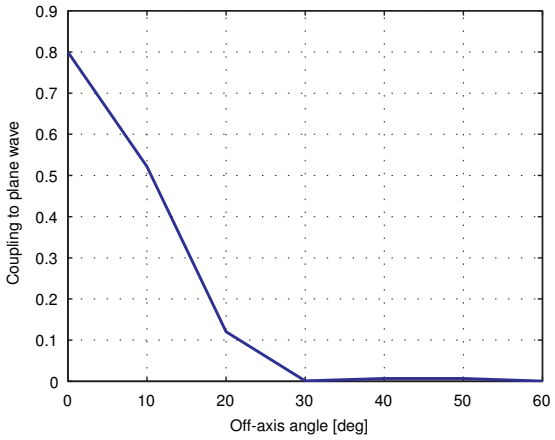

Figure 6.7: Coupling of the plane wave to the single-moded back to back corrugated horn system for various off-axis angles of illumination

as expected. The coupling is high for an on-axis plane wave, and then steadily drops off as the angle of illumination increases. At 30 degrees the coupling is approximately 0. This is reflected in the corresponding beam patterns shown in figure 6.6, where the intensity of the beam pattern drops off with increasing angle of illumination, up until 30 degrees. A moderate increase in coupling to slightly above zero is noted beyond 30 degrees. This is due to the fact that the coupling of the plane wave is now increased, due to the fact that it is now incident on the sidelobes present in the field distribution of the mode. This is again reflected in the beam patterns for these angles which show elevated levels relative to 30 degrees, however these levels are of the order of -55 dB and so are effectively zero as would be expected for high angles of illumination. Angels of illumination beyond 60 degrees were not considered for clarity of illustration and also as such angles would be difficult to experimentally achieve.

6.2.2 Back to Back Horn, Run Few-Moded

The horn geometry used in the previous section was examined again at 100 GHz, but with six azimuthal orders considered ($n \in [0, 5]$). At 100 GHz, the radius of the throat (filter) section is sufficient for the horn to support the TM_{01} and TE_{21} modes in addition to the TE_{11} mode. The modes supported for the different azimuthal orders will excite their associated hybrid modes along the length of the horn. The farfields (E-modes) of the modes supported in this system are shown in figure 6.8a, where again the orthogonal polarisations are included, as they are excited by an unpolarised plane wave. Figure 6.8b shows the analytic form of the farfield distribution of the farfield forms of the TM_{01} and TE_{21}, normalised as in SCATTER, outlined in chapter 2.

Calculating the S_{21} matrix for each of the azimuthal orders and performing an SVD on each of

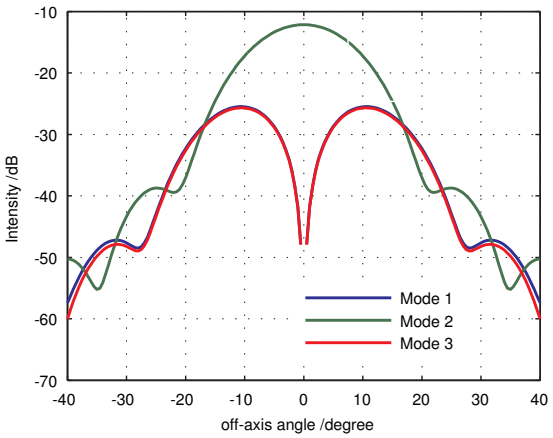

(a) Modes supported by system, from SVD

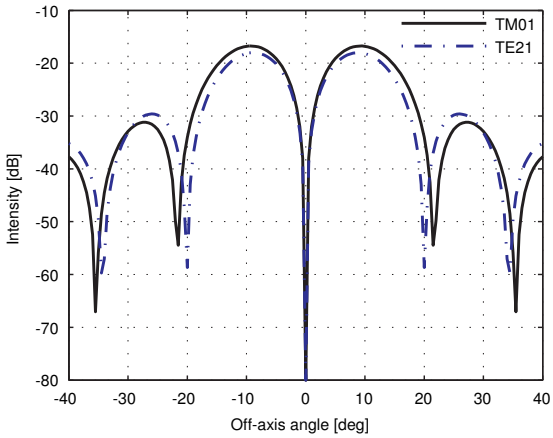

(b) Analytic farfield distribution of TM_{01} and TE_{21} associated hybrid modes

Figure 6.8: Comparison of modes supported in the back to back system and the analytic distributions of the theoretically supported modes, in order to identify the modes calculated using SVD

these matrices revealed that over the six azimuthal orders the horn supported three modes with non-zero singular values, thus the horn essentially supports three independent (hybrid) modes. The modes obtained using the SVD analysis are arranged such that there is one mode located in each of the azimuthal orders $n = 0, 1, 2$. There are no modes supported for $n = 3, 4, 5$. This is precisely as is expected from the prediction made using the cut-off frequencies of the waveguide modes in the throat section of this horn. This means that each of the modes shown in figure 6.8a belong to a different azimuthal order. They are numbered consecutively, starting with the $n = 0$ azimuthal order, so mode 1 is associated with the $n = 0$ azimuthal order and thus represents the hybrid mode corresponding to the TM_{01} mode. Mode 2 is associated with the $n = 1$ azimuthal order and so represents the TE_{11} (HE_{11}) mode, i.e. this is the mode that was used in the calculations in section 6.2.1. This is correct, as the frequency and geometry have remained the same, with the only change being that now additional azimuthal orders are being included in the calculation. This does not impact the performance of the $n = 1$ azimuthal order, as the azimuthal orders are independent of each other (incoherent), so this mode is precisely the same as that in the previous section. Mode 3 is associated with the $n = 2$ azimuthal order and therefore the hybrid mode that represents the TE_{21} mode. It is understood that the waveguide and hybrid waveguide modes will now be referred to equivalently for the sake of clarity as the system to be considered in chapter 7 uses this analysis to consider smooth-walled horns, i.e. standard waveguide modes are used. Therefore when for example the TM_{01} mode is mentioned, the hybrid mode excited by this in the 0^{th} azimuthal order is in question. The analytic forms of the $n = 0$ and $n = 1$ modes shown in figure 6.8b compare well with the corresponding modes in panel (a), confirming that the modes calculated using SVD do actually

correspond to the stated waveguide modes.

Figure 6.9 shows the farfield response of the back to back horn system as a function of off-axis illumination from zero to sixty degrees. The overall trend is precisely what would be expected. As the angle of illumination increases, the overall power contained in the field decreases. The only exception to this is the relative power levels of the 40 and 50 degree cuts. These fields show that the illumination of the horn at fifty degrees gives rise to a farfield pattern

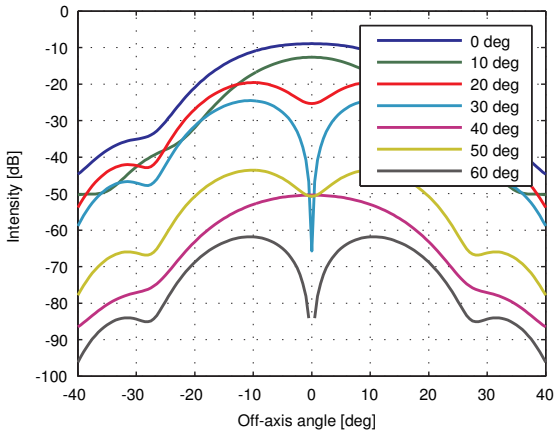

Figure 6.9: Response of the farfield of the few-moded back to back corrugated horn system for various off-axis angles of illumination

containing more power than is contained in the equivalent pattern at an angle of illumination of 40 degrees. The reason for this likely lies in the precise amount by which each of the three modes is excited at this frequency, a point which will be considered momentarily. Considering the overall field distributions as a function of excitation angle, the shape can clearly be seen to change with excitation angle. This was not observed in the single-moded example, as in this case there was only one mode propagating in the system. This meant that varying the excitation angle merely changed the amount by which the mode was excited, but this had no impact on the spatial distribution of the mode. Thus, the same shape mode was excited in the system as a function of angle, with merely the amplitude varying. In this example there are three modes in the system. The hybrid modes representing the TM_{01} and TE_{21} modes have approximately the same distribution and so will likely couple to the plane wave by the same amount, but this will potentially be very different to the amount by which the mode representing the TE_{11} mode will couple to the plane wave. As the excitation angle changes, these modes will therefore couple by different amounts to the plane wave and so will be excited in different proportions as a function of excitation angle. Since there are two distinct shapes amongst the three supported modes, as the relative amount of each mode present changes then the composite field in which the modes have been added incoherently in order to obtain the overall farfield, should display a varying shape in addition to a variation in amplitude. This is what is observed, and so the results obtained are sensible.

Figure 6.10 shows the coupling of each of the three modes to the plane wave source as a function of off-axis illumination angle. For on-axis illumination, corresponding to an excitation angle of zero degrees, all three modes couple to the plane wave by a factor of approximately 0.82. This high coupling of all modes means that they will all be excited by a similar amount and so the composite farfield will consist of an approximately equal weighting of all three modes, again due to the symmetry of the system about the filter section. The addition of the modes of azimuthal order $n = 0, 2$ has a visible effect on the farfield pattern in this case, with the sidelobe structure varying between the few and single-moded cases. It can be clearly seen in the few-moded case that the addition of the extra modes results in the suppression of some sidelobe and shouldering structure.

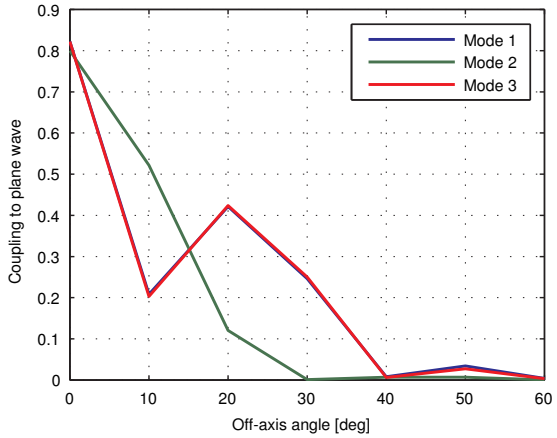

Figure 6.10: Coupling of the plane wave to the few-moded back to back corrugated horn system for various off-axis angles of illumination

As the angle of illumination increases to ten degrees, the coupling of all modes decreases sharply, with the coupling of the TM_{01} and TE_{21} associated hybrid modes matching each other closely due to the similar spatial distributions. The coupling of the TE_{11} mode is higher at this angle, the implications of which are clearly visible in the farfield pattern in figure 6.9 as the on-axis power is still high with the composite field more closely resembling the TE_{11} mode than the other modes. Increasing the excitation angle to 20 degrees reveals that the TM_{01} and TE_{21} modes show increased coupling relative to the TE_{11} mode. This is because the former modes have a greater off-axis presence at this angle than the TE_{11} mode, and so they will have higher coupling. Examining the corresponding composite field shows this contribution to emerge clearly with an on-axis dip emerging. The field very obviously contains contributions from both of the mode distributions.

Increasing the angle of illumination to thirty degrees shows that all modes suffer a drop in coupling to the planes wave, however the TM_{01} and TE_{21} modes are still dominant, and since the contribution of the TE_{11} mode at this angle is effectively zero, the composite farfield contains very little on-axis power and closely resembles the distribution of the TM_{01} and TE_{21}

modes, as would be expected given the relative couplings. Further increasing the angle of illumination to 40 degrees reveals that there is very little coupling between the plane wave and any of the modes. This is clear from the very low, difficult to measure, power levels displayed for the composite field. The farfield shows on-axis power and has a distribution that resembles the TE_{11} mode and so it can be concluded from this that although the coupling of all modes to the plane wave is very low, in relative terms the TE_{11} mode is dominant and has the highest coupling.

For an angle of illumination of 50 degrees, it is seen that the coupling of the TE_{11} mode has remained approximately zero as before, however due to the sidelobe structure of the TM_{01} and TE_{21} modes, their coupling has moderately increased and so these modes are dominant at this angle, a fact which is clear to see from the composite farfield distribution. The indicated coupling also explains the elevated power levels of the 50 degree illumination case relative to the 40 degree case as it is clear from figure 6.10 that at this angle the total coupling is greater due to the sidelobe structure of the modes of azimuthal order $n = 0, 2$. For an angle of illumination of sixty degrees, the coupling of all modes has again dropped off to approximately zero, with the power levels of the composite farfield falling in response to this to levels that are comparable with the 40 degree case. The farfield reveals however, that unlike the 40 degree case the TM_{01} and TE_{21} modes are dominant and so must couple marginally better to the plane wave at this angle than the TE_{11} mode.

6.3 Conclusions

In this chapter, one possible method to simulate the off-axis illumination of a waveguide system by a plane wave in SCATTER was considered. Although the approach applies to any such system, the specific cases considered were identical horns mounted back to back, meaning that the entrance and exit apertures of the system were identical. The method works by calculating the modes supported by the structure by means of a standard SCATTER calculation and singular value decomposition. The distribution of these modes at the entrance aperture of the horn, the aperture being illuminated, is then calculated and an overlap integral is carried out between each of the modes and a plane wave incident on the aperture at some angle. The supported modes are then scaled by their coupling factors and used to calculate a suitable input scattering matrix which SCATTER then uses to excite the waveguide structure and then proceeds as normal to calculate the performance of the system subject to this modified excitation.

The examples considered were a back to back version of the corrugated horn (at 100 GHz) used in NUIM. The results obtained were sensible. For the single-moded case, only one mode was present (TE_{11}) and so the overall trend in the farfield of the system as the angle of illumination was increased was a reduction in intensity with the spatial distribution remaining

the same, as expected. Any deviations from this trend can be accounted for by considering the sidelobe structure of the mode and realising that this will cause an increase in coupling at particular angles of illumination relative to the overall trend. The few-moded case was also considered, and here there were two additional modes supported (TM_{01} and TE_{21}). The overall field levels reduced with increasing illumination angle, as well as the overall field shape changing as the relative contributions of each mode changed, since there were two very different distributions contributing to the composite farfield. Again, any deviations from the trend of decreasing intensity could be accounted for by considering sidelobe coupling. This is all expected, verifying the functionality of the code.

In the next chapter, an experimental representation of a system such as that described here will be considered, in collaboration with the University of Manchester

It is intended to extend the code to allow illumination of the horn with an arbitrary field, both on and off-axis. It is also intended to make the code more efficient during the consideration of multiple angles (in the back to back case) by making use of the symmetry between the input and output fields outlined in the previous section. This will have to be investigated further. The possibility of including this off-axis functionality in the pyramidal SCATTER software will also be investigated.

Chapter 7

Measurement of Individual Modes in a Multi-Moded Horn

In this chapter, the initial development of a system with the ability to measure the individual modes present in a multi-moded horn is discussed.

The principle in question is, as described in the previous chapter, that if a horn is illuminated with a plane wave from a certain angle then each mode won't be excited with unity power (as per the standard blackbody illumination used in SCATTER), rather it will be excited by an amount proportional to the coupling between each mode and a plane wave of the indicated angle. If a back to back version of the horn is constructed, then the source that is exciting the horn excites the exit aperture of the original horn. The modes present in the horn normally are all supported at the aperture and the excitation source now couples to this plane instead of the throat waveguide feeding the normal horn geometry. In the best case (blackbody) illumination scenario, the source will excite all of these modes. They will then propagate through the horn, passing through the various filtering sections (both of the throat sections) in the back to back horn arrangement. Following the throat, the modes will propagate through the second portion of the back to back horn, having only been excited by an amount proportional to the coupling of the aperture field to the excitation source. The modes will arrive at the aperture in a distribution that matches the distribution at the other aperture. Assuming illumination from only one angle, as the excitation angle is changed, the source will couple to the various modes of the horn by differing amounts, depending on the off-axis distribution of the fields associated with these modes. For a given angle, any mode that couples to such a plane wave will thus be excited and enter the horn. It is possible that for some angles of illumination, some modes will not be sufficiently excited to contribute, and so for these angles of illumination those modes are effectively not excited.

Using this principle, the individual modes supported by the horn could potentially be measured individually. The underlying principle is the same as that which the SCATTER modification illustrated in the previous chapter is based on, namely that due to the off-axis

field distribution of each mode, an off-axis source won't couple to certain modes for given angles. Each mode therefore has certain angles at which the mode is not sufficiently excited to make a meaningful contribution, and at other angles it will contribute. These angles depend on the sidelobe and main beam structure of each of the modes. Therefore, for a given angle, potentially only a subset of the various modes supported will be excited. By knowing the amount of each mode that should be present for a given excitation angle, the contributions from all but one mode could be subtracted from the measured pattern, potentially resulting in the ability to measure the remaining mode in isolation. This would be extremely useful in the design of horns and optical systems, as well as in the verification and testing of already manufactured horns.

The equipment, experimental setup, control software and initial tests are now discussed.

7.1 Equipment and Experimental Setup

The experimental setup is shown in figure 7.1. This experimental arrangement was developed to compliment the phase and amplitude Vector Network Analyser (VNA) scanning system in the Department of Experimental Physics, NUI Maynooth. The source is a 100 GHz Gunn diode which draws a current not exceeding 40mA from a power supply(BIAS FDB-F8). The source feeds a pyramidal smooth walled 100 GHz feed horn, of internal dimensions 2.7 mm$\times$ 2 mm. The detector is an Artas Golay cell detector with a tapered waveguide section attached to the front in order to couple better to the field being measured. This arrangement is mounted on an x,y,z stage manufactured by Zaber (In order to protect the Gunn diode, rather than rotating the source horn to achieve off-axis illumination, the back to back horn is rotated about the axis indicated in figure 7.1. This shifts the location in the scanning plane in which the beam pattern from the back to back horn will fall. In order to compensate for this, the scanner is re-centred on the new location of the beam). The Golay cell functions best when the signal being measured is modulated at a frequency of 20 Hz, and so to achieve this the modulation function on the source power supply was used to modulate this supply with a 20 Hz square wave. A Tenma Jupiter 500 function generator supplied this square wave. This modulation resulted in the Gunn diode, and hence the output of the feed horn, being modulated at the frequency required by the Golay cell. The output signal from the detector was passed to a model 5104 Princeton Applied Research lock-in amplifier, which was also supplied with the function generator signal as a reference. The output of the lock-in amplifier was passed to the computer using the ADC functionality of an NI USB-6008 (input range set to 0-10 V) data acquisition card. On the computer, software was written using NI LabWindows which controlled the scanner position and recorded the output voltage from the system, that of lock-in amplifier, as a function of this position. The actual setup is shown in figure 7.2, offering a dynamic range of circa -30 dB.

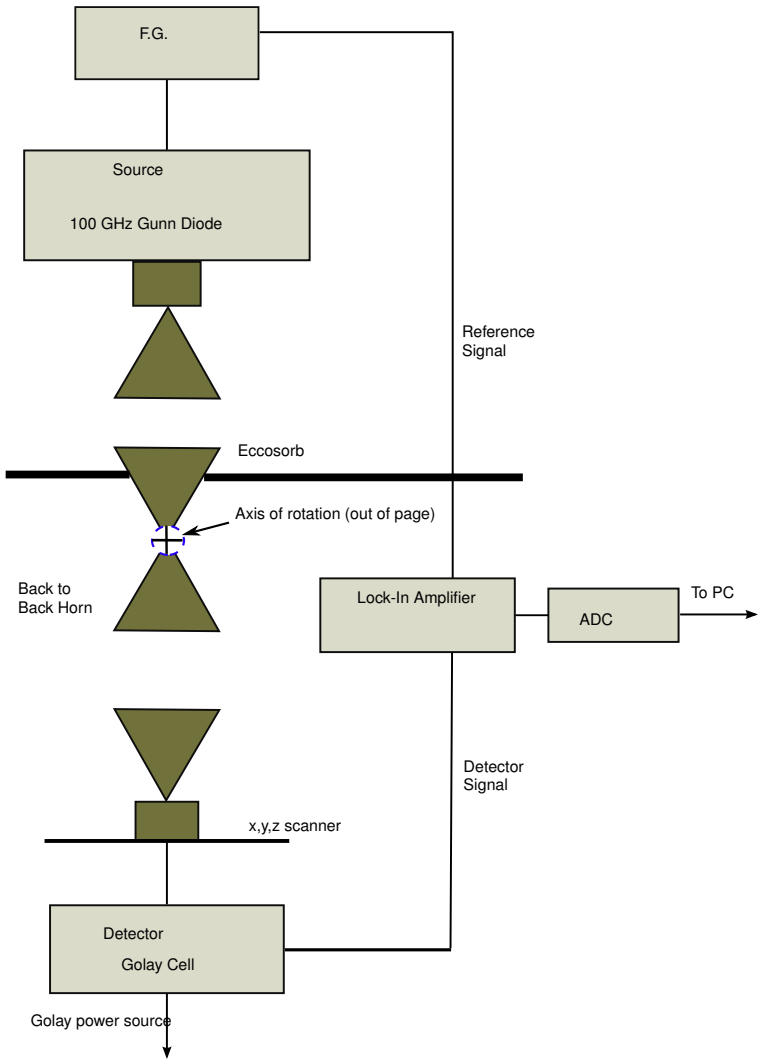

Figure 7.1: Schematic of the setup for the measurement of the individual modes present in a horn

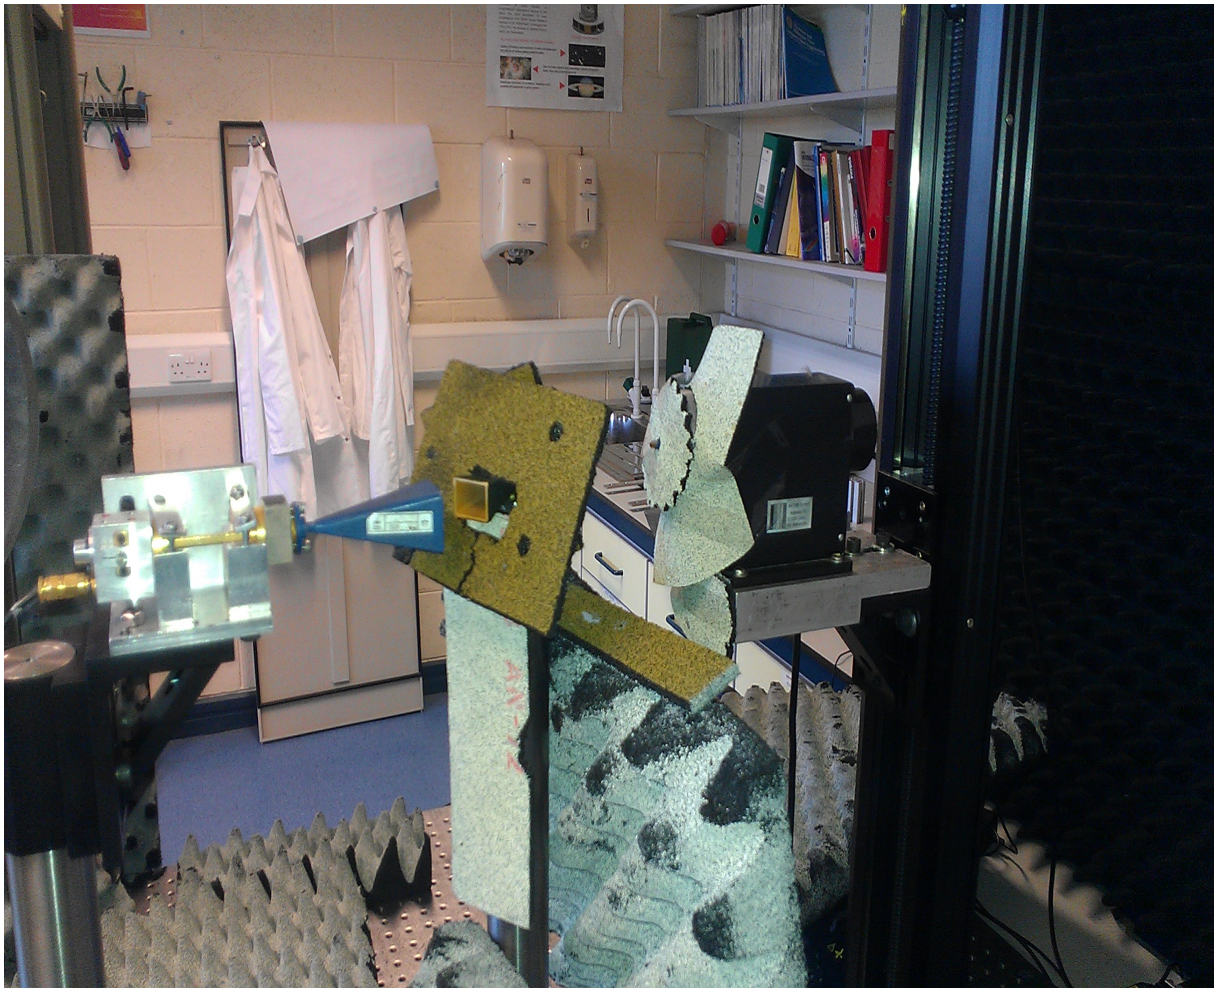

Figure 7.2: Experimental setup. In this image a pyramidal horn is used as the source, however in the measurements presented in this thesis the source was a conical corrugated horn.

The Golay cell functions as follows. It is made up of two primary components: a gas filled enclosure that is primarily encased by a material that is sensitive to a wide wavelength range that the cell is designed to detect, and a deformable membrane that is attached to the back of the enclosure. When radiation of the correct wavelength is incident upon the detector, it becomes absorbed by the material surrounding the gas enclosure, resulting in the gas heating up. This heating results in the expansion of the gas, which causes the membrane at the back of the enclosure to deform, as the gas presses on it. The power supply to the Golay Cell either produces a potential difference across the membrane that alters in response to changes in the shape of the membrane (the more deformation the higher the potential), or is used to illuminate the membrane via a light. The light reflected from the membrane is angled onto a photodiode which produces an output. As the membrane is deformed, more or less light is incident on the photodiode, and so the output is varied. Either way, the amount of radiation incident on the detector determines the degree to which the membrane is deformed, and this in turn determines the output of the detector, which is proportional to the intensity of the incident radiation. The greater the deformation, the higher the temperature of the gas, the more intense the radiation is at the wavelengths that the detector is designed to operate at.

In this setup, Eccosorb is used to minimise reflections in the system and to isolate the Golay cell from any additional sources of radiation external to the experiment. The cell is sensitive to infrared radiation, and since this is present naturally in large amounts, great care must be taken to insulate the system. Also to be considered is the fact that the Golay cell is by its nature highly dependent on heat variation, and so its positioning must be carefully planned so as it is not in close proximity to heaters or components that may produce heat, causing a substantial and unpredictable/uncontrollable temperature shift.

With the system in place, it was necessary to determine the optimal settings for the lock-in amplifier. By simply placing the detector in front of the source horn and observing the output voltage from the lock-in amplifier, the following settings were determined, giving maximum dynamic range and stability.

- **Range/Phase:** Adjusted depending on power arriving at detector to keep the amplifier below the 20 V output saturation level.
- **Time Constant:**0.1 s (determines how long the value read out by the instrument is integrated and averaged over)
- **Roll off:**6 dB (determines how steep the transmission function of the filter used by the instrument is)
- **Filter:**None

These settings were used in each of the subsequent scans, as they gave the best compromise between the stability of the output and the subsequent necessary dwell time during the scans,

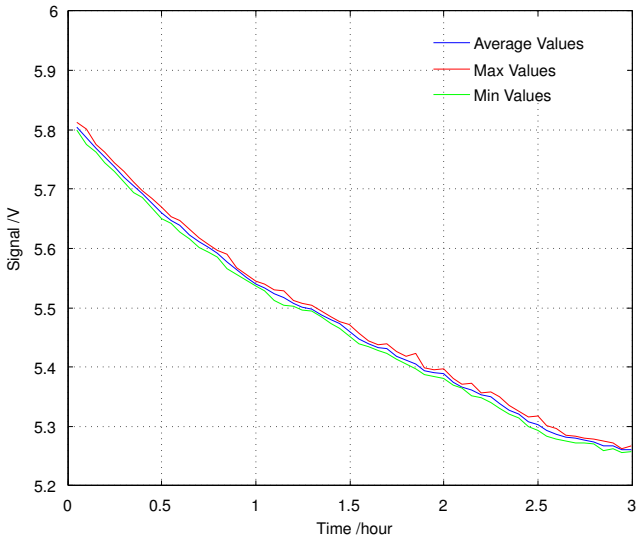

Figure 7.3: Stability test of the system output. The max/min/average values refer to the system output in 2.92 minute periods

considering the typical response time of Golay cells of approximately 10 ms. Checks were carried out by covering the Golay cell with Eccosorb and removing it again, and checking the response time of the lock-in amplifier. The time taken for a stable output to be achieved was well within reasonable limits for experimental work. The stability of the output signal was then tested.

Over several three hour periods, the stability of the system output was tested. The range and phase of the lock-in amplifier were adjusted so as to give a maximum voltage output whilst the detector was directly in front of the source. the central position of the scan. The scanner was then left in this position, with the output of the system recorded every second. The data was then divided into 2.92 minute intervals, and the average, maximum and minimum values of each period were plotted as shown in figure 7.3. In theory each plot should be approximately constant (a horizontal line), however an exponentially decaying drift equating to a total of 0.6 V is noted over each three hour period. Regardless of the starting voltage values or the time that the scan was started at, a decay of similar magnitude was always observed. Examining the time derivatives of these plots showed that the decay was physically very slow, and so for a given measurement (dwell time of a few seconds at most), the integrated system output could be time averaged to give an accurate value. This also meant that were a high resolution scan being made, which would be a lengthy process, the on-axis value would need to be periodically re-measured and all measurements taken in the corresponding period would need to be re-normalised to this so as to allow valid relative comparisons to be made between all measurements in a given scan grid.

Following this, the dynamic range of the detector was tested. To achieve this, a cut was taken along the vertical and horizontal axes of the source horn, using 300 steps in either di-

rection. The results are shown in figure 7.4. The detector was found to have, on average, a

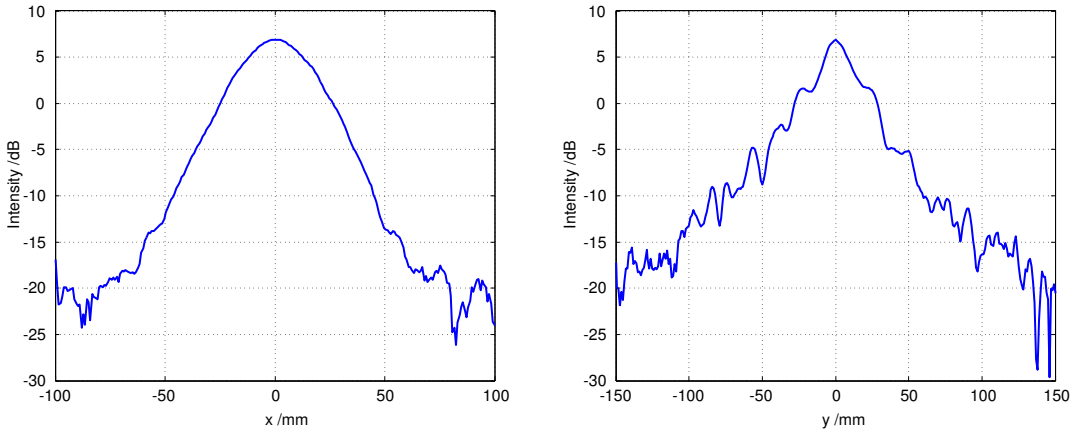

(a) x-cut. y=0 (b) y-cut. x=0

Figure 7.4: Cuts in the x and y principal directions to determine the dynamic range of the detector

dynamic range of approximately 30 dB, which is sufficient for the measurements to be undertaken. The shape of the measured beam is also to be noted. The horn used was a pyramidal horn, and so it is expected, [60], that there should be sidelobes present in one orthogonal direction only. The x-cut is a very smooth Gaussian distribution, with the only structure that appears being some noise towards the bottom end of the dynamic range of the detector. The y-cut has clear sidelobe structure however, and this is consistent with what would be expected. The detector was therefore functioning correctly. This was further confirmed by averaging the output at the various critical points (0 dB, -3 dB. -5 dB etc.) on the curve for fifteen minutes each, and plotting the results. These positions were found by noting the on-axis voltage and calculating the voltage which would therefore be expected at each of the critical points (for example at the -3 dB point the voltage would be reduced to half of the on-axis value), with the probe then moved to the point at which this voltage was recorded. Table 7.1 shows the location on the curve of the measurement (0 dB, -3 dB etc.) and the corresponding average voltage value. The deviation of the maximum and minimum values from each average was of the order of 0.001 dB, and so was not considered significant. The results were taken for one half of the bell curve only. Figure 7.5 shows the mean voltage for such a set of readings plotted as a function of position on the curve. This should therefore resemble half of a clean bell curve, which is precisely what is observed. All voltages quoted are DC components, as this is what is to be physically measured and recorded, the steady output of the lock-in amplifier. The AC component of each set of measurements was also noted, however this was insignificantly small, as required.

The above tests confirmed the correct operation of all components of the system, and the system as a whole. The dynamic range of the detector was found to be approximately 30 dB (see figure 7.4), and the lock-in settings optimising its performance were also determined. Although

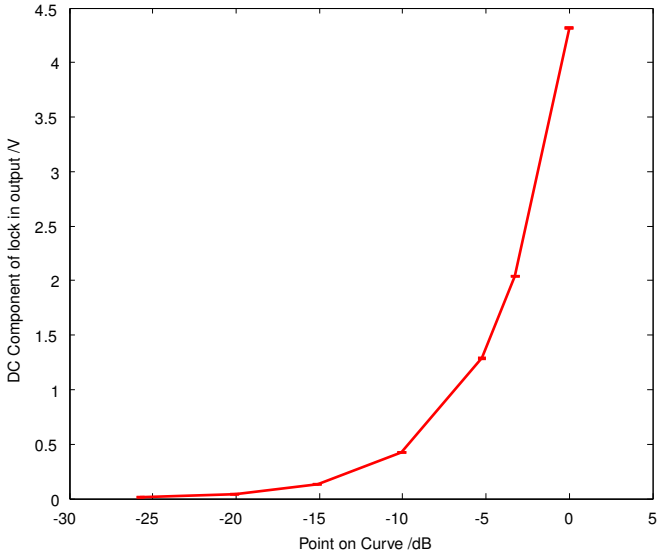

Figure 7.5: Mean Voltage as a function of location on curve

the dynamic range achievable with this arrangement is more limited than the VNA setup, it allows a reasonable range with which to explore the response of many devices or antennas.

7.2 Control Software

It was required that a customised piece of software was developed to control the scanning and measurement system. The ultimate measurements are highly precise and so the scanning parameters must be highly customisable to allow the best possible performance to be achieved. The software consists of two primary parts: the section to control the measurements (data acquisition and processing) and the section to control the scanning system. The former section was developed by Dr. Massimo Candotti, and the latter by the author, with the two sections being integrated at a later stage to form one overall measurement control system. These two sections will now be described. The control panel, which will be referred to in the description, is shown in figure 7.6. The control software was developed in NI LabWindows, a C based

Location on Curve	**Mean Voltage** [V]
0 dB	4.3194
-3 dB	2.040
-5 dB	1.289
-10 dB	0.426
-15 dB	0.133
-20 dB	0.042
-25 dB	0.012

Table 7.1: Mean DC voltage values recorded at various locations on the bell curve

integrated development environment produced by National Instruments.

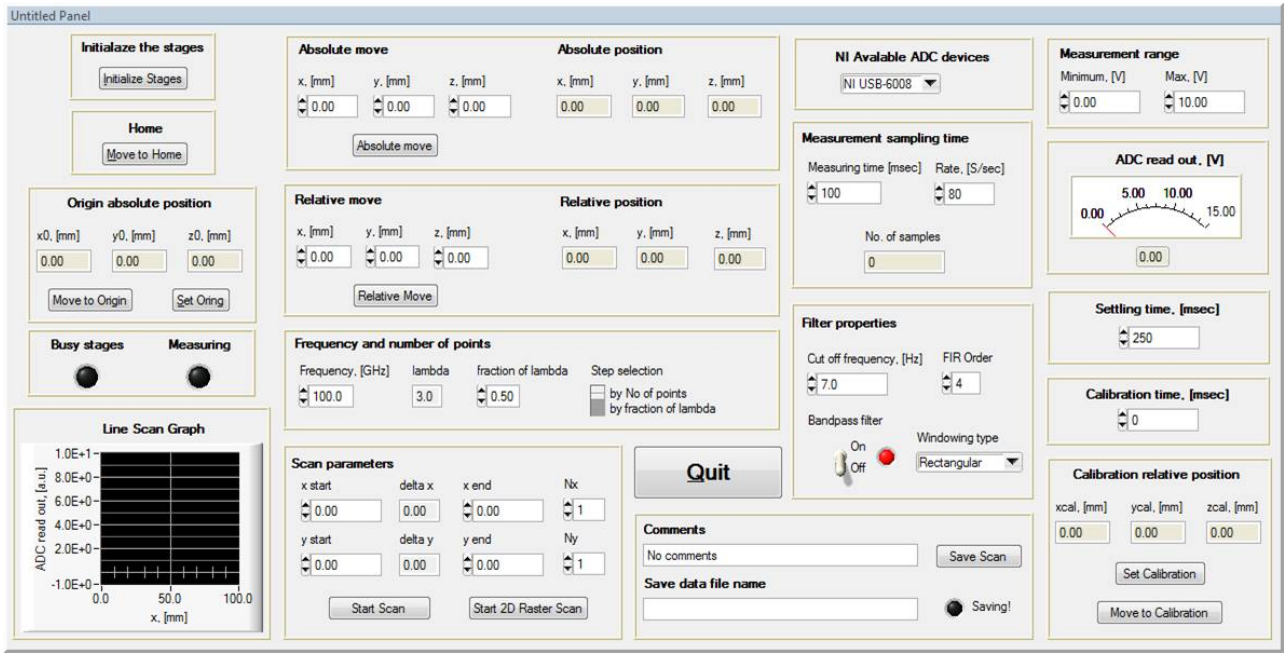

Figure 7.6: Control panel for scanner software

7.2.1 Data Acquisition and Processing

In this section, the functionality of the portions of the control panel that control the data acquisition and processing will be discussed. There are two data acquisition cards attached to the computer (each providing different resolution), and the one being used must be initially selected. The NI USB-6008 is used in all measurements here, and this is the default option in the **available ADC devices** dialogue box. The minimum and maximum input voltages are then selected in the **measurement range** section. The respective default values are 0 V and 10 V, as a negative output from the system is not normal with the ideal lock-in amplifier configuration. 10 V is the maximum allowable input to the 6008 card in this measurement configuration, and so this determines the upper limit.

In the **measurement sampling time** section, the measuring time (amount of time that the output of the lock-in amplifier is read in at a given point in the scanning plane) can be selected in ms, along with the rate in samples per second. Using these values, the resulting number of samples taken at a given point in the scanning plane can be calculated and is indicated to the user. The acquisition logic treats each individual sampled value as an isolated measurement. The ideal values for these parameters, determined by Dr. Candotti, are a 400 ms sampling rate, sampling 150 times per second, giving a total of 60 samples. The averaged output of each sampling period is given opposite in the **ADC read out** box in volts. This allows real time monitoring of the measurements during a scan.

The **filter properties** box allows the filter characteristics to be adjusted. The filter cut-off frequency can be set in Hz (optimum value of 7 Hz), the FIR order (Finite Impulse Response-a filter with the property that its response to an input pulse settles to zero in a finite time as opposed to instantaneously) can be set (optimum value of 35) and the type of windowing used can be set, with the ideal regime being the Hamming filter. There is also a switch for selecting the filter on or off. The filter is used to smooth the measured result and return an accurate value for the average, and so should ideally be left on as this is critical in terms of achieving an accurate overall scan result. Central to this also is the settling time, in ms, which is an amount of time that the scanner is allowed to settle for after moving positions before sampling begins. This serves to reduce the impact of vibrations in the scanning structure on the system.

This portion of the control panel also shows a real time readout of the current scan cut, allowing the measured voltage at each point to be displayed as a function of position, which enables real time monitoring of the scan. If there is an issue then it will be clear here, and the scan can be stopped and adjusted accordingly. This means that time is not wasted waiting for a scan to end before being able to see if the result is looking correct.

The file can also be saved. Once a file name (and extension) are inserted in the **save data file name** window, the file is automatically saved, however a manual save can be initiated through the save scan button. A comment can also be added to the file. The file contains information on the scan such as the scan type (1-D or 2-D raster), the start/stop times, the scan frequency and the scan range/number of points. This allows a complete record of the scan to be kept, and allows a Matlab file to plot the data on the appropriate graph type for the scan (1-D or 2-D).

The parts of the panel which control the scanner position are now discussed.

7.2.2 Scanner Control

The remainder of the panel is concerned with the control of the physical position of the scanner, the underlying code for which the author developed.

The stages must first be initialised, or reset, as required by the internal software of the stages. The stages contain absolute encoders, and so their absolute position is always known as opposed to their relative displacement from some variable position. The initialisation essentially initialises these encoders to their current position. A *C* function is supplied by Zaber to affect this, and so pressing this button calls this function resulting in the initialisation of the stages. The home position (all axes to their 0 position, corresponding to the centre of the scanning plane) can be realised quickly by pressing the move to home button. The stages contain predefined memory blocks for preset and user defined positions. Home is one such preset position. Zaber supply a function to move the stages to one of these memory block positions, and so pushing this button supplies the memory block location of the home position to this

function, and the axes move to the home position.

The **origin absolute position** box allows the origin of the scan to be set and saved in absolute coordinates, meaning that the centre of the scanning pattern can be set to a user defined location which does not necessarily coincide with the centre of the scanning plane. The absolute coordinates are entered in this box and these values are passed to the stages, one axis at a time, so that the axes move to the defined position. Pressing the set as origin button saves the current absolute position of the stages as the origin in a user defined memory block within the stages. At any time, the move to origin button can be pushed and the code will call the function which moves the stages to a pre-defined memory block position, while passing the origin memory block location to this function, resulting in the stages moving to the origin.

The **absolute move** and **relative move** boxes allow the user to move the axes to an absolute position (in terms of mm) or by some amount relative to the current position. This allows the user, for example, to set the origin of the scan or to move the scanner to a certain location for a manual measurement. Each box provides a readout of the current absolute or relative position.

The **frequency and number of points** dialogue box allows the stepsize of the scan to be determined by the user manually, either in terms of a fixed scanning distance divided into a user defined number of points or in terms of wavelength. With the step selection switch in No. of points, this box is ignored, and the user manually enters the number of points, as will be discussed below. With the switch set to fraction of lambda, the stepsize is determined by the frequency at which the scan is to be carried out. The user enters the frequency in GHz, and the wavelength is displayed in mm. The fraction of lambda is then entered by the user, with the default being 0.5. Thus, the code will step between the maximum and minimum x and y values, defined in the **scan parameters** box, in steps of 0.5 lambda (typically 1.5 mm). Alternatively, the scan parameters box allows the number of points desired to be entered. If fraction of lambda is selected, this box is unusable. In either case, the size of the steps, *delta* x and *delta* y is displayed to the user. If a 1-D scan is desired, the appropriate bounds are inserted and the start scan button is used. If a 2D scan is desired, the start 2D raster scan button is used. The remaining two controls follow from this.

The **calibration** options are required at this stage. The calibration position is generally taken to be the origin of the scan, however it is possible to re-define the calibration position relative to the origin. The relative coordinates of the new calibration position can be set, and the calibration position is saved in the allocated memory block address. It is possible to move to the calibration position using the *move to cal* button, to check that it is correctly defined. The calibration time, in ms, is the time that the measurement is integrated and averaged over in the calibration position. 250 ms is a standard example. All measurements in a scan are normalised with respect to the calibration position value, so an accurate value is essential.

In the case of a 1-D scan, the calibration value is taken once at the start, and all values are normalised with respect to this. In the case of a 2-D scan, a series of cuts, a new calibration value is taken at the start of each cut, and the cut is normalised to this value. This allows the effects of the drift illustrated earlier to be cancelled out, as all measurements are normalised with respect to the on-axis value that was present at the time that the measurement was taken.

The remaining option is the **settling time** option, in ms. A typical value would be 250 ms. This option defines how long the system waits, having moved to a measurement position, before beginning a measurement. This allows the effects of vibration in the stages etc. to be removed, as with a suitably high time then the vibrations will have ceased. In order to keep measurement times reasonable, a compromise must be reached between allowing all vibrations to cease and a short settling time. A low settling time returns usable results as the stages are securely mounted on the optical bench and the damping is high, so vibrations will be quickly eliminated. Therefore, the quoted time of 250 ms is sufficient for a highly accurate result, whilst not increasing the scan time to unacceptable levels.

7.3 Initial Testing of the System

A small selection of initial results will be presented that were obtained using this testbed and software. The first set of results that will be presented relate to the corrugated horn used in NUIM. The field pattern expected of this horn is well understood, [37], and so this scan was used simply to test the operation of the scanning system. The horn was mounted on the source (Gunn oscillator) such that the detector was in its farfield, and numerous measurements were taken for various system configurations in order to obtain the optimum results. Figure 7.7 shows the measured co-polar field associated with the corrugated horn, performed without using Eccosorb to minimise reflections and outside interference. The measured field is precisely what is expected to be measured from this horn, [37, 61]. The Gaussian distribution is well defined in all directions, with low level sidelobes which are accurately represented as they fall within the limits of the previously measured dynamic range of the system. The field is clearly quite symmetric down to this level, as expected, and so the detection system and measurement software can both be seen to be functioning correctly. This is confirmed by figure 7.8, which shows the horizontal and vertical cuts through the 2D field. Of interest is panel (b) in figure 7.7, where the calibration value is plotted as a function of each calibration event. Since this was a 2D raster scan, the calibration value was re-measured at the start of each scan cut. This calibration value is seen to drift by approximately 0.1 V over the course of the scan. This would be an issue in the final result if no action were taken, however each cut is normalised to its respective calibration value, and so the cuts can be directly compared, resulting in an accurate reconstruction of the field despite the time period of the scan being sufficiently long so as to allow drift in the output value of the system at a fixed measurement point to occur. The

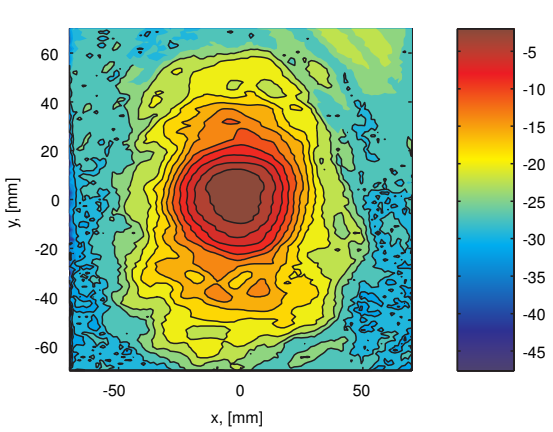

(a) Co-polar field of the NUIM corrugated horn

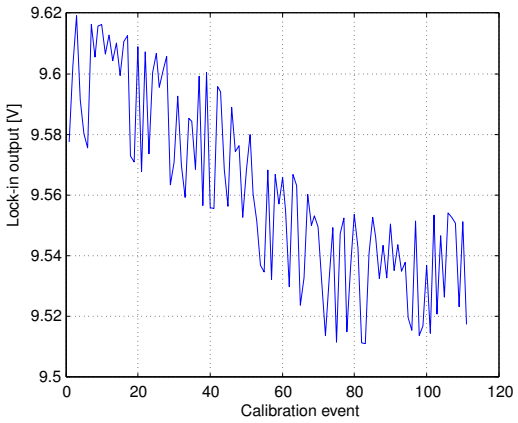

(b) Calibration value as a function of calibration event

Figure 7.7: Co-polar field of the corrugated horn in NUIM, and the calibration value as a function of calibration event

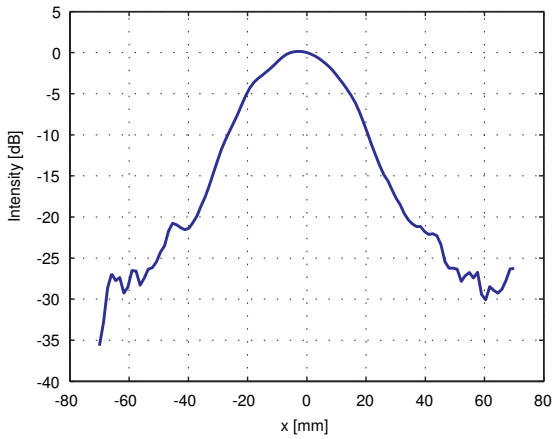

(a) X-cut

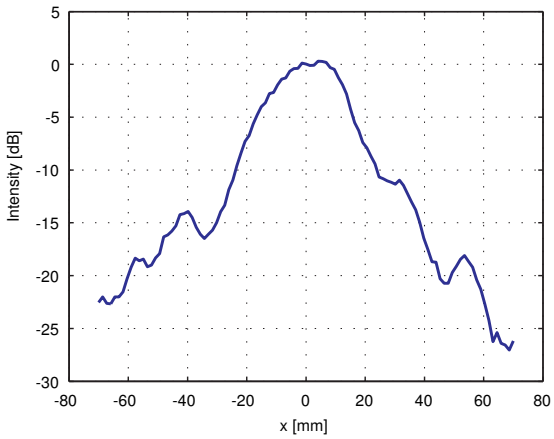

(b) Y-cut

Figure 7.8: Cuts through the measured 2D grid of the corrugated horn

calibration occurs within the data acquisition software, so the uncalibrated form of the data presented for the corrugated horn is not available, however the effect of the calibration can still be considered. Figure 7.7b shows the calibration value as a function of calibration event. This value corresponds to the output of the lock-in amplifier at the origin of the scan, i.e. pointing directly at the horn. This value is seen to fluctuate about a general downward trend, dropping by approximately 0.1 V overall. The scan is a raster scan, and thus starts at one side of the scanning grid and works its way across in vertical cuts, with the output value of the lock-in amplifier at each point being normalised relative to the value at the origin of the scan, as this is taken to be the calibration value of the system. If the same value were to be retained throughout the entire scan, then although the values for the first few cuts at the starting side of the scan would be normalised correctly, the true value of the output would be dropping slightly as the scan proceeds across, resulting in the output value being underestimated to a

progressively higher degree as the scan proceeds. This would result in asymmetry appearing in the beam in the form of a slant, with the starting side of the scan showing higher levels than the end side. The constant re-calibration of the system eliminates this error.

Figure 7.9 shows the results for the same horn, but with the system insulated with Eccosorb to minimise contamination from outside sources, and also to minimise reflections within the system. It is clear that the field is much more well defined and symmetric. This is emphasised

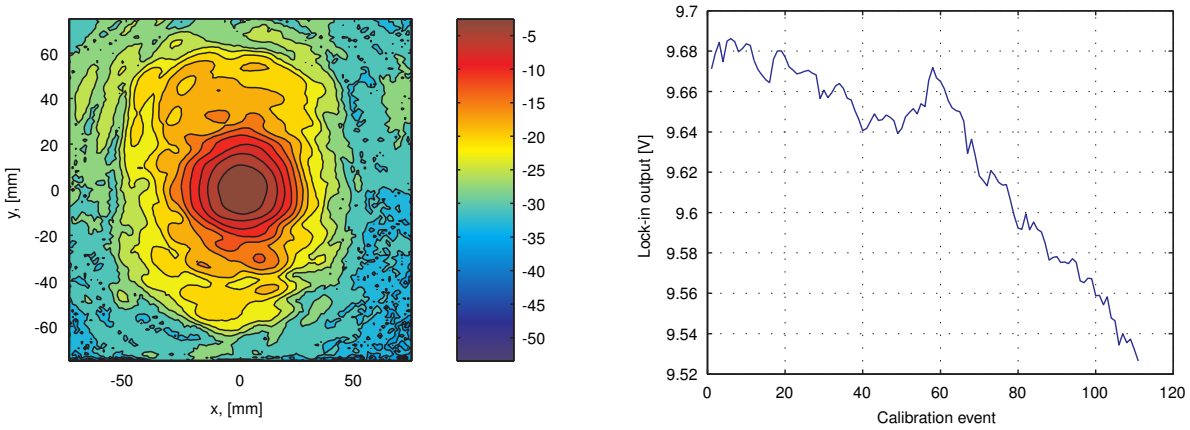

(a) Co-polar field of the NUIM corrugated horn

(b) Calibration value as a function of calibration event

Figure 7.9: Co-polar field of the corrugated horn in NUIM, and the calibration value as a function of calibration event, with Eccosorb

by examining figure 7.10, which shows the vertical and horizontal cuts through the 2D grid, as before. It is immediately clear upon examining the 2D field and the cuts that they are much

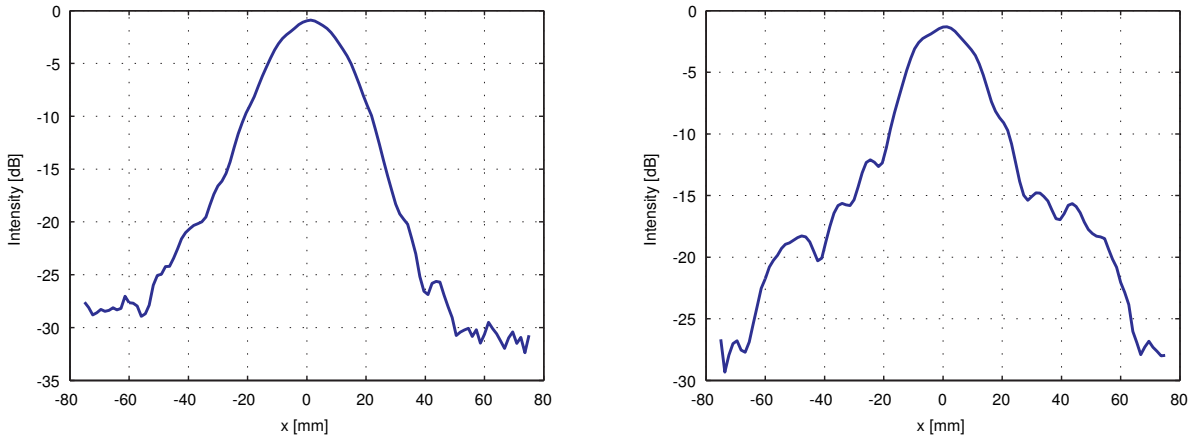

(a) X-cut

(b) Y-cut

Figure 7.10: Cuts through the measured 2D grid of the corrugated horn with Eccosorb

more well defined when the Eccosorb is used, as there is less jitter on the measurement which may have previously arisen from reflections within the system (from the metallic optical bench

and surroundings) and contamination from external infrared sources. This results in a higher degree of beam symmetry and the field being clearly defined down to lower levels than before. The sidelobes are therefore also much better defined. The Eccosorb clearly has an effect at lower intensity levels that results in a more clearly defined field distribution being measured at these levels (thus increasing the dynamic range of the setup), but also has an affect at higher levels where the measurement appears more stable. A similar amount of drift was noted in the calibration values compared to the case without the Eccosorb, however there is far less oscillation once the absorber is used, implying greater system stability.

These measurements yielded results in line with what is expected for pyramidal horns and conical corrugated horns, showing that the detector and data acquisition/processing software function correctly. These measurements, particularly those performed on the corrugated horn, allowed the performance of the system to be investigated with respect to, for example, dynamic range and the benefit of using Eccosorb versus not using it.

7.4 Back-to-Back Cylindrical Horn

In this section, results will be presented for a cylindrical smooth-walled horn that was manufactured at the University of Manchester. Two units were manufactured, enabling them to be mounted together at the throat sections giving rise to a symmetric back to back system of the type considered in chapter 6. The results presented were obtained at NUIM using the 100 GHz Golay cell measurement testbed outlined above. In each case, only the co-polar field was measured, as the radiation is coupled to the Golay cell via a section of rectangular WR10 waveguide which supports only one polarisation.

The horn in question was designed by the author in collaboration with Dr. Bruno Maffei of the Jodrell Bank Centre for Astrophysics at the University of Manchester, with the horn being manufactured in their workshop. The horn has an aperture radius of 10 mm, a throat radius of 1.4 mm and a total length of 43.44 mm, including a throat of length 0.24 mm. The horns were made from copper. Figure 7.11 shows profiles of the horn in both single and back to back configurations. The modes supported by the horn as a function of frequency in the 80-115 GHz range are given in table 7.2. The frequencies shown are effectively the frequencies at which the 'new' mode indicated became supported. Table 7.2 shows that at 100 GHz the horn is few-moded, supporting both the TE_{11} and TM_{01} modes. Illuminating all possible modes at 100 GHz and carrying out a singular value decomposition on the resulting scattering matrices reveals that there is one mode supported in each of the $n = 0, 1$ azimuthal orders. Plotting these modes, as in figure 7.12 reveals that they are the TE_{11} and TM_{01} modes, as was the case for the corrugated horn, as outlined in section 6.2.2. The Golay cell is sensitive to the overall composite field and so will sense the net field that is produced by the sum of the modes supported by the back to back system. Although the Golay cell is coupled to the incident

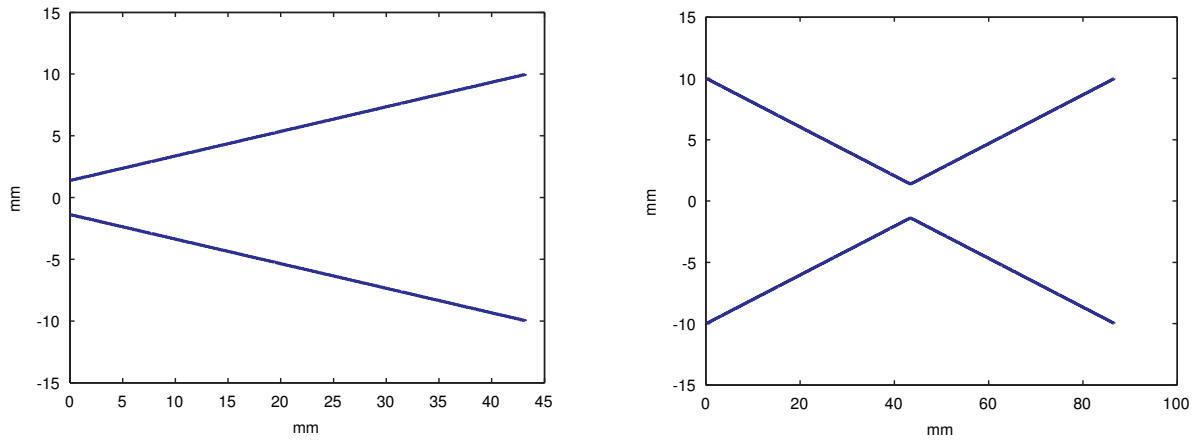

(a) Single horn

(b) Back to back arrangement

Figure 7.11: Schematic of the horn designed by the author and built at the University of Manchester, in both single and back to back configurations

Frequency [GHz]	**Modes Supported**
80	TE_{11}
85	TE_{11}, TM_{01}
105	TE_{11}, TM_{01}, TE_{21}
115	TE_{11}, TM_{01}, TE_{21}

Table 7.2: Modal content of the Manchester horn as a function of frequency

radiation via a section of rectangular waveguide, all of the incident field will couple into the waveguide and hence the detector, as the field is now in free space as opposed to waveguide modes. Additionally, although the rectangular waveguide section is sensitive to polarisation, the excitation source is polarised in the same manner and so all of the radiation should be detected.

The back to back horn system was mounted in such a way as to allow it to be rotated through some angle relative to the axial line of the source corrugated horn. When the main axis of the back to back system was collinear with the axis of the source horn, the two horns were facing each other. This corresponded to a zero degree off-axis angle of illumination. This orientation defines the optical axis of the system. The origin of the scan was defined such that it lay along the optical axis of the system at the same height as back to back and source horns. Scans were carried out for a number of angles of off-axis illumination. The full two dimensional grids are presented for each scan that was undertaken, along with the corresponding x and y cuts through these grids.

Prior to this, it is necessary to consider two factors.

Excitation Source: The code assumes plane wave illumination of the back to back horn arrangement, while the experimental setup uses a corrugated horn as the illumination source.

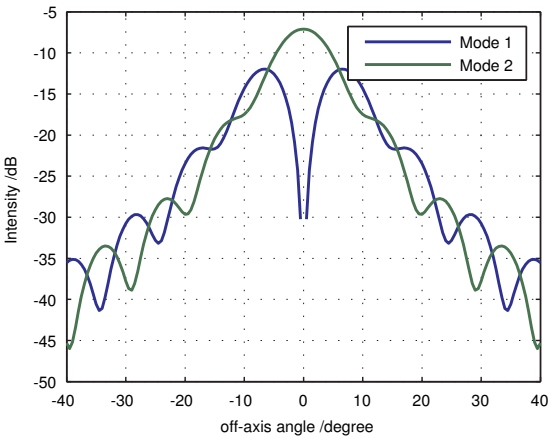

Figure 7.12: Modes supported by the back to back Manchester horn system, TM_{01} and TE_{11} respectively

Clearly a Gaussian field will emerge from the source horn as opposed to a plane wave, as shown in figure 7.9. This will result in theory, for a given illumination angle, differing amounts of each mode being excited relative to the plane wave excitation case, and so the resulting output field from the back to back system will be different. In the experimental setup described here, the source horn is placed very close to the entrance aperture of the back to back horn system, and so the back to back system will only intercept a small portion of the field from the corrugated horn which can be approximated as being planar. Although there will be discrepancies between the simulated and measured cases due to this, the systems will be comparable and so similar trends should be observed.

It is possible to experimentally generate a plane wave. For example, positioning the source horn a very large distance from the back to back arrangement, thus giving rise to a large propagation distance. This would result however in the power being very low by the time the field had propagated to the back to back horns. Additionally, spillover of the field would occur and illuminate the area around the back to back horns, distorting the results as the Golay cell would also detect this signal. It would also be possible to generate plane waves using combinations of lenses and/or mirrors. Such work will, in the future, be carried out at the University of Manchester.

Free Space Propagation: The Golay cell is not necessarily in the farfield of the back to back arrangement, thus in order to make comparisons with simulations, it may be necessary to process the simulated results as follows. The scattering matrices defining the horn, subject to plane wave illumination, are calculated and the aperture field of the system calculated according to this. Following the procedure detailed in sections 8.2.1 and 8.2.2, the aperture field is transformed from waveguide modes to free space modes and propagated through free space to the detector plane. This allows a comparison to be made between the simulated and measured beam patterns. In the example presented below, the detector is in the farfield of the

back to back arrangement and so it is possible to directly use the farfield output of SCATTER.

Figures 7.13 to 7.23 show the measured 2D farfield pattern of the back to back configuration of the Manchester horns for various off-axis angles of illumination. The measured voltages in a given cut were calibrated relative to the value at the origin for that cut, as outlined earlier, so as to account for the voltage drift observed in the Golay cell. These plots were then normalised to the on-axis voltage value in order to allow ready comparison between the measured and simulated cuts. The dynamic range of the system when a back to back horn is inserted between the source and the detector is naturally reduced. This reduction will be clear when cuts of the measured 2D grid are examined. The dynamic range reduces, as there is less power propagating through the system to the detector in this configuration, as opposed to the direct illumination of the detector by a source horn. This effect will be increased when off-axis angles of illumination are used, as even less power will now enter the back to back system and propagate to the detector. The dynamic range of the system is thus expected to decrease with increasing angle of illumination and this will be mentioned for each angle of illumination presented.

Table 7.3 shows the coupling of both of the supported modes in this system to a plane wave of varying off-axis illumination angle.

Off-axis angle [deg]	**TM_{01} coupling**	**TE_{11} coupling**
0	0.8699	0.9677
5	0.3002	0.6109
12	0.5321	0.0665
23	0.1427	0.0560
35	0.0651	0.0261

Table 7.3: Coupling of the TM_{01} and TE_{11} modes of the back to back Manchester horn to a plane wave for varying off-axis angle of illumination

Figure 7.13 displays the measured 2D field pattern for an illumination angle of zero degrees, i.e. on-axis illumination. A regular Gaussian profile is observed, as expected due to the modal content of the horn. This is emphasised by examining the cuts in figure 7.14a. As in the case of the corrugated horn, which has the same modal content, both modes couple strongly to an on-axis plane wave (see table 7.3) meaning that the composite farfield will have on-axis power. The simulated composite field (horizontal cut) is shown in figure 7.14, alongside a horizontal cut through the measured data, where the farfield pattern (simulated data) has been projected onto the xy axis, allowing the off-axis distance equivalent to the off-axis angle to be calculated and displayed for the purposes of comparison. Given the differences in the illumination source of the system, it is not expected that perfect agreement be noted, rather it is hoped that the trends predicted by the simulation are observed. Despite this, excellent agreement is noted

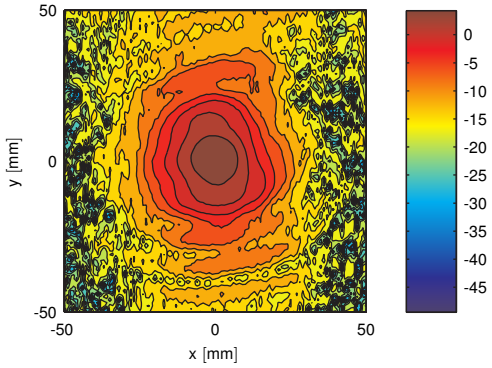

Figure 7.13: Measured field pattern for 0 degree illumination of the Manchester back to back horns

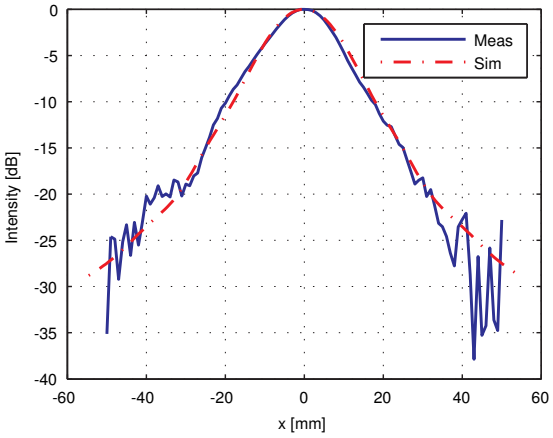

Figure 7.14: Comparison of measured (meas) and simulated (sim) horizontal cut through the co-polar field for 0° illumination of the Manchester back to back horn

between the simulated and measured field patterns in the case where the back to back horn is illuminated by an on-axis source. This is an excellent validation of the code. Given the levels where noise begins to appear, the dynamic range of the system in the on-axis illumination configuration is approximately -25 dB. The system will of course detect below this, however portions of the field pattern with power levels below the dynamic range should not be trusted and should be regarded as noise.

Increasing the angle of illumination to five degrees resulted in the field pattern shown in figure 7.15. The horizontal cut of this field for both the simulated and measured cases are shown in figure 7.16. The field is still of a regular shape, with on-axis power and a Gaussian distribution, something that is predicted by the simulations and shown by measurement. Referring to table 7.3, the TE_{11} couples to the plane wave twice as much as the TM_{01} mode, and so will be dominant in the composite field, giving rise to on-axis power. This large relative coupling is understood by noting that five degrees is not a very large off-axis angle of illumination, and so the plane wave will couple most strongly to the mode with the greatest amount of power about

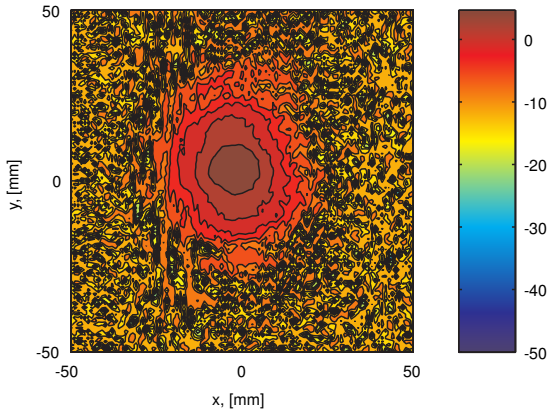

Figure 7.15: Measured field pattern for 5 degree illumination of the Manchester back to back horns

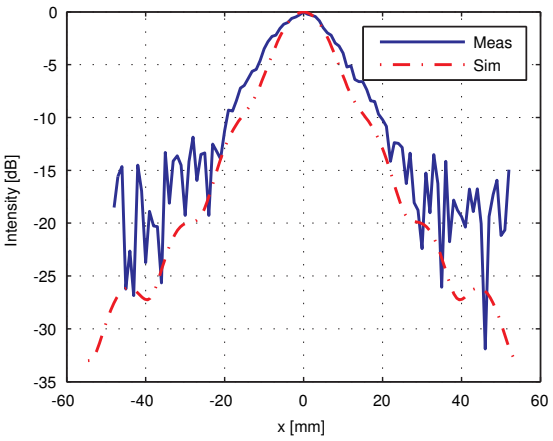

Figure 7.16: Comparison of measured and simulated horizontal cut through the co-polar field for 5° illumination of the Manchester back to back horn

the axis, which is the TE_{11} mode. Since the overall coupling of each mode is, as expected, less than in the on-axis case, the amount of power coupled into the system in total is also less, resulting in the intensity of the simulated composite field being lower. This is obscured by the fact that the comparison between the simulated and measured cases are normalised to 0 dB for the purposes of comparison, however the full 2D plot is not normalised in this manner and so the decrease in power is seen. A decrease in power is also noticed in the simulated results when they are not normalised to the maximum value in the field pattern. Of note is the fact that the measured field appears broader than the simulated field. This is due to the fact that in order to protect the Gunn diode, the back to back horn arrangement was rotated about its central vertical axis with respect to the source horn, as opposed to the source horn being rotated. Although necessary to protect the measurement apparatus, this is not ideal as now the Golay cell does not measure the field in a plane perpendicular to the direction of propagation, but rather it measures the projection of the field onto the scanner plane. With the scanning plane viewed from the exit aperture of the back to back horn, the back to back arrangement

was angled towards the right of the scanning plane. This will result in an apparent widening of the beam, however the result will not be linear as different parts of the projected beam (the parts furthest away from the centre of the scanner origin) will have further to travel, and so a smearing effect will be observed on this side of the beam, making it appear lopsided, which is what is observed. If the back to back horn is rotated then the resulting beam pattern will therefore be centred on a different part of the scanning plane relative to the centre. Thus, for each angle of illumination it was necessary to undertake alignment scans to locate the new centre of the measurement pattern and define this point as the new origin of the scan. Taking into account these effects, the agreement is still good, and the predicted trends match the measured ones, i.e. on-axis power is measured. For an angle of illumination of five degrees, the dynamic range of the system is approximately 15-16 dB

Figure 7.17 shows the measured 2D field pattern when the off-axis angle of illumination is increased to twelve degrees. The dynamic range of the system in this configuration is approximately 10-11 dB. Increasing the angle of illumination to this higher value results in an

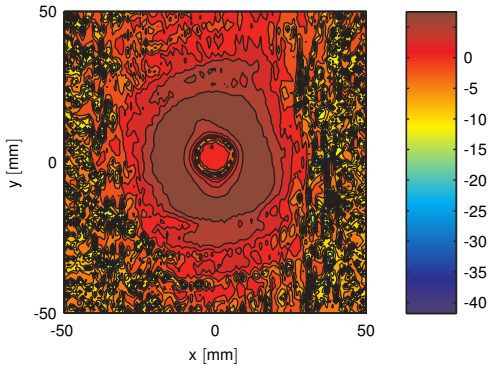

Figure 7.17: Measured field pattern for 12 degree illumination of the Manchester back to back horns

on-axis dip. The peak that was located on axis for zero and five degree angles of illumination has now been removed, a feature which can be noted by examining the cut presented in figure 7.18. This result implies that for this angle of illumination the relative balance between the modes has shifted, with the TM_{01} mode now dominant, resulting in the on-axis null. Examining table 7.3, this is seen to be the case. The TM_{01} mode couples to this plane to a significantly higher degree than the TE_{11} mode (which contributes relatively little), resulting in the predicted composite field more closely resembling the distribution of the TM_{01} mode. The measured field pattern matches this predicted trend with the on-axis dip emerging, and so the agreement between model and experiment is again satisfactory, taking into account the now greater projection effect. The measured on-axis dip extends approximately 20 dB down from the maximum value, taking an approximate average of the bottom value on either side of the dip. The predicted dip only extends to $\approx$-13.5 dB. It is difficult to measure the experimental angle of

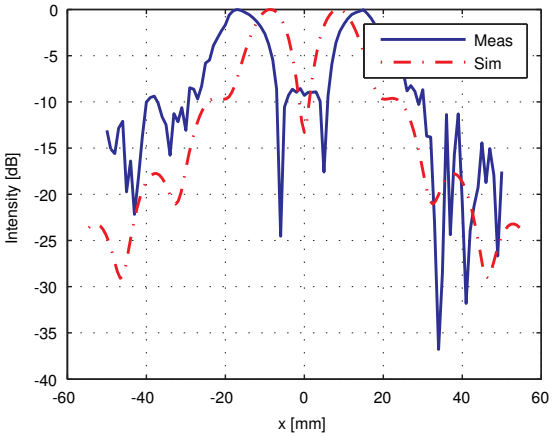

Figure 7.18: Comparison of measured and simulated horizontal cut through the co-polar field for 12° illumination of the Manchester back to back horn

offset accurately, and so it is entirely possible that what was measured to be a 12 degree angle of illumination could be a different angle to within one or two degrees either side of 12. In order to attempt to improve the agreement between the measured and simulated cuts, the simulation was repeated for angles of illumination of 11 and 13 degrees with the results shown in figure 7.19. It can be clearly seen that for an angle of illumination of 13 degrees, the simulated dip

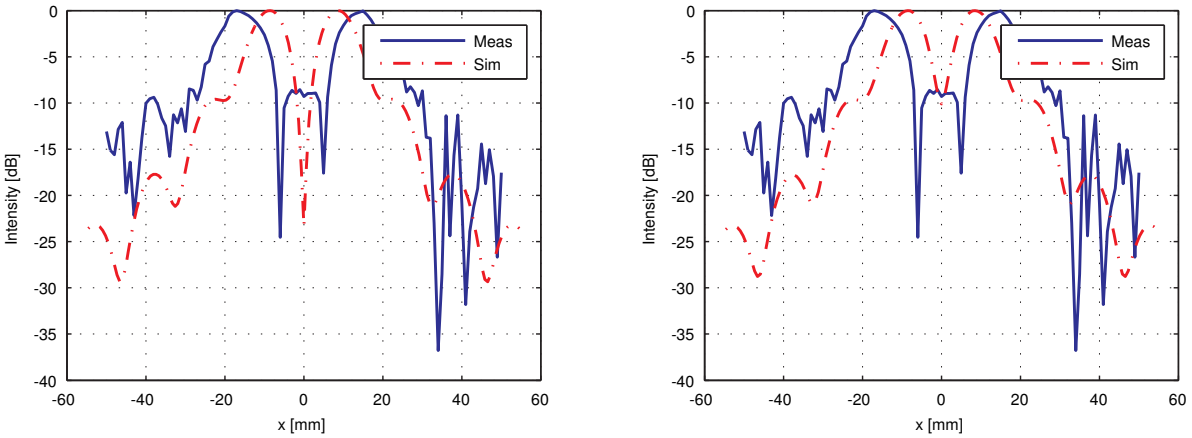

(a) 11°

(b) 13°

Figure 7.19: Comparison of measured and simulated horizontal cut through the co-polar field for 11° and 13° illumination of the Manchester back to back horn

reduces in depth, thus increasing the difference between the measured and simulated cuts. The simulated dip is seen to increase to a depth of ≈-23 dB when an angle of illumination of 11 degrees is used, which provides agreement with the measured data. Thus, by allowing for error in the measured angle of illumination and by also allowing for the projection effects described previously, excellent agreement is seen between the measured and simulated cuts, in terms of the code developed by the author predicting the effects of off-axis illumination.

Increasing the angle of illumination to twenty three degrees resulted in the 2D field in figure 7.20 being measured, with the corresponding cut shown in figure 7.21. The dynamic range of the system in this configuration is approximately 10 dB. Examining the coupling of each of

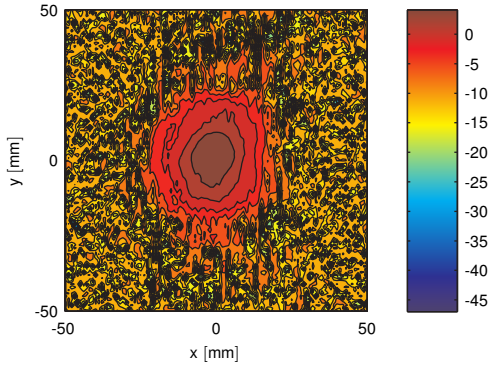

Figure 7.20: Measured field pattern for 23 degree illumination of the Manchester back to back horns

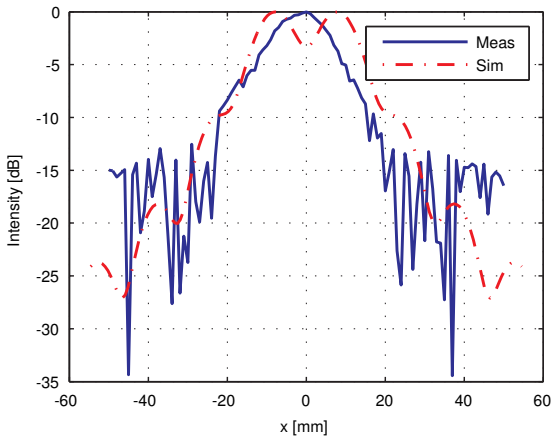

Figure 7.21: Comparison of measured and simulated horizontal cut through the co-polar field for 23° illumination of the Manchester back to back horn

the modes to the plane wave, again shown in table 7.3, shows that again the TM_{01} mode is dominant, although the power entering the system is severely reduced once again, resulting in a lower level field distribution. The expected low power levels and TM_{01} dominance are clearly shown in figure 7.21, where an on-axis dip is noted. It is not as extreme as in the previous case however, due to the fact that the difference in coupling between the two modes has now been severely reduced. Although the TM_{01} mode is still dominant, it is not dominant by as much, and so the increased balance between the modes yields a less extreme on-axis dip. The simulated cut therefore makes sense. Examining the measured 2D field, the lower power levels are clear, with an increased amount of noise clearly present in the measurement, which is visible from the cut. This is expected for such large off-axis angles of illumination and so structure will become less clear at this point. Thus, the slight on-axis dip predicted

by simulations is not necessarily measurable by this system. It is possible that at this angle the contribution from the TM_{01} mode, expressed in free space modes, does not couple as well into the rectangular WR10 waveguide that feeds the Golay cell. The coupling of the modes into rectangular waveguide modes in an effect that is not accounted for here. As before, angles of illumination a degree either side of 23 degrees were simulated in order to account for the error in measuring the experimental angle of illumination to see if the on-axis peak that was measured would appear. This is shown in figure 7.22. Examining these plots reveals that in

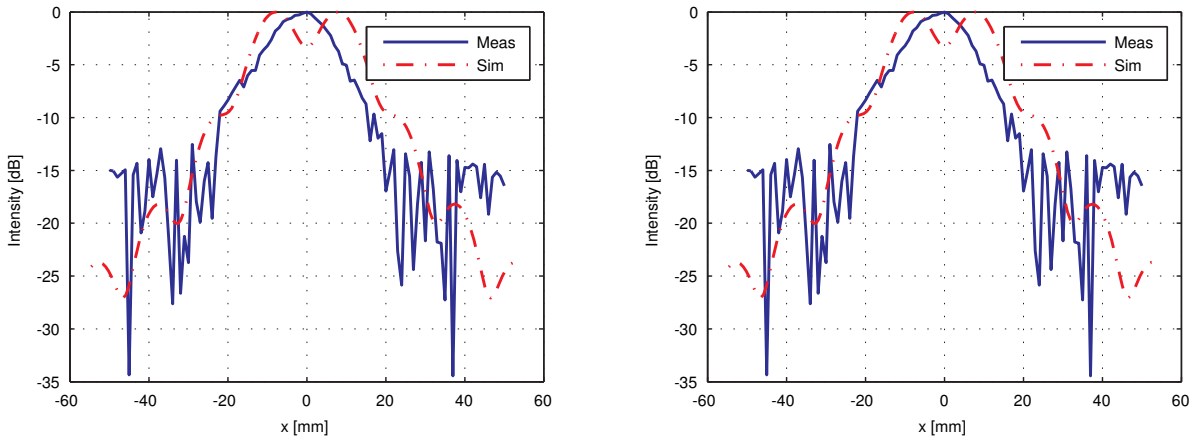

(a) 22° (b) 24°

Figure 7.22: Comparison of measured and simulated horizontal cut through the co-polar field for 22° and 24° illumination of the Manchester back to back horn

this case the on-axis maximum that was measured is not reproduced by simulations. An angle of twenty degrees was also examined, yielding no further improvement. This is not particularly an issue as the simulated dip is only of the order of 2.5 dB in depth, and so at the low levels of power in question for such a large off-axis illumination, it is entirely possible that this feature would not be observed.

Further increasing the angle of illumination to thirty five degrees resulted in the measured plots shown in figures 7.23 and 7.24. The large off-axis angle results in the TM_{01} mode still being dominant due to its off-axis power distribution, in addition to the overall coupling of each mode being extremely low, resulting in low levels of power entering and therefore emerging from the system. The balance in the coupling between the modes is also more even in this case. The predicted field should therefore be of a low level with a slight but noticeable on-axis dip. This is what is observed to be the case, as shown in figure 7.24. In terms of the measurement system, this angle of illumination is too large for sufficient power to enter the system in its current form, and so low level noise (normalised to 0 dB) is measured as the output, and so no comparison can be made with the simulated data.

For the various angles of off-axis illumination presented, the simulated results make intuitive sense. As the angle is increased, the TM_{01} mode becomes dominant due to its off-axis

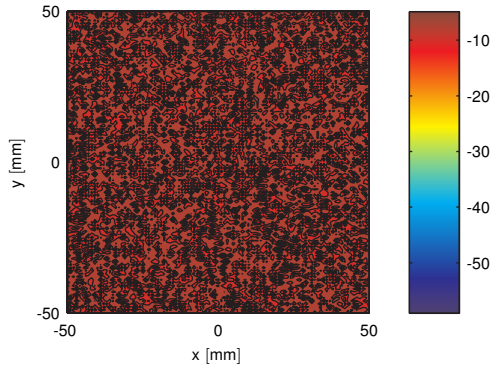

Figure 7.23: Measured field pattern for 35 degree illumination of the Manchester back to back horns

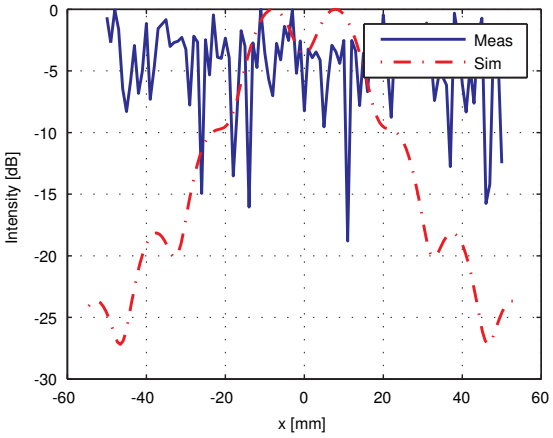

Figure 7.24: Comparison of measured and simulated horizontal cut through the co-polar field for 35° illumination of the Manchester back to back horn

distribution, and the varying balance between the modes is clear from examining the composite farfield pattern. The amount of power predicted to emerge from the system also decreases with increasing angle of illumination, as would be expected. The simulated data is therefore sensible. The same trends are seen to emerge from the measured data, showing that the model makes the correct predictions and that as a proof of concept, using the off-axis illumination of a back to back system of identical few-moded horns is a good way of measuring the performance of horns that are not single-moded.

7.5 Conclusions

In this chapter, the development and testing of a Golay cell based 100 GHz testbed for measuring the individual modes present in a back to back horn was presented as a proof of concept. A testbed was developed with associated software for controlling the entire measurement process. A corrugated conical horn and a pyramidal horn were both tested using the measurement

system, and the measured fields agreed with what would be expected for such horns. This validated the performance of the measurement system.

The testing of a pair of horns constructed at the University of Manchester (few-moded at 100 GHz), arranged in a back to back manner, was also carried out using the Golay cell testbed. The results clearly showed that as a function of illumination angle, the balance between the two modes supported by the horn shifted, with the TM_{01} mode becoming dominant over the TE_{11} mode for large off-axis angles due to its off-axis distribution. This resulted in on-axis dips appearing in the measured data, with this trend agreeing very well with those predicted using the corresponding model, as presented in the previous chapter.

The measured and experimentally predicted trends were shown to agree well, providing further verification of the modelling process, despite it being an initial step towards modelling off-axis illumination. The testbed was also shown to function very well, giving a dynamic range of approximately 30 dB. As a proof of concept, the presented work shows that it is possible to use such a back to back system to measure the performance of the various modes in few/multi-moded horn, a critical task in order to be able to better analyse few/multi-moded systems. In the future, VNA based measurements coupled with models of the system will allow the experimental retrieval of the scattering matrices of multi-moded horns and therefore allow their characterisation. The VNA measurements would also allow for the possibilities of both better system alignment and increased dynamic range.

Chapter 8

Including a Free Space Gap in SCATTER

8.1 Introduction

In this chapter, the implementation of a free space gap in SCATTER will be discussed, which will then be applied to a prototype SAFARI pixel, in order to quantify the optical efficiency of the system, including potential losses due to the presence of the gap. Such capabilities are also potentially capable of estimating the crosstalk between neighbouring pixels in the detector array that SAFARI will make use of. It is of course possible to model such systems, including the lossy materials from which the system is constructed, in full electromagnetic solving packages such as CST Microwave Studio, however systems such as these are electrically large, and even using distributed computing that includes several high power computer with an abundance of memory, such problems can be impossible to solve. To this end, it is therefore useful and practical to implement the capability to model such systems within SCATTER.

Using the framework that exists within SCATTER, it is possible to model many waveguide based systems, including those with bolometer detectors (modelled using absorbers placed in the waveguide) and cavities. It was necessary, for the purpose of modelling SAFARI like systems, as outlined in section 1.4, to develop a method to integrate a free space section within a waveguide system in SCATTER, thus considering this free space section to be a section within the waveguide system. A more detailed description of a SAFARI like pixel will be given in section 8.4 when the tools developed in the first few sections are applied in order to model it, however a brief description of the overall pixel will now be provided for context. The SAFARI system couples the radiation being measured to a feed horn, whose throat feeds a detector cavity. In reality pyramidal feed horns will be used, however conical horns are used in the prototypes and so they will be used in this thesis. There is a free space gap between the throat of the horn and the cavity which serves two purposes. The first is for mechanical reasons. The horn antenna array is cut from a solid section of metal, however the detector array is cut from

a silicon wafer. Since two materials are involved it is not possible to manufacture both arrays monolithically. It is also not possible to press them together, so they must be bolted together to form the pixel. The mechanical connections that enable this result in the gap emerging. The second purpose of the gap relates to the fact that to minimise damage to the absorber that is used in the detection system (as will be outlined in section 8.4) that would otherwise be experienced during mission launch, it is necessary to mount the absorber just outside the cavity, although for the purposes of this thesis the absorber will be considered to reside in the plane of the cavity entrance. The gap is sufficiently wide to achieve this protection, however its impact on the system must be investigated.

SAFARI like focal plane array systems are currently being experimentally investigated by SRON Netherlands Institute for Space Research (SRON). The focal plane array is described in section 1.4, while the pixels compromising the array will be described in section 8.4. In the current setup, the section which will be realised using this work as a free space gap is actually comprised of a section of waveguide of very large radius. Considering the spatial extent of the waveguide modes propagating in this section however, it may be considered to be effectively free space. This approximation is used, as the number of waveguide modes required to accurately model a section of waveguide of this radius would be extremely large. The results obtained from this very precise setup must be compared with models of the system in order that they may be fully understood and explained, and so that the models can be validated. Both of these requirements are critical so as the system as a whole can be understood and so as future predictions can be made. Software packages such as CST and COMSOL are being used to model the system in an attempt to explain the results, however in order to further validate these results, and to verify and make the system available in the more efficient mode matching framework, it shall be modelled in SCATTER. It is of course possible to model the system in SCATTER as is, despite the number of waveguide modes required, using a continuous waveguide structure, however the system experiences losses due to this gap and these will not be implemented correctly in SCATTER due to the fact that the waveguide structures are perfectly electrically conducting. As stated, it is valid to consider this section as a free space section, and so in implementing it in this manner, it is possible to introduce losses to the system. This shall now be discussed.

8.2 Implementing a Free Space Gap in SCATTER

In order to recreate the experimental setup that exists in SRON using the mode-matching technique, it is necessary to calculate the full set of scattering matrices for each part of the experimental setup and to then cascade these in sequence with the overall scattering matrices in order to represent the overall system. The system represents one pixel of a proposed focal plane array. Each horn in the array is connected to the next, as they are all cut from one piece of

material. Again, the SAFARI proposal makes use of pyramidal horns in the array, however the prototype tests carried out at SRON are conducted on conical horns. This material (into which the horns are cut) presents a reflecting surface to the radiation propagating within the system, as it impinges on the pixel, defined by the hemispherical cavity. This will be shown later in figure 8.11. It is therefore important to model this section. The free space gap will potentially allow power to leak from the pixel, which could give rise to crosstalk between the pixels. Since SCATTER assumes that complex power is conserved, a mechanism must be implemented in order to allow power to leave the system. This will be accomplished via truncating absorber sections, that absorb all power incident upon them at all frequencies from the free space mode set. The cavity also contains an ohmic sheet absorber such as that detailed in section 4.5.1 to model the detector.

Once the scattering matrices for the various parts of the system have been found and cascaded together, the system can be understood by considering the S_{11} matrix. The elements required to model such a system, in addition to the standard waveguide and ohmic sheet functionality already available in the standard SCATTER implementation are

- Waveguide to free space mode coupling (and vice versa)
- Horn array reflector section (simulates the reflective walls of the array in the plane of the horn throat)
- Free space loss section (to implement a loss mechanism, as SCATTER conserves complex power by default)
- Propagation through free space
- Cavity

The locations of these sections are shown in figure 8.1, which shows them in the context of a generic pixel of the type used in the conical SAFARI prototype being tested at SRON. The steps necessary to model the outlined parts of the system above, including how to cascade the matrices in order to achieve the overall matrices for the system, are now examined in detail.

8.2.1 Waveguide to Free Space Coupling

In this section, the coupling of waveguide modes into free space, and vice versa, will be considered, as outlined in [62, 63]. This occurs at the interface between the throat of the horn and the free space gap, and the interface between the free space gap and the cavity. The mode set selected to represent the free space modes is the Associated Laguerre Gaussian (LG) beam mode set. This is the most natural mode set to select for cylindrically symmetric systems such as the SRON SAFARI pixel prototype. This mode set is a solution to the wave equation in

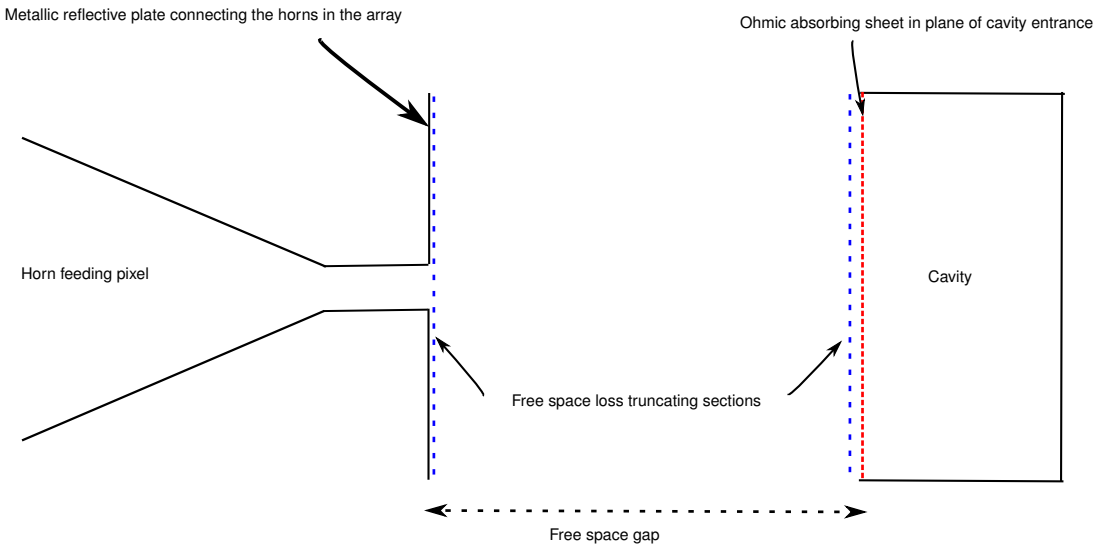

Figure 8.1: Locations of the sections needed to model a prototype SAFARI pixel. Not to scale.

the paraxial limit, which assumes that the variation in amplitude along the direction of propagation due to diffraction will be small over wavelength sized distances, when compared to the variation in the perpendicular direction, [45]. This applies at the wavelengths in question. The formalism invoked to describe and work with the fields as free space modes follows that laid out in [45]. The LG modes of order m and degree α can be expressed according to equation 8.1

$$\begin{pmatrix} \Psi_m^{\alpha,\cos}(r,\phi,z) \\ \Psi_m^{\alpha,\sin}(r,\phi,z) \end{pmatrix} = \sqrt{\frac{2\left(2-\delta_{0\alpha}\right)m!}{\pi W^2 (m+\alpha)!}} \exp\left(\frac{-r^2}{W^2}\right)\left(\frac{2r^2}{W^2}\right)^{\frac{\alpha}{2}} L_m^{\alpha}\left(\frac{2r^2}{W^2}\right) \exp\left(-jk\left(\frac{r^2}{2R}\right)\right) \times \exp\left(-j\left(kz-\phi_{m,\alpha(z)}\right)\right)\begin{pmatrix}\cos\alpha\phi \\ \sin\alpha\phi\end{pmatrix}, \quad (8.1)$$

where $L_m^{\alpha}(x)$ are the Associated Laguerre polynomials of order m and degree α. W, R and $\phi_{m,\alpha}$ are the width, phase radius of curvature and the phase slippage of each mode, each being a function of the distance propagated from the beam waist, the point at which the beam has minimum width. Typically $z = 0$ at this point. These parameters are defined, for the field relative to a horn aperture, as follows:

$$W(z) = \sqrt{W_0^2\left(1+\left(\frac{\lambda(z+\Delta)}{\pi W_0^2}\right)^2\right)} \quad (8.2)$$

$$R(z) = (z+\Delta)\left(1+\left(\frac{\pi W_0^2}{\lambda(z+\Delta)}\right)^2\right) \quad (8.3)$$

$$\phi_{m,\alpha}(z) = (2m + \alpha + 1)\left(\arctan\left(\frac{\lambda(z+\Delta)}{\pi W_0^2}\right) - \arctan\left(\frac{\lambda\Delta}{\pi W_0^2}\right)\right), \tag{8.4}$$

where the phase slippage is taken to be zero at the point where $z = 0$. W_0 is the width of the beam at the waist, i.e. the waist radius. Δ is the position of the waist behind the aperture of the horn, assuming that the mode set is being used for coupling between a horn and free space. This is not necessarily coincident with the phase centre of the horn. Δ and W_0 are given by

$$\Delta = \frac{R_a}{\left(1 + \left(\frac{\lambda R_a}{\pi W_a^2}\right)^2\right)} \tag{8.5}$$

and

$$W_0 = \sqrt{\frac{W_a^2}{1 + \left(\frac{\pi W_a^2}{\lambda R_a}\right)^2}}, \tag{8.6}$$

with W_a and R_a being the beam width and radius of curvature at the aperture of the horn. The forms of equations 8.2 to 8.4 differ slightly from the standard form of the equations typically encountered in Gaussian beam mode analysis. This is due to the fact that these equations take into account the fact that the waist of the beam is not necessarily coincident with the horn aperture by using equations 8.5 and 8.6.

Values for R_a and W_a can be calculated by fitting to the field being represented. For various specific horn designs their values can be readily calculated from standard formulas, for example for a single-moded conical corrugated horn, it is known that $W_a = 0.6435a$, where a is the aperture radius of the horn, and $R_a = L$, where L is the slant length of the horn. Exact values are not required however, as the choice of R_a and W_a is arbitrary. Their proximity to the correct values influences the efficiency of the LG modes. Essentially, the closer these values are to their correct values, the fewer modes that are required to represent the field at the aperture, thus reducing computational requirements.

For some choice of $W = W_a$ and $R = R_a$, an infinite set of these free space modes therefore exists such that a summation can be used to describe either the x or y component of any complex scalar field $E(r, \phi, z)$. Computationally, a finite number of modes must be used. The number of modes used determines the off-axis distance to which the reproduction of the field using the free space modes is accurate. Assuming N modes are used and r is the off-axis distance, then the field can be represented according to

$$E(r, \phi, z) = \sum_{\alpha,m}^{N} A_{\alpha,m} \Psi_m^{\alpha\cos}(W, R, r, \phi, z) + \sum_{\alpha,m}^{N} B_{\alpha,m} \Psi_m^{\alpha\sin}(W, R, r, \phi, z). \tag{8.7}$$

$A_{\alpha,m}$ and $B_{\alpha,m}$ are the modal coefficients, the amount of each mode present, and are given by

$$A_{\alpha,m} = \int_A E\left(r,\phi,z\right)\left[\Psi_m^{\alpha\cos}\left(W,R,r,\phi,z\right)\right]^* r\mathrm{d}r\mathrm{d}\phi \tag{8.8}$$

and

$$B_{\alpha,m} = \int_A E\left(r,\phi,z\right)\left[\Psi_m^{\alpha\sin}\left(W,R,r,\phi,z\right)\right]^* r\mathrm{d}r\mathrm{d}\phi, \tag{8.9}$$

where A is the surface area of the aperture that the transformation is occurring over. For a paraxial vector field it is necessary to calculate these modal coefficients for each of the x and y components, in order to retain the vector nature of the field.

It is also possible to represent the aperture field of the horn antenna by using the waveguide modes that are utilised in SCATTER. The LG modes display $\cos n\phi$ and $\sin n\phi$ angular dependence and will only couple to the waveguide modes that display the same angular dependence, [36, 37]. This reduces the number of coupling integrals required, and so to implement this it is useful to rewrite the standard waveguide mode equations in their $\hat{\mathbf{i}}$ and $\hat{\mathbf{j}}$ components in such a way so that terms containing trigonometric functions of $(n-1)\,\phi$ are grouped with Bessel functions of order $(n-1)\,\phi$, and similarly for terms with $(n+1)\,\phi$ appearing. This allows the waveguide modes to be expressed according to $\mathbf{e}_{nl}^{TM} = \mathbf{e}_{nl}^{TM(n-1)} + \mathbf{e}_{nl}^{TM(n+1)}$ and $\mathbf{e}_{nl}^{TE} = \mathbf{e}_{nl}^{TE(n-1)} + \mathbf{e}_{nl}^{TE(n+1)}$, where

$$\mathbf{e}_{nl}^{TM(n-1)} = \sqrt{\frac{2-\delta_{n0}}{\pi a^2 J_{n+1}^2\left(p_{nl}\right)}} J_{n-1}\left(p_{nl}\frac{r}{a}\right)\left(\begin{bmatrix}\cos\left(n-1\right)\phi\\ -\sin\left(n-1\right)\phi\end{bmatrix}\hat{\mathbf{i}} - \begin{bmatrix}\sin\left(n-1\right)\phi\\ -\cos\left(n-1\right)\phi\end{bmatrix}\hat{\mathbf{j}}\right) \tag{8.10}$$

$$\mathbf{e}_{nl}^{TM(n+1)} = \sqrt{\frac{2-\delta_{n0}}{\pi a^2 J_{n+1}^2\left(p_{nl}\right)}} J_{n+1}\left(p_{nl}\frac{r}{a}\right)\left(\begin{bmatrix}\cos\left(n+1\right)\phi\\ -\sin\left(n+1\right)\phi\end{bmatrix}\hat{\mathbf{i}} + \begin{bmatrix}\sin\left(n+1\right)\phi\\ -\cos\left(n+1\right)\phi\end{bmatrix}\hat{\mathbf{j}}\right) \tag{8.11}$$

$$\mathbf{e}_{nl}^{TE(n-1)} = \sqrt{\frac{2-\delta_{n0}}{\pi a^2\left(1-\left(\frac{n}{q_{nl}}\right)^2\right) J_n^2\left(q_{nl}\right)}} J_{n-1}\left(q_{nl}\frac{r}{a}\right)\left(\begin{bmatrix}\cos\left(n-1\right)\phi\\ -\sin\left(n-1\right)\phi\end{bmatrix}\hat{\mathbf{i}} - \begin{bmatrix}\sin\left(n-1\right)\phi\\ -\cos\left(n-1\right)\phi\end{bmatrix}\hat{\mathbf{j}}\right) \tag{8.12}$$

$$\mathbf{e}_{nl}^{TE(n+1)} = \sqrt{\frac{2-\delta_{n0}}{\pi a^2\left(1-\left(\frac{n}{q_{nl}}\right)^2\right) J_n^2\left(q_{nl}\right)}} J_{n+1}\left(q_{nl}\frac{r}{a}\right)\left(\begin{bmatrix}\cos\left(n+1\right)\phi\\ -\sin\left(n+1\right)\phi\end{bmatrix}\hat{\mathbf{i}} + \begin{bmatrix}\sin\left(n+1\right)\phi\\ -\cos\left(n+1\right)\phi\end{bmatrix}\hat{\mathbf{j}}\right), \tag{8.13}$$

and p_{nl} is the l^{th} root of $J_n\left(z\right)$ and q_{nl} is the equivalent for the derivative of the Bessel functions.

Due to the fact that the components of these waveguide modes that have a ϕ dependence are orthogonal, the waveguide modes only couple into vector combinations of the Laguerre modes given in equation 8.1.

When calculating the coupling between the mode sets and $n \neq 0$, the $\mathbf{e}_{nl}^{TM(n-1)}$ and $\mathbf{e}_{nl}^{TE(n-1)}$ components only couple to $\mathbf{\Psi}_m^{n-1}(r,\phi,z)$ and the $\mathbf{e}_{nl}^{TM(n+1)}$ and $\mathbf{e}_{nl}^{TE(n+1)}$ only couple to $\mathbf{\Psi}_m^{n+1}(r,\phi,z)$. For the case where $n=0$, the waveguide combinations $\mathbf{e}_{nl}^{TM(-1)}$, $\mathbf{e}_{nl}^{TE(-1)}$, $\mathbf{e}_{nl}^{TM(1)}$ and $\mathbf{e}_{nl}^{TE(1)}$ will only couple to $\mathbf{\Psi}_m^{1}(r,\phi,z)$, due to the fact that Laguerre modes of negative order are not meaningful. Using this, the waveguide modes can be written as

$$\mathbf{e}_{nl}^{TE} = \sum_{m}^{M} \left(T_{m,l}^{(n)} \mathbf{\Psi}_m^{(n-1)} + T_{m+M,l}^{(n)} \mathbf{\Psi}_m^{(n+1)} \right) \tag{8.14}$$

and

$$\mathbf{e}_{nl}^{TM} = \sum_{m}^{M} \left(T_{m,l+L}^{(n)} \mathbf{\Psi}_m^{(n-1)} + T_{m+M,l+L}^{(n)} \mathbf{\Psi}_m^{(n+1)} \right), \tag{8.15}$$

where the $T_{i,j}^{(n)}$ form the elements of a $2M \times 2L$ matrix, $T^{(n)}$, that is a transformation matrix for waveguide modes of azimuthal order n. $2M$ is the number of Laguerre modes used, and $2L$ is the total number of waveguide modes. The elements of the matrix are given by

$$T_{m,l}^{(n)} = \int_A \left(\mathbf{e}_{nl}^{TE(n-1)} + \mathbf{e}_{nl}^{TE(n+1)} \delta_{0n} \right) . \left[\Psi_m^{n-1(-)}(W,R,r,0,0) \right]^* r \mathrm{d}r\mathrm{d}\phi, \tag{8.16}$$

$$T_{m+M,l}^{(n)} = \int_A \left(\pm \mathbf{e}_{nl}^{TE(n-1)} \delta_{0n} + \mathbf{e}_{nl}^{TE(n+1)} \right) . \left[\Psi_m^{n+1(+)}(W,R,r,0,0) \right]^* r \mathrm{d}r\mathrm{d}\phi, \tag{8.17}$$

$$T_{m,l+L}^{(n)} = \int_A \left(\mathbf{e}_{nl}^{TM(n-1)} - \mathbf{e}_{nl}^{TM(n+1)} \delta_{0n} \right) . \left[\Psi_m^{n-1(-)}(W,R,r,0,0) \right]^* r \mathrm{d}r\mathrm{d}\phi, \tag{8.18}$$

$$T_{m+M,l+L}^{(n)} = \int_A \left(\pm \mathbf{e}_{nl}^{TM(n-1)} \delta_{0n} + \mathbf{e}_{nl}^{TM(n+1)} \right) . \left[\Psi_m^{n+1(+)}(W,R,r,0,0) \right]^* r \mathrm{d}r\mathrm{d}\phi, \tag{8.19}$$

where $T_{m,l}^{(0)} = T_{m,l+L}^{(0)} = 0$ and the $T_{m+M,l}^{(0)}$ and $T_{m+M,l+L}^{(0)}$ terms have a $\pm$ term with each sign representing a set of orthogonal modes. This reflects the fact that the waveguide modes of azimuthal order zero have zero degeneracy, whereas for a non-zero azimuthal order the waveguide modes have a degeneracy of two, accounted for by the fact that both terms, when the $\pm$ is allowed for, are non-zero. The transmission coefficient for corresponding orthogonal (degenerate) modes is the same.

Using this, the aperture field at the horn for each independent spatially coherent field within

the waveguide can be written according to

$$\bar{\mathbf{e}}^i = \sum_j^{2L} A_j^i \mathbf{e}_j^G = \sum_j^{2L} \left(\sum_m^M T_{m,j}^{(n)} A_j^i \mathbf{\Psi}_m^{n-1} + T_{m+M,j}^{(n)} A_j^i \mathbf{\Psi}_m^{n+1} \right), \tag{8.20}$$

with j being equal to either l or $l + L$, depending on whether the mode in question is TE or TM, respectively. With this analysis complete, it is now possible to construct a scattering matrix that transforms waveguide modes into free space modes.

Consider for example, the $n = 1$ azimuthal order of a system that contains 3 TE and 3 TM modes that are to be coupled into 12 Laguerre modes, split evenly between modes of order 0 and order 2 (This is due to the fact that waveguide modes couple into Laguerre modes of order $\alpha = n \pm 1$, as discussed previously). When the transmission coefficients are calculated, this results in the 12×6 matrix shown,

$$\begin{array}{c|cccccc}
 & TE_{11} & TE_{12} & TE_{13} & TM_{11} & TM_{12} & TM_{13} \\
\hline
\Psi_1^{(0)} & T_{1,1}^{(1)} & T_{1,2}^{(1)} & T_{1,3}^{(1)} & T_{1,4}^{(1)} & T_{1,5}^{(1)} & T_{1,6}^{(1)} \\
\Psi_2^{(0)} & T_{2,1}^{(1)} & T_{2,2}^{(1)} & T_{2,3}^{(1)} & T_{2,4}^{(1)} & T_{2,5}^{(1)} & T_{2,6}^{(1)} \\
\Psi_3^{(0)} & T_{3,1}^{(1)} & T_{3,2}^{(1)} & T_{3,3}^{(1)} & T_{3,4}^{(1)} & T_{3,5}^{(1)} & T_{3,6}^{(1)} \\
\Psi_4^{(0)} & T_{4,1}^{(1)} & T_{4,2}^{(1)} & T_{4,3}^{(1)} & T_{4,4}^{(1)} & T_{4,5}^{(1)} & T_{4,6}^{(1)} \\
\Psi_5^{(0)} & T_{5,1}^{(1)} & T_{5,2}^{(1)} & T_{5,3}^{(1)} & T_{5,4}^{(1)} & T_{5,5}^{(1)} & T_{5,6}^{(1)} \\
\Psi_6^{(0)} & T_{6,1}^{(1)} & T_{6,2}^{(1)} & T_{6,3}^{(1)} & T_{6,4}^{(1)} & T_{6,5}^{(1)} & T_{6,6}^{(1)} \\
\hline
\Psi_1^{(2)} & T_{7,1}^{(1)} & T_{7,2}^{(1)} & T_{7,3}^{(1)} & T_{7,4}^{(1)} & T_{7,5}^{(1)} & T_{7,6}^{(1)} \\
\Psi_2^{(2)} & T_{8,1}^{(1)} & T_{8,2}^{(1)} & T_{8,3}^{(1)} & T_{8,4}^{(1)} & T_{8,5}^{(1)} & T_{8,6}^{(1)} \\
\Psi_3^{(2)} & T_{9,1}^{(1)} & T_{9,2}^{(1)} & T_{9,3}^{(1)} & T_{9,4}^{(1)} & T_{9,5}^{(1)} & T_{9,6}^{(1)} \\
\Psi_4^{(2)} & T_{10,1}^{(1)} & T_{10,2}^{(1)} & T_{10,3}^{(1)} & T_{10,4}^{(1)} & T_{10,5}^{(1)} & T_{10,6}^{(1)} \\
\Psi_5^{(2)} & T_{11,1}^{(1)} & T_{11,2}^{(1)} & T_{11,3}^{(1)} & T_{11,4}^{(1)} & T_{11,5}^{(1)} & T_{11,6}^{(1)} \\
\Psi_6^{(2)} & T_{12,1}^{(1)} & T_{12,2}^{(1)} & T_{12,3}^{(1)} & T_{12,4}^{(1)} & T_{12,5}^{(1)} & T_{12,6}^{(1)}
\end{array}.$$

The top half of the first column represents how the TE_{11} mode couples into the Laguerre modes of order 0, and the bottom half of the column gives the coupling into the Laguerre modes of order 2. The second column gives the same information for the TE_{12} mode, and so on until the fourth column is reached. At this point the TM modes are considered. In this way, the coupling of each waveguide mode into the entire set of free space modes is defined within the matrix. A separate matrix exists for each azimuthal order of waveguide modes being considered in the system. Since each column represents how much of each free space mode is excited by a given waveguide mode, the entries in a given column are coherent, but columns are incoherent with respect to each other, as in the standard scattering matrix formulation for waveguide systems.

This matrix therefore contains the coefficients for each Laguerre mode, and so knowing this each mode can be propagated to some plane of interest and the field at this point calculated

by summing the modes in this plane, or the field can be coupled back into some waveguide structure. This latter point will now be considered.

The transformation matrix governing the transformation from free space modes to waveguide modes of azimuthal order n is denoted by $R^{(n)}$, with entries given by $R^{(n)}_{i,j}$, according to

$$R^{(n)}_{l,m} = \int_A \left(\mathbf{e}^{TE(n-1)}_{nl} + \mathbf{e}^{TE(n+1)}_{nl}\delta_{0n}\right)^* . \left[\Psi^{n-1(-)}_m (W, R, r, \phi, z)\right] r\mathrm{d}r\mathrm{d}\phi, \tag{8.21}$$

$$R^{(n)}_{l,m+M} = \int_A \left(\pm\mathbf{e}^{TE(n-1)}_{nl}\delta_{0n} + \mathbf{e}^{TE(n+1)}_{nl}\right)^* . \left[\Psi^{n+1(+)}_m (W, R, r, \phi, z)\right] r\mathrm{d}r\mathrm{d}\phi, \tag{8.22}$$

$$R^{(n)}_{l+L,m} = \int_A \left(\mathbf{e}^{TM(n-1)}_{nl} - \mathbf{e}^{TM(n+1)}_{nl}\delta_{0n}\right)^* . \left[\Psi^{n-1(-)}_m (W, R, r, \phi, z)\right] r\mathrm{d}r\mathrm{d}\phi, \tag{8.23}$$

$$R^{(n)}_{l+L,m+M} = \int_A \left(\pm\mathbf{e}^{TM(n-1)}_{nl}\delta_{0n} + \mathbf{e}^{TM(n+1)}_{nl}\right)^* . \left[\Psi^{n+1(+)}_m (W, R, r, \phi, z)\right] r\mathrm{d}r\mathrm{d}\phi, \tag{8.24}$$

where z is the distance from the waist (where the waveguide modes were originally transformed to free space modes) to the plane at which the free space modes are to be transformed back into waveguide modes, i.e. the propagation distance. ϕ is the on-axis phase slippage at this plane relative to an on-axis plane wave.

The waveguide modes can be normalised in terms of the power contained in the electric field such that $\int |\mathbf{E}|\,\mathrm{d}\mathbf{S} = 1$ or in terms of the complex power contained in the modes, such that $\int \mathbf{E} \times \mathbf{H}^* \mathrm{d}\mathbf{S} = 1$. In the former case, with the propagation distance set to zero, the transformation matrix for the free space to waveguide transition is given by $R^{(n)}_{i,j} = T^{(n)*}_{j,i}$. In the latter case, where the normalisation is carried out with respect to the complex power contained in the modes and the propagation distance is again zero, the transformation matrix is given by $R^{(n)}_{i,j} = T^{(n)*}_{j,i} / |Z|$, where $|Z|$ is the impedance of the appropriate waveguide mode and $|Z| \approx Z_0$ for large apertures, where Z_0 is the impedance of free space. This is also the impedance of each of the free space modes, according to the paraxial approximation. Since SCATTER uses the conservation of complex power, the latter normalisation is used for the waveguide modes and so in order to retain normalisation between the free space and waveguide modes, the transformation matrix is scaled according to the impedance of the free space modes. Since an equivalent normalisation is used for the two mode sets, this allows the losses due to diffraction to be estimated directly.

The reason for the different forms of the transformation matrix is as follows. Consider a waveguide mode $\bar{\mathbf{e}}_{nl}$ that exhibits $\cos n\phi$ and $\sin n\phi$ dependence. This can be expressed in terms of Laguerre modes according to

$$\bar{\mathbf{e}}_{nl} = \sum_m T^{(n)}_{m,l}\psi^n_m. \tag{8.25}$$

Due to the orthogonality of the Laguerre mode set, it is clear that $T^{(n)}_{m',l} = \int \psi^{n*}_{m'}.\bar{\mathbf{e}}_{nl}\mathrm{d}A$. Expressing the Laguerre modes in terms of waveguide modes yields

$$\psi^n_m = \sum_l R^{(n)}_{l,m}\bar{\mathbf{e}}_{nl}, \tag{8.26}$$

with $R^{(n)}_{l',m} = \frac{\int \psi^n_m.\bar{\mathbf{e}}^*_{nl'}\mathrm{d}A}{\int \bar{\mathbf{e}}^*_{nl'}.\bar{\mathbf{e}}_{nl'}\mathrm{d}A}$. Therefore, if the modes are normalised using $\int \bar{\mathbf{e}}^*_{nl}.\bar{\mathbf{e}}_{nl}\mathrm{d}A = 1$, then $R^{(n)}_{l',m} = \int \psi^n_m.\bar{\mathbf{e}}^*_{nl'}\mathrm{d}A$. If the modes are normalised by $\int \bar{\mathbf{e}}_{nl} \times \bar{\mathbf{h}}^*_{nl}\mathrm{d}A = 1$, then $R^{(n)}_{l',m} = \int \psi^n_m.\bar{\mathbf{e}}^*_{nl'}\mathrm{d}A/\left|Z_{nl}\right|$. Respectively, this is equivalent to the expressions encountered previously; $R^{(n)}_{i,j} = T^{(n)*}_{j,i}$ and $R^{(n)}_{i,j} = T^{(n)*}_{j,i}/\left|Z\right|$, with the assumption that the modes are not propagated from their original position at $z = 0$.

The scattering matrices are defined such that $S^{i,j}_{21} = T_{i,j}$, $S^{i,j}_{12} = R_{i,j}$ and $S_{11} = S_{22} = 0$. The four scattering matrices required to represent the interface between a horn aperture and free space have now been obtained, so this interface can be represented as a step within SCATTER. If the field then propagates through free space and arrives at some other component such as a detector cavity, then the same formulation as above is applied to model that transition as a step within SCATTER, however in this case the S_{21} matrix will represent the transition from free space to waveguide modes and the S_{12} matrix will represent the transition from waveguide to free space modes.

8.2.2 Propagation Through Free Space

Consider a free space field that is represented as a sum of Laguerre Gaussian modes, each scaled by their corresponding mode coefficient. If the beam width (W_{in}), radius of curvature (R_{in}) and on-axis phase slippage (ϕ_{in}) are known at some (input) plane, then the field at that plane can be reconstructed according to

$$E_{in} = \sum A_n \Psi_n \left(W_{in}, R_{in}, \phi_{in}\right). \tag{8.27}$$

If it is desired to compute this field at a new plane that is located a distance z away, i.e. to compute the field after having propagated it a distance z, then it is necessary to track the evolution of the beam width, radius and phase slippage as the beam propagates. Carrying this out yields the values of these parameters at the output plane, namely W_{out}, R_{out} and ϕ_{out} respectively. This allows the field at the output plane to be reconstructed using an equation of the form

$$E_{out} = \sum A_n \Psi_n \left(W_{out}, R_{out}, \phi_{out}\right), \tag{8.28}$$

where the mode coefficients, A_n, are the same as for the input set, with the new values for the beam parameters accounting for the phase and spatial evolution of the beam over the propagation distance, [45].

Standard geometrical rays can be propagated through a system using the ABCD matrix method. There are similarities between the propagation of these rays and the propagation of Gaussian beams through the same system, within the paraxial limit. The ABCD matrix method can therefore be adapted to long wavelength systems by invoking the idea of complex rays.

In Gaussian beam systems, the waist position of the beam corresponding to the point at which the beam width is a minimum, has an on-axis maximum for the intensity. The fundamental Gaussian beam mode can be considered to be a spherical wave that is generated from a complex source. This is implemented by allowing the field to have a complex radius of curvature, q. If the point source is approximated to have a field distribution given by $E = \exp\left(-jkr^2/2q\right)$, then the complex beam parameter can be written, [45], as

$$\frac{1}{q} = \left(\frac{1}{q}\right) - j\left(\frac{1}{q}\right). \tag{8.29}$$

To allow this representation of the complex beam parameter to remain consistent with that of the fundamental Gaussian beam mode, equation 8.29 is written so that the real portion is equal to the reciprocal of the phase radius of curvature of the Gaussian beam and so that the complex portion represents a Gaussian variation in the off-axis amplitude. Using this, the complex beam parameter can therefore be written as

$$\frac{1}{q} = \frac{1}{R} - j\left(\frac{\lambda}{\pi W\left(z\right)^2}\right). \tag{8.30}$$

Combining this with the equation representing the field distribution of the point source yields the following expression for the full field distribution,

$$E = \exp\left(\frac{-jkr^2}{2q}\right) = \exp\left(\frac{-jkr^2}{2R}\right)\exp\left(\frac{-\lambda}{\pi W\left(z\right)^2}\right). \tag{8.31}$$

The complex beam parameter therefore tracks the key parameters of a Gaussian beam, and so by tracking its evolution through a system, the beam parameters are also tracked. Since the beam parameter is directly analogous to the radius of curvature in geometrical optics, the ABCD matrix approach will first be considered within the framework of geometrical optics.

Figure 8.2 illustrates the propagation of a beam between an input and an output plane, through a uniform medium. At the input plane, the beam, emanating from the complex point source, is fully described by considering its vertical position x and its angle of incidence i, relative to the optical axis and input plane respectively. This gives rise to a radius of curvature at the input plane of R_{in}. The beam is similarly described at the output plane (relative to the image of the complex point source), using the corresponding parameters x_o, t, and R_{out}. It is assumed that the output parameters (which allow the calculation of R_{out}), are linear functions

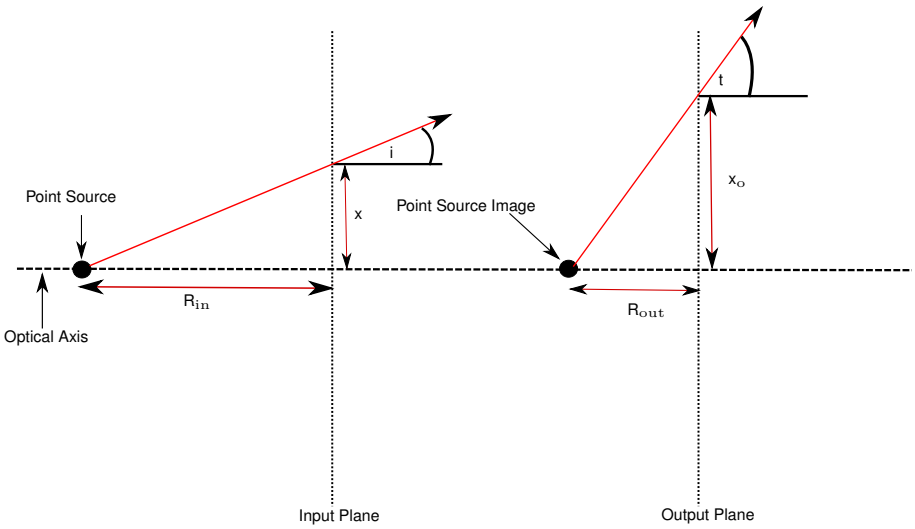

Figure 8.2: Propagation of a beam between two planes in a uniform medium

of the input parameters, and so can be related to the inputs using a 2×2 matrix such that

$$\begin{bmatrix} x_o \\ t \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} x \\ i \end{bmatrix}. \tag{8.32}$$

For an on-axis point source producing an on-axis image point source, as is the situation in figure 8.2, $R_{\text{in}} = \frac{x}{i}$ and $R_{\text{out}} = \frac{x_o}{t}$, if the paraxial approximation is used; $\tan i \approx i$. These expressions for the input and output and output radii can be used with equation 8.32 to yield an expression for the output radius of curvature in terms of the input radius and the components of the ABCD matrix such that

$$R_{\text{out}} = \frac{AR_{\text{in}} + B}{CR_{\text{in}} + D}. \tag{8.33}$$

If the components of the ABCD matrix act so as to describe the affect of the optical system on the portions of the beam that they act on, then if their values are known, the system can be described fully by using this approach.

To apply this to long wavelength systems, the geometric radius of curvature is replaced by its long wavelength equivalent, the complex beam parameter q, and the optical system in question is represented by a suitable ABCD matrix such that

$$q_{\text{out}} = \frac{Aq_{\text{in}} + B}{Cq_{\text{in}} + D}. \tag{8.34}$$

Recalling the definition of the complex beam parameter given in equation 8.30, the phase radius of curvature of the beam at the output plane can therefore be found using

$$\frac{1}{R_{\text{out}}} = \Re\left(\frac{1}{q_{\text{out}}}\right) = \Re\left(\frac{Cq_{\text{in}} + D}{Aq_{\text{in}} + B}\right) = \Re\left(\frac{C + D/q_{\text{in}}}{A + B/q_{\text{in}}}\right), \tag{8.35}$$

which can be further expanded to

$$\frac{1}{R_{\text{out}}} = \Re\left(\frac{C + D/R_{\text{in}} - jD\lambda/\pi W_{\text{in}}^2}{A + B/R_{\text{in}} - jB\lambda/\pi W_{\text{in}}^2}\right). \tag{8.36}$$

Considering the imaginary part of the full definition for q allows an expression to be found, [45], for the beam radius at the output plane,

$$W_{\text{out}} = \sqrt{\frac{-\lambda}{\pi\Im\left(\frac{C+D/R_{\text{in}}-jD\lambda/\pi W_{\text{in}}^2}{A+B/R_{\text{in}}-jB\lambda/\pi W_{\text{in}}^2}\right)}} \tag{8.37}$$

It is also critical to calculate the phase slippage between the modes as they propagate through the system, and this is also accomplished by means of the ABCD matrix method. For a known phase slippage at the input plane, ϕ_{in}, the phase slippage in an optical system is given by

$$\phi_{\text{out}} - \phi_{\text{in}} = -\text{Arg}\left(A + B\left(1/q_{\text{in}}\right)\right). \tag{8.38}$$

For a given set of Gaussian beam input parameters, R_{in}, W_{in} and ϕ_{in}, it is therefore possible using the above method to calculate their equivalent values at the output of some quasi-optical system that is represented by the components of an ABCD matrix. All that is required are the values of these components. Two simple examples of ABCD matrices are those for propagation through a uniform medium and for a thin lens. If the beam is to be propagated a distance d through a uniform medium, then the components of the ABCD matrix are given by

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} = \begin{bmatrix} 1 & d \\ 0 & 1 \end{bmatrix}, \tag{8.39}$$

and for propagation through a thin lens of focal length f,

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ -f^{-1} & 1 \end{bmatrix}. \tag{8.40}$$

For a system consisting of several components, an overall ABCD matrix is calculated by obtaining the matrix governing each individual element in the system and starting with the first element on the input side, pre-multiplying this matrix for each successive element. For example, consider a system consisting of a lens of focal length f, that is located a distance d_1 from the input plane of the system and a distance d_2 from the output plane, with the system remaining in free space at all times other than at the lens. The system therefore consists of three elements;

1. Free space propagation from the input plane to the thin lens, a distance d_1

2. A thin lens of focal length f

3. Free space propagation from the thin lens to the output plane, a distance d_2

The overall ABCD matrix for this system, M_{system}, is the sequential product of the matrices governing the three individual steps outlined above, and is given by

$$M_{\text{system}} = \begin{bmatrix} 1 & d_2 \\ 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 \\ -1/f & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & d_1 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 - d_2/f & d_1 + d_2 - d_2 d_1/f \\ -1/f & 1 - d_1/f \end{bmatrix}. \tag{8.41}$$

The components of M_{system} yield A, B, C and D, and so for a given set of input parameters, the output parameters can be readily computed. As stated previously, this allows the calculation of the output field of the system, according to equation 8.28.

For the purposes of modelling the free space gap between the throat of the horn and the input of the detector cavity, as in the SRON setup, the gap is assumed to be a uniform medium, namely that of free space. Therefore, the optical system being considered consists of only one section and so only one matrix is required to describe it, given by equation 8.39, where d is the length of the free space gap. The values for the beam parameters at the input plane (the throat of the horn), are known as they were assigned when the waveguide modes were coupled into free space. Using the ABCD matrix method, the parameter values can be calculated at the output plane which is coincident with the plane in which the opening of the detector cavity resides. Using these new parameter values, the free space modes can be coupled back into waveguide modes at this plane, yielding a set of scattering matrices in the formulation of SCATTER, as previously described. Due to the fact that the coupling into the cavity is carried out using the beam parameters at the output, the coupling takes into account the propagation of the free space modes across the free space gap and so the free space to cavity scattering matrices include the propagation contribution.

The free space loss truncating section and array plate reflector (to model the presence of the material into which the horns are cut) sections are now considered. The backing reflector exists beyond the extent of the horn exit aperture and cavity entrance aperture and serves to represent the connection between the different parts of material that the horns are mounted through. Considering one pixel in isolation, it is necessary to model this section, as it serves to trap and reflect radiation within the pixel by means of multiple reflections between the plates. After a certain number of reflections, some power will encroach on space beyond which the spatial extent of the pixel, and this power is effectively lost to the system. This loss must also be modelled. The loss is implemented by means of a free space loss truncating section, and the backing reflector by means of a perfect reflector. The implementation of these sections will now be analysed in detail.

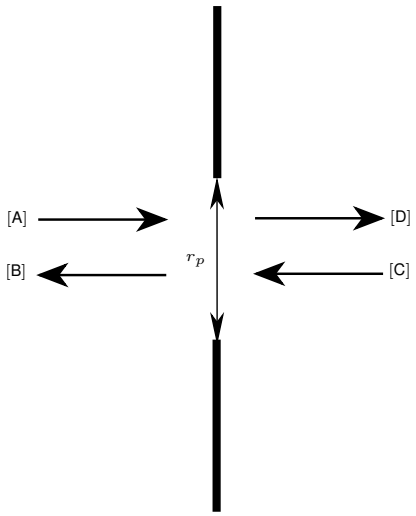

Figure 8.3: Free space loss section, simulating power loss beyond pixel extent

8.2.3 Free Space Loss Truncating Section

Truncating loss sections are placed on the free space side of the transition planes between the free space and waveguide mode sets and allow power to be removed from the system (giving two in total), that is they truncate the free space mode set beyond the pixel radius and suppress any reflections beyond this radius. This effectively simulates diffractive loss through the free space gap in the pixel. Consider the situation illustrated in figure 8.3. The value r_p is the radius of a single pixel of the system being modelled. Thus, if any mode extends beyond the pixel, the power contained within that mode beyond the extent of the pixel will be lost.

The mode sets on either side of the sheet are assumed to be identical free space modes, i.e. they are defined over the same spatial extent such that

$$\int_{-\infty}^{\infty} e_m . e_n^* \mathrm{d}r = \delta_{mn}. \tag{8.42}$$

The field can however travel either forward or backwards, where in figure 8.3, left to right is taken as the forward direction. At a given spatial point, the forward and backwards propagating fields differ by a phase curvature term such that

$$e^- = e^+ \exp\left(\frac{jkr^2}{2R}\right), \tag{8.43}$$

where k is the free space wavenumber, r is the off-axis distance, and R is the phase radius of curvature at the point in space being considered. Using the mode matching technique and

assuming a total of N free space modes, in the region where $r < a$ we have that

$$\sum_{n}^{N} \left(A_n e_n^+ + B_n e_n^- \right) = \sum_{n}^{N} \left(C_n e_n^- + D_n e_n^+ \right) \tag{8.44}$$

for the electric field. Laguerre-Gaussian mode sets are scalar, and so for the magnetic field the same mode set can be used, such that in the region $r < a$ it is also true that

$$\sum_{n}^{N} \left(-A_n e_n^+ + B_n e_n^- \right) = \sum_{n}^{N} \left(C_n e_n^- - D_n e_n^+ \right). \tag{8.45}$$

If r_p is the pixel radius, then in the region where $r < r_p$, equation 8.44 is valid for the electric field, and equation 8.45 is valid for the magnetic field component. For $r > r_p$ then all incident radiation is absorbed, meaning that no reflections occur. Equation 8.46 represents this situation. Since no reflections occur, this equation includes the effects of both the electric and magnetic fields.

$$\begin{aligned} \sum_{n}^{N} B_n e_n^- &= 0 \\ \sum_{n}^{N} D_n e_n^+ &= 0. \end{aligned} \tag{8.46}$$

At this point, several expressions are introduced.

$$P_1 = \int_0^{r_p} e_n^+ \cdot e_n^+ 2\pi r \mathrm{d}r, \tag{8.47}$$

$$P_2 = \int_0^{r_p} e_n^- \exp\left(\frac{jkr^2}{2R} \right) \cdot e_n^- 2\pi r \mathrm{d}r, \tag{8.48}$$

$$Q_1 = \int_{r_p}^{r_p} e_n^+ \cdot e_n^+ 2\pi r \mathrm{d}r \tag{8.49}$$

and

$$Q_2 = \int_{r_p}^{r_p} e_n^- \exp\left(\frac{jkr^2}{2R} \right) \cdot e_n^+ 2\pi r \mathrm{d}r. \tag{8.50}$$

The P_n integrals represent the coupling of modes on one side of the section to the modes on the other side of the section. As such, they represent the transmission of the mode sets. The Q_n integrals represent the self-coupling of the mode sets, or the reflections.

Considering equations 8.44, 8.45 and 8.46 and taking their dot product with e_n^+ and using

equation 8.43, then the following is found to be true;

$$\begin{aligned} P_1 A + P_2 B &= P_2 C + P_1 D \\ -P_1 A + P_2 B &= P_2 C - P_1 D \end{aligned} \tag{8.51}$$

and

$$\begin{aligned} Q_2 B &= 0 \\ Q_1 D &= 0. \end{aligned} \tag{8.52}$$

Solving 8.51 simultaneously yields the solutions

$$P_2 B = P_2 C, \tag{8.53}$$

and

$$P_1 A = P_1 D. \tag{8.54}$$

By conservation, we must have that $P_1 + Q_1 = \mathbb{I}$ and $P_2 + Q_2 = \mathbb{I}$. Using the identities in equations 8.53 and 8.54, along with 8.52, gives

$$D = P_1 A, \tag{8.55}$$

and

$$B = P_2 C. \tag{8.56}$$

When the scattering matrix is applied to any general system, the system can be expressed in terms of the input and output mode coefficients according to

$$\begin{bmatrix} [B] \\ [D] \end{bmatrix} = \begin{bmatrix} [S_{11}] & [S_{12}] \\ [S_{21}] & [S_{22}] \end{bmatrix} \begin{bmatrix} [A] \\ [C] \end{bmatrix}, \tag{8.57}$$

which, when solved, gives

$$[B] = [S_{11}]\, A + [S_{12}]\, C \tag{8.58}$$

$$[D] = [S_{21}]\, A + [S_{22}]\, C. \tag{8.59}$$

Comparing these solutions to equations 8.55 and 8.56 gives that for the truncating absorber section:

$$\begin{aligned} S_{11} &= 0 \\ S_{12} &= P_2 \\ S_{21} &= P_1 \\ S_{22} &= 0. \end{aligned} \tag{8.60}$$

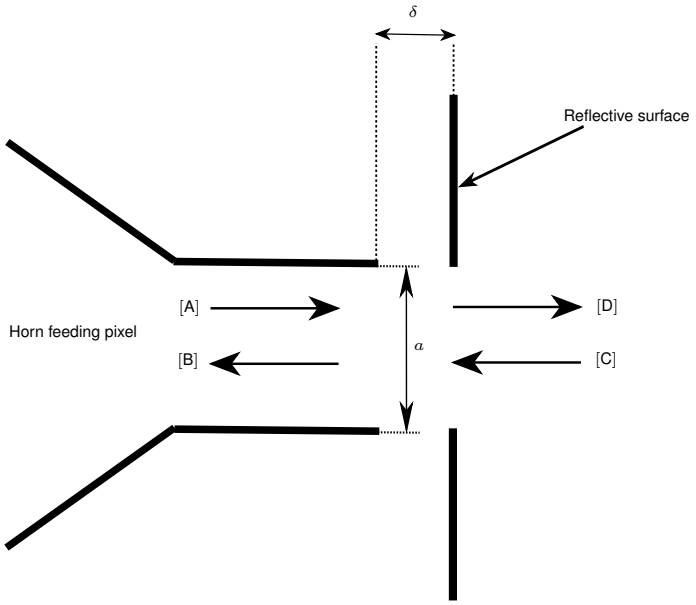

Figure 8.4: Reflective plate section simulating the portion of the material that the horn array is formed from that impinges on the section of the pixel between the horn throat and the pixel limit

8.2.4 Horn Array Reflective Plate

The metal plate that the horn array is cut from presents a reflective surface to the pixel between the outer radius of the pixel and the horn throat, and this will now be examined. Consider figure 8.4, a general case in which a waveguide forming the throat of a horn meets a reflective surface that extends from the entrance aperture of the horn in a direction that is perpendicular to the axis of the horn. In this case, a waveguide of radius a (the throat of the horn feeding the system) is separated from a reflective surface by a distance δ. The reflective surface has a pupil of radius a, centred on the waveguide, where a is the radius of the exit aperture of the horn. In the limit that $\delta \to 0$, the reflective plate is in contact with the waveguide, representing the setup in the pixel. In this case, no radiation incident upon the sheet from the left exists beyond a distance a off-axis. Thus, $[A]$ and $[B]$ exist only in the region $r < a$. In order to implement the reflective plate in SCATTER, the waveguide modes at the opening of the throat are transformed into free space modes as previously described, and these modes are immediately applied to the plate. If the plate is assumed to be infinitely thin, then this process results in the modes propagating a zero distance in free space, with the plate effectively truncating the field, expressed in free space modes. In the SCATTER implementation, a waist is assumed to be present at the throat of the horn and hence at the plate, however the more general situation when the sheet is located away from a waist will be considered.

In the region where $r < a$, equations 8.44 and 8.45 are both still valid. For the region where $r > a$, $[A] = [B] = 0$ on the waveguide side. On the free space side, the sheet is assumed to be

a perfect reflector, giving

$$\sum_{n}^{N} D_n e_n^+ = -\sum_{n}^{N} C_n e_n^-. \tag{8.61}$$

Following the same procedure as that outlined for the absorber, the scattering matrices for the reflector are found to be given by

$$\begin{aligned} S_{11} &= 0 \\ S_{12} &= P_2 \\ S_{21} &= P_1 \\ S_{22} &= -Q_2, \end{aligned} \tag{8.62}$$

which is intuitive. In this case, the P integrals are evaluated from $r : 0 \rightarrow a$ and the Q integrals from $r : a \rightarrow r_p$. A second reflector and absorber are located at the entrance to the cavity. The absorber section is symmetric with respect to either side of the expression, and so the scattering matrices calculated above are correct in either case. The reflector section is not symmetric however. Applying the same analysis as outlined previously, to the case where the waveguide is located to the right of the reflector and the free space section to the left, yields the set of solutions given by equation 8.63;

$$\begin{aligned} S_{11} &= -R_1 \\ S_{12} &= P_2 \\ S_{21} &= P_1 \\ S_{22} &= 0, \end{aligned} \tag{8.63}$$

where R_1 is the same as Q_1, simply calculated for the other side of the junction. For this situation, a is no longer the radius of the exit aperture of the horn, but the cavity radius.

The final step of the system will now be considered, the implementation of a cavity.

8.2.5 Implementation of a Cavity in SCATTER

The cavities utilised in this context consist of a waveguide structure with the end closed such that the surface area vector of the closing surface is co-linear with the propagation directions of the waveguide modes within the cavity structure. In the case where the cavity is a waveguide of uniform radius, this is implemented within SCATTER by considering a geometry that consists of two sections;

- **Section 1:** Radius and length equal to those of the cavity, R_{cav} and L_{cav}.
- **Section 2:** Radius < 0.00001 mm and length $> \lambda$, so as to avoid evanescent modes as this is the final step in the overall system.

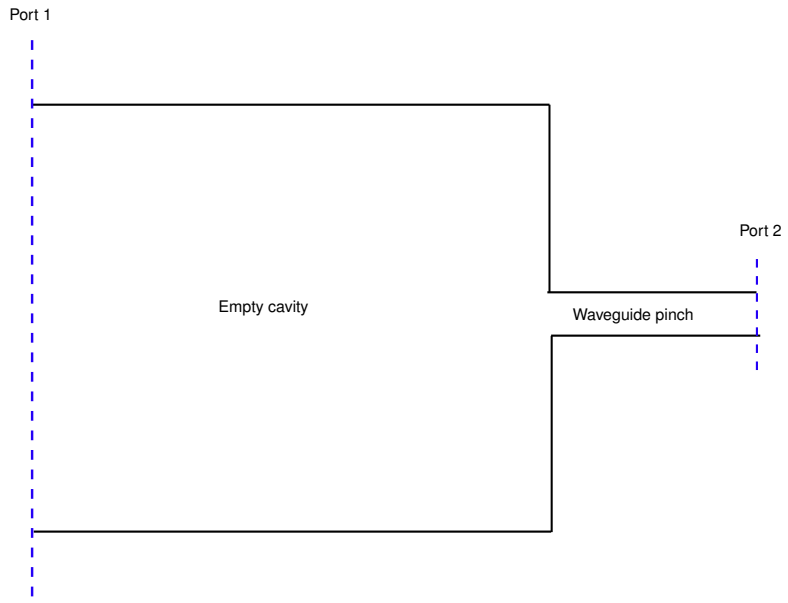

Figure 8.5: Empty cavity, showing the physical cavity and the waveguide pinch section which simulates the closing of the cavity

The first section serves to model the cavity itself in terms of mode content and phase shift as the modes propagate along its length. The second section models the closing of the back of the cavity. It achieves this by exploiting the fact that at a given frequency each waveguide mode has a critical value for the radius of a section of waveguide, that, if the radius falls below this value, the mode will no longer be supported and so will not propagate. The radius for this section is chosen so that it is well below the value required for the first waveguide mode, TE_{11}, and so neither this mode nor any other mode are supported in the final section. To ensure that no modes are excited in this section, the radius is chosen to be of order 0.00001mm. The length is chosen to be $> \lambda$ so that if any modes were to be excited, evanescence would not be an issue. Since 'pinching' the waveguide to this small value forbids modes from entering it from the cavity section, and also forbids modes to enter the cavity when this small section is excited as occurs in SCATTER, it acts as an effective backshort, as though the back of the cavity were closed as is the case in reality.

The scattering matrices for the second section can be defined manually as a result of this, or SCATTER can be allowed to calculate them. The results in either case are the same, and in this work, SCATTER will be allowed to calculate the scattering matrices. Calling the input on the cavity aperture side port 1, and the input on the pinched waveguide side port 2, as shown in figure 8.5, then the following is observed. Since no modes are supported in the pinched waveguide section, none of the power incident at port 2 can be coupled into the system, and so no power can propagate through to port 1 or be reflected within the system and be returned to port 2. Thus,

$$S_{12}^{i,j} = S_{22}^{i,j} = 0 \forall i, j \in [1, N]\,, \tag{8.64}$$

where N is the total number of waveguide modes ($\frac{N}{2}$ TE and $\frac{N}{2}$ TM modes), and so the scattering matrices are of dimension $N \times N$.

Considering the illumination of port 1, since no modes can propagate to port 2, it is found that S_{21} is also a zero matrix of dimension $N \times N$. Thus, by the conservation of power used in SCATTER, all of the power that enters port 1 must be reflected within the cavity and arrive back at port 1, therefore being accounted for by the S_{11} matrix. Naturally, various resonances will be present in the cavity due to standing waves, however this manifests itself in the complex phase information, and so if power is considered it will be found that all power that enters the cavity will return to the entrance to the cavity.

These four matrices govern the overall behaviour of the cavity, namely that all power that goes into the cavity must come back out again, and that no power can enter from the closed side. If one assumes that a real cavity were made from lossless material, then this is exactly what is expected in terms of cavity behaviour and so the model is valid. It has been analysed and verified in detail in [57].

The cavities encountered in the analysis of a SAFARI pixel are hemispherical. This is realised by reading in a geometry file that contains the radii and lengths of the sections required to build up the hemispherical cavity structure. but with a final extra section that is not of sufficient radius to support any modes. This effectively closes the back of the cavity, as outlined above.

8.2.6 Considering a SAFARI Like Pixel

The previous sections detailed the calculation of the scattering matrices for each of the elements required to model a SAFARI like pixel within SCATTER. In order to model the entire pixel, it is necessary to combine these elements. This is realised in the same manner as when combining two adjacent waveguide sections in SCATTER, i.e. by cascading the matrices for each element in the correct order. This means that each new element effectively corresponds to a step in SCATTER. The order in which the elements must be combined is as follows:

- Propagate the waveguide modes through the horn antenna in the standard manner
- At the horn exit aperture to the cavity region:
  - Couple the waveguide modes to the free space modes
  - Implement a reflective plate out to the boundary of the pixel to simulate the horn array plate
  - Implement a free space loss truncating section to remove power beyond the boundary of the pixel
- Propagate the free space modes across the length of the free space gap (50 μm)

- At the plane of the entrance to the hemispherical backshort cavity:
  - Implement a reflective plate out to the boundary of the pixel (if the cavity aperture does not extend to the pixel radius)
  - Implement a free space loss truncating section beyond the boundary of the pixel
  - Couple the free space modes to the waveguide modes in the cavity
- Propagate the waveguide modes through the ohmic sheet absorber (located in the plane of the cavity entrance) and the cavity in the standard manner

In order to cascade the matrices in the correct order, it is necessary to work through the system in the order indicated above. Consider two adjacent elements in a system, a and b, such that the radiation is nominally propagating from section a into section b. The scattering matrices for each section are represented as

$$[S^a] = \begin{bmatrix} [S^a_{11}] & [S^a_{12}] \\ [S^a_{21}] & [S^a_{22}] \end{bmatrix} \tag{8.65}$$

and

$$\left[S^b\right] = \begin{bmatrix} \left[S^b_{11}\right] & \left[S^b_{12}\right] \\ \left[S^b_{21}\right] & \left[S^b_{22}\right] \end{bmatrix}. \tag{8.66}$$

When the matrices for the two sections are cascaded, they result in the matrix $[S^c]$, given by

$$[S^a] = \begin{bmatrix} [S^c_{11}] & [S^c_{12}] \\ [S^c_{21}] & [S^c_{22}] \end{bmatrix}, \tag{8.67}$$

where this matrix represents the entire system from the input up to the output of section b. The elements of $[S^c]$ are defined as

$$\begin{aligned} [S^c_{11}] &= [S^a_{12}] \left[[I] - \left[S^b_{11}\right] [S^a_{22}]\right]^{-1} \left[S^b_{11}\right] [S^a_{21}] + [S^a_{11}] \\ [S^c_{12}] &= [S^a_{12}] \left[[I] - \left[S^b_{11}\right] [S^a_{22}]\right]^{-1} \left[S^b_{12}\right] \\ [S^c_{21}] &= \left[S^b_{21}\right] \left[[I] - [S^a_{22}] \left[S^b_{11}\right]\right]^{-1} [S^a_{21}] \\ [S^c_{22}] &= \left[S^b_{21}\right] \left[[I] - [S^a_{22}] \left[S^b_{11}\right]\right]^{-1} [S^a_{22}] \left[S^b_{12}\right] + \left[S^b_{22}\right]. \end{aligned} \tag{8.68}$$

$[S^c]$ then becomes $[S^a]$ (i.e. the entire system to that point becomes section a), and this is then cascaded with the matrices for the next section. This iterative process is repeated until all sections of the system have been considered, at which point $[S^c]$ gives the scattering matrices that govern the entire system.

Since for the final section of the cavity portion of a pixel, all scattering matrices are populated with zeros apart from the S_{11} matrix, when these get cascaded with the rest of the system at the final step, it will result in the overall set of scattering matrices governing the system

exhibiting this property. This makes sense, as when the pixel is considered as a whole, port 1 is the horn aperture and port 2 is the rear of the cavity. Since power cannot enter or exit the system through the rear of the cavity, all scattering matrices relating to port two will consist of zeros. Thus, the only matrix containing any power will be the S_{11} matrix.

One final consideration in implementing the pixel in SCATTER is dealing with evanescent modes in the system. These are modes whose cut-off frequency is above the frequency being considered in the simulation. These modes can be excited in SCATTER (with unity power), however as they are not supported by the system they are attenuated if the initial system is long in terms of wavelength. Although evanescent modes do physically exist throughout the system, they do not couple well to free space and so they must be suppressed, otherwise erroneous results may occur. This is achieved in SCATTER by extending the first section in the geometry file to a length that is several wavelengths long ($\approx$10 λ for example) but disregarding the phase shift that occurs along the length of this section. Thus, the evanescent modes propagate along a sufficiently long section such that they are attenuated prior to reaching the first horn junction, but the phase shift is not accounted for, thus the physical system being simulated has not altered.

8.3 Verification of Transition and Propagation Steps

In this section, the free space gap modelling process will be verified. This will be achieved by calculating and comparing the nearfield and farfield beam patterns of a horn antenna in both waveguide and free space modes. A system containing a free space gap will also be constructed in CST MWS (Computer Simulation Technology Microwave Studio) and the results obtained from this will be compared to those obtained from the author's SCATTER development, with the free space gap implemented, allowing the overall calculation including all components to be validated.

8.3.1 Verification of Waveguide/Free space Transition

The horn under consideration in both of these systems is a standard corrugated horn, designed for operation at 93 GHz. This specific horn geometry is used, as it is well understood from previous analysis, [36, 37], and is therefore used here for verification. A schematic of the horn is shown in figure 8.3.1. The key parameters of this horn are listed in table 8.3.1.

In order to test the transition between the waveguide and free space modes, the scattering matrices governing the horn were calculated, at a frequency of 90 GHz, and using 40 waveguide modes which were split evenly between TE and TM modes. The aperture field of the horn was calculated in the standard manner using these scattering matrices.

If the transition to free space modes is implemented correctly, then following the calculation

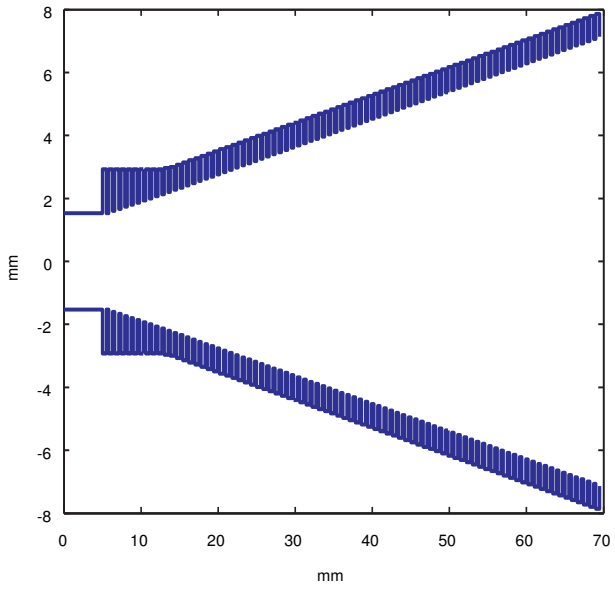

Figure 8.6: Horn antenna used in the testing process

Throat Radius [mm]	1.535
Aperture Radius [mm]	7.1345
Slant Length [mm]	≈65
Throat Length [mm]	5
λ [mm]	3.3

Table 8.1: Physical parameters of horn used to verify waveguide/free space transition code

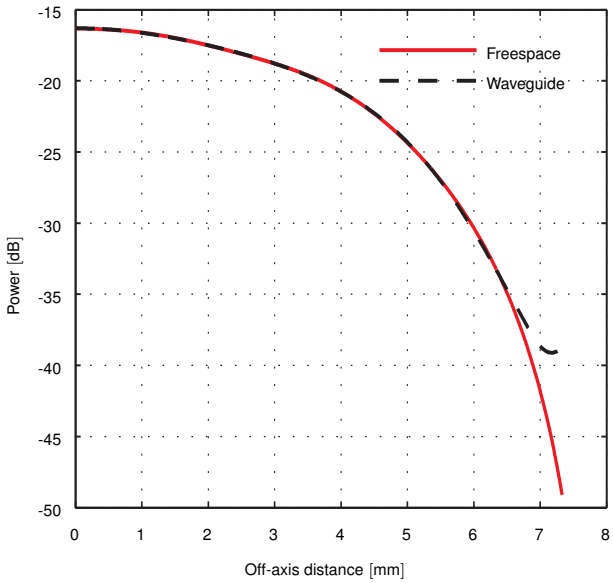

Figure 8.7: Aperture field of the horn antenna, calculated using waveguide and free space modes

of the transformation matrix and the subsequent cascading of this with the matrices for the horn, then the field reconstructed from this using free space modes should match the waveguide aperture field, assuming that the aperture of the horn is large in terms of wavelength, giving little impedance mismatch between the aperture and free space.

The transformation matrix was calculated as previously outlined, one for each azimuthal order being considered in the calculation (six azimuthal orders were considered, ranging from $n = 0$ to $n = 5$). In order to retain square matrices at all points in the calculation, 40 free space modes were used that were split evenly between Laguerre Gaussian functions of degree $n - 1$ and $n + 1$. Where $n = 0$ is being considered, the portion of the transformation matrix dealing with the LG functions of degree 1 are calculated as normal, however the portion dealing with the functions of order -1 is populated with zeros. The resulting fields are compared in figure 8.7. It is clear that there is excellent agreement between the field produced in SCATTER using waveguide modes, and the field produced when free space modes are used. Divergence only occurs between the fields once the power drops to below -35 dB relative to the on-axis power. This is at a negligible level, and divergence would be expected at this point. Laguerre-Gaussian modes also only retain an accurate description of the beam being represented up to some off-axis distance that is an increasing function of the number of free space modes used to describe the beam, [38]. Were divergence noted for higher levels of power then this would likely be the cause and so it could be rectified by increasing the number of free space modes used.

The agreement between the two fields illustrates the fact that the original S_{21} matrix for the waveguide modes at the horn aperture has been successfully transformed into a scattering matrix that represents the same information, albeit using a difference basis set, the free space

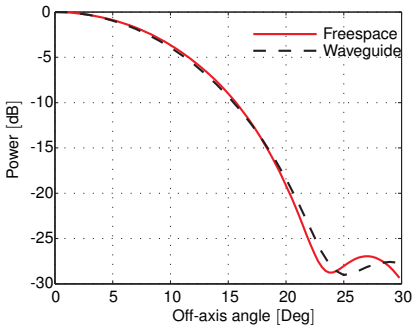

Figure 8.8: Farfield, calculated using SCATTER Fourier transform method and by free space propagation (FS prop)

modes. This serves to validate the transformation technique. When the waveguide aperture field is propagated to the farfield (accomplished by using a Fourier transformed version of the Bessel functions) and the free space field is propagated a large distance using the ABCD matrix method, the two fields give excellent agreement, as shown in figure 8.8. There is a high level of agreement between the two methods. It would be necessary to use an infinite number of modes to get perfect agreement, however this level of agreement obtained using 40 modes is more than sufficient. This result validates the propagation code (which will be further validated when CST is used to analyse an entire pixel) and serves to again validate the waveguide/free space transformation code.

8.3.2 Verification of Power Loss Calculation

In order to verify the approach taken to estimate the power lost from the system due to the gap, which could reduce the detected power, it was necessary to implement a waveguide fed cavity in CST MWS and use a field monitor around the structure in order to measure the power leaked from the structure by means of the gap between the waveguide and the cavity. An equivalent system was then implemented in SCATTER using the steps outlined previously in section 8.2 and compared to the results obtained from CST.

A cavity of length 1 mm and radius 2 mm was implemented in CST. This was separated by some variable distance from a waveguide of radius 1.535 mm and length 5 mm. A backshort was added to the end of the waveguide facing the cavity, such that this portion of the system was of total radius 2 mm, matching the spatial extent of the cavity. This simulates the structure that would surround the waveguide were it mounted on a focal plane. All structures were constructed from perfect electrically conducting (PEC) material, and the leaked power was measured at 70 GHz ($\lambda \approx 4.29$ mm). A CAD rendering of the system, as constructed in CST, is shown in figure 8.3.2 from two angles, allowing the waveguide and the cavity to both be clearly seen.

The equivalent system was implemented in SCATTER. Since the material used in CST is lossless, then like SCATTER, the total power lost from the system will be due to power

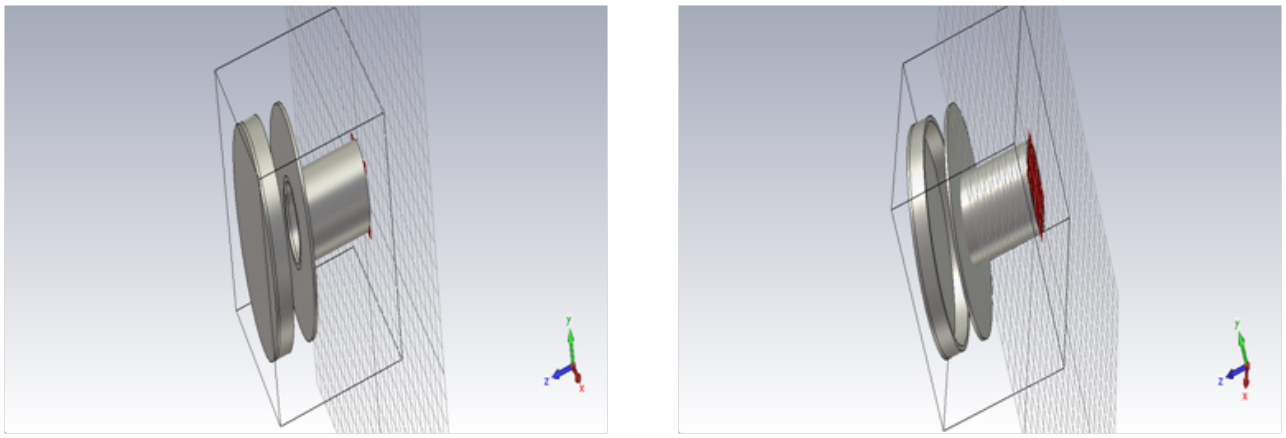

Figure 8.9: CAD views of the cavity test system, as in CST

leaking out via the free space gap. Since the mechanism for power loss is the same in either case, a comparison can be made between the figures. At a frequency of 70 GHz, the waveguide feeding the cavity is single-moded, however the cavity itself is sufficiently large to support the TM_{01} mode in addition to the TE_{11} mode. Since the waveguide is feeding the system, and no inter-azimuthal order scattering is possible when dealing with cylindrical waveguides, only modes with $n = 1$ can enter the cavity, and so $n = 0$ need not be considered in the calculations as the TM_{01} mode will not be excited at any point. It is therefore correct to consider only the $n = 1$ azimuthal for this calculation.

The leaked power can be read directly from CST using the power monitor that was placed around the structure. It is however necessary to calculate the leaked power in SCATTER by means of examining the scattering matrices. As previously outlined, in a closed cavity system, if the cavity is fed through port one, then the S_{11} is the only scattering matrix that will contain any power, as no power can enter or leave the system through port two. In the absence of a free space gap, then each mode supported by the system in the S_{11} matrix would contain unity power. This means that all of the power that enters the system through port one exits the system through port one, as SCATTER excites each mode with unity power. This is correct, as due to conservation of power,an explicit assumption in the current version of SCATTER, there is no mechanism for power to be lost, and so it all must be reflected.

As outlined in chapter 3, if a singular value decomposition (SVD) is performed on the S_{11} matrix, then it is decomposed into the product of three matrices, such that

$$S_{11} = U\Sigma V^{\dagger}. \tag{8.69}$$

If N is the number of waveguide modes considered in the calculation (i.e. S_{11} is an $N \times N$ matrix), then Σ is an $N \times N$ diagonal matrix with non-negative real numbers on the main diagonal (the singular values of the matrix) and zeros elsewhere. The non-zero values are

arranged in descending order starting with the entry in row one, column one, and only the first β values are non-zero. The columns of U that have an associated non-zero singular value form a new compact basis set for expressing the fields on the output side of the side, and likewise for V on the input side, where each mode is excited with unity power. Only the first β columns are required to fully represent the original scattering matrix.

The matrix of interest in this case is the S_{11} matrix, which means that the input and output sides of this matrix both refer to port one, the excitation port. Thus, the columns of U represent the modes that are supported on a round trip through the entire system. If the system supports m modes, i.e. only the first m singular values are non-zero, then a list of the singular values corresponding to supported modes can be found, $\sigma_i = \Sigma_{i,i}, i \leq m$. The square of each singular value represents the power in the associated mode relative to unity, as each mode is assumed to be excited with unity power. If the free space gap is removed, then there is no mechanism for power loss and so each mode must contain the same amount of power as it was excited with, in this case unity power, so $\sigma_i^2 = 1 \implies \sigma_i = 1 \forall i \leq m$. In the example presented above, only the TE_{11} mode will be supported on a round trip through the system, as the waveguide acts as a filter. Thus, the system is single-moded and $m = 1$.

If a free space gap is now added to the system, then it is expected that a leakage of power should occur across this gap. There is now therefore a mechanism for power to be lost from the system. For the power that enters such a system (unity power times the number of modes supported in a round trip of the system), it must all be accounted for either at the throat of the horn following a round trip through the system (and therefore it will appear in the S_{11} matrix) or via the lost (radiated) power. There is no other mechanism to account for power due to port two being closed, and as a result of this it may be concluded that any power that is not accounted for in the S_{11} matrix must have been radiated away from the system due to the free space gap. If the system is analysed using the SVD technique, then an estimate of the leaked power, P_{leaked}, is therefore given by

$$P_{\text{leaked}} = (1) \sum_{1}^{\beta} i - \sum_{i}^{\beta} \sigma_i^2. \tag{8.70}$$

For each value for the length of the free space gap, this calculation can be carried out and the power lost due to the gap can be readily found. If a multi-moded system were in question, then since each mode (regardless of azimuthal order) is independent and therefore incoherent, the total power lost can be found by adding the individual power losses directly. Since the accuracy of the multi-moded calculation depends merely on the SVD approach working when applied to multi-moded systems, something that was verified in chapter 3, if the power loss calculation is verified in the single-moded case then it can be assumed to work in the multi-moded case. This fact is exploited in the example above that was used in the verification

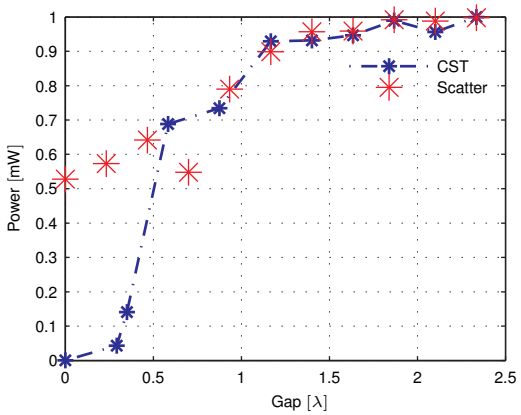

Figure 8.10: Comparison of leaked power as a function of free space gap length (expressed in terms of wavelength ($\lambda = 4.28275$mm)) found using CST and SCATTER

process, as a single-moded computation greatly reduces the computational load encountered when analysing a structure in CST. The dimensions were also selected with computational loads in mind, as the system was kept relatively electrically small.

In the case of the single-moded cavity/waveguide system outlined above, the length of the free space gap was varied between 0 and 10 mm in 1 mm gaps using a parameter sweep within CST. The power recorded by the field monitor surrounding the structure (i.e. the leaked power) was recorded for each gap. Physically, CST measures the root mean square power, and so in order to convert this to peak power to make comparisons with SCATTER, it is necessary to multiply the power figures from CST by a factor of 2. Normally a factor of $\sqrt{2}$ would be used, however due to the manner in which CST carries out the calculations, a factor of 2 is used. A comparison of the results obtained from CST and SCATTER for the leaked power as a function of free space gap length at 70 GHz is shown in figure 8.10. As expected, for a given frequency, the larger the gap, the more power is lost due to each mode having a larger distance over which to diffract. This means that the spatial extent of the mode will exceed the width of the pixel by a greater amount as distance increases. The spatial extent of a mode can be thought of as the area that a mode occupies where its power is non-negligible. In this case, where the TE_{11} mode is in question, then due to conservation of power, as the spatial extent of a mode increases the power on-axis (near the cavity opening) will decrease and so less power will couple into the cavity and will instead be lost from the system.

In the SRON testbed implementation of a SAFARI pixel, the wavelength varies from 35 to 70 μm with a free space gap of length 50 μm. This corresponds to the gap varying in terms of wavelength from $\approx 0.71 - 1.42\ \lambda$. Considering figure 8.10, there is excellent agreement between CST and SCATTER when the length of the gap is of this order, and so the SCATTER approach can be applied with confidence to the SRON measurement system.

The disagreement between CST and SCATTER for lower wavelength values of the free

space gap is expected. For small distances the system is in the reactive nearfield, whose effects are not accounted for in SCATTER, and so discrepancies shall arise. Once this region is exited, excellent agreement results. There is a large disagreement for a gap of 0 mm due to the fact that in CST this represents electrical contact between the waveguide, backshort and cavity. No power will be lost due to this, as a continuous perfectly conducting surface is presented to the system. In SCATTER, a gap of 0 mm corresponds to transitioning from waveguide modes to free space modes at the waveguide aperture, propagating zero distance, and transitioning back into waveguide modes at the cavity aperture. Due to the impedance mismatch between the two mode sets, there is an inherent loss of power once transitioning occurs between them, and this is clearly seen as an unphysical loss of power from the system. Due to the fact that the cavity is large, there is little impedance mismatch between it and free space, however the waveguide section is smaller, giving rise to the TE_{11} mode having an impedance of $\approx$655.6 Ω/sq, versus 377 Ω/sq for the free space modes. The waveguide mode thus has an impedance of $\approx$1.74 times that of the free space modes. This mismatch which is experienced on transition from the waveguide section to free space is sufficient to cause a noticeable loss of power.

The agreement between CST and SCATTER in the region of interest indicates that the SAFARI pixel represented by the experimental setup in SRON can be analysed with confidence using the developed SCATTER code. Although the leaked power calculation was verified explicitly for a single-moded system, this also validates its performance for multi-moded systems, as the only potential sources of error are in the multi-moded analogue of the transition between mode sets and the SVD analysis, however both of these steps were verified for multi-moded systems. This approach allows the investigation of the performance of a pixel both with and without an absorber in the cavity, to evaluate the effect of the free space gap.

8.4 Analysis of a Single SAFARI Pixel

In this section, a single SAFARI S-band pixel (as experimentally setup in SRON) will be considered, with and without an absorber present, in order to examine the effect of the free space gap. A schematic of the pixel is shown in figure 8.11, with a detailed diagram of the measurement system shown later in figure 8.12. One such pixel consists of a few-moded smooth walled conical feed horn that is of length 4.5 mm with exit and entrance aperture diameters of 46 μm and 450 μm respectively. The horn is backed by a cavity with a hemispherical backshort. A Transition Edge Sensor (TES) detector, and associated absorber, sits immediately in front of the cavity entrance, with the cavity being fed by the exit aperture of the horn, however in the models presented in this thesis, the absorber (modelled by an ohmic absorbing sheet as outlined in section 4.5.1) will be placed in the same plane as the cavity entrance. SRON have made use of blackbody illuminators to construct an optical test facility with ultra-low background noise in order to determine the optical response of this pixel, effectively allowing the detector

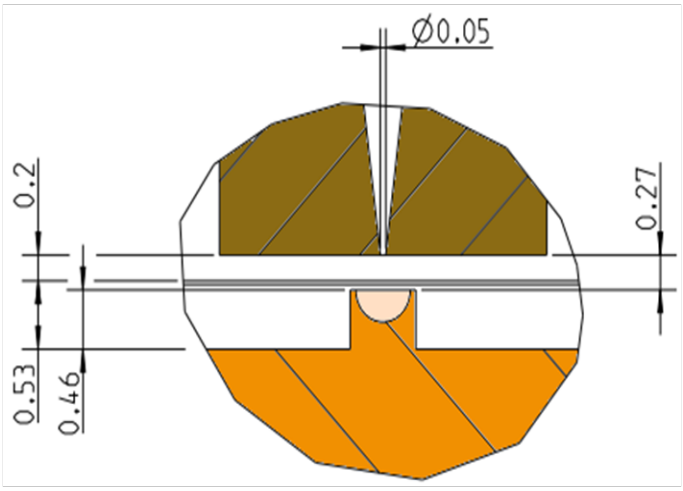

Figure 8.11: Schematic of a single SAFARI pixel

efficiency to be measured. As shown in figure 8.12, several filters are placed in front of the entrance aperture of the pixel. Their frequency response is known, and their transmission as a function of wavelength is given by $\tau(\lambda)$, [64].

It is clear from figure 8.11 that the horn antenna does not make contact with the backshort of the pixel. This is by design, in order to minimise the impact of vibrations in the system during mission launch. The resulting gap was designed to be of the order of the central wavelength of the system, 50 μm, but in reality it is known only to within 10 μm.

As shown in figure 8.12, the system is illuminated by a conical blackbody illuminator that is located a known distance from the front of the pixel. Using this, one can spatially define the source as either the aperture of the 3 K shield of the blackbody or by the size of smaller pinhole apertures that are placed between the illuminator and the horn. The implications of this for the modelling process will be considered in section 8.4.1.

In order to draw valid comparisons with the pixel efficiencies measured at SRON, which give an indication of the power lost in the system, if is necessary to couple the pixels whose scattering matrices will be found, to the optical measurement system in use at SRON. This process will now be examined.

8.4.1 Far-Field Pixel Efficiency Calculations

The coupling of the pixel to the optical measurement system used at SRON will be accomplished by means of calculating an aperture efficiency for the pixel and combining this with a model of a blackbody for a given experimental run (pixel setup). This enables the absorbed power to be found and a direct comparison to be drawn with the measured data. This approach was developed by Dr. Stephen Doherty, (as outlined in [64], which also illustrates how this approach compliments the measurement technique employed at SRON), and was applied to the

free space gap model developed by the author of this thesis in order to include the effects of the free space gap and absorber in a system that yielded results that could be readily compared to the measured data.

Figure 8.12 shows the SRON measurement setup, as outlined in the previous section. The pixel efficiency is naturally highly dependent on this arrangement, and is defined as the aperture efficiency of the horn antenna with the cavity backing, for given values of the diameter of the aperture pinhole, D_{BB}, and the axial separation distance between the horn aperture and the pinhole aperture, L_{BB}. The aperture efficiency of the horn is defined as the ratio of the total absorbed power to the total incident power that arrives from the pinhole aperture. The source can be modelled as a blackbody source that is located behind a circular aperture of diameter D_{BB}, at a distance L_{BB} from the entrance aperture of cavity backed horn antenna. Throughout this chapter, various cooldowns will be referred to. These refer to measurements

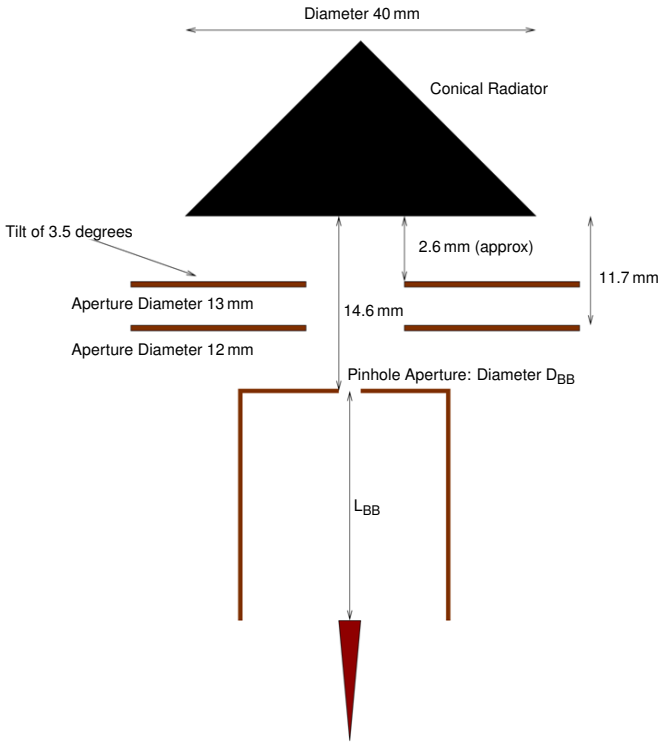

Figure 8.12: Measurement setup at SRON for measuring the efficiency of a single SAFARI pixel

and the associated simulations that were carried out for different values of D_{BB} and L_{BB}. Four such cooldowns are of interest in this thesis, and the values for D_{BB} and L_{BB} that are associated with these cooldowns are given in table 8.2.

If a singular value decomposition is carried out on the S_{11} matrix of the horn/cavity system, then the eigenfields for all of the modes supported by the horn are found. In the Fraunhofer (farfield) region, these fields can be normalised in such a manner as to represent the fraction of power coupled to each of the fields at a given angle of incidence, relative to the total power

that would couple to the antenna in the case of isotropic illumination. The fraction of the total power that arrives at the entrance aperture of the pixel for each eigenfield can then be found by integrating over the angle subtended by the farfield pattern of the horn on the pinhole aperture. This fraction of power, Γ_i, is defined for the farfield pattern of the i^{th} eigenfield, F_i, according to

$$\Gamma_i = \frac{\int_0^{2\pi} \int_{-\chi}^{\chi} F_i(\phi, \theta)\, \mathrm{d}\theta \mathrm{d}\phi}{\int_0^{2\pi} \int_{-\pi/2}^{\pi/2} F_i(\phi, \theta)\, \mathrm{d}\theta \mathrm{d}\phi}, \tag{8.71}$$

where χ is the farfield angle that is associated with a particular pinhole aperture diameter (D_{BB}) and axial distance (L_{BB}) and is defined as

$$\chi = \arctan\left(\frac{D_{BB}}{2L_{BB}}\right). \tag{8.72}$$

With this arrangement in place, the efficiency of the pixel, η_{pixel}, is defined as the fraction of the total power incident on the aperture that is actually absorbed, which is what is measured by the detector that is located within the cavity.

Recall that σ_i, the singular value of the i^{th} eigenfield, gives an indication of the power present at the output of the system relative to the power present in that mode at the input, and that each mode is excited with unity power. In a lossy system, it was established that for each mode, $1 - \sigma_i^2$ gives the power absorbed from that mode by the system. Summing this expression over all eigenmodes therefore yields the total absorbed power, and so η_{pixel} can be expressed as

$$\eta_{\text{pixel}} = \frac{\sum_i (1 - \sigma_i^2)\, \Gamma_i}{\sum_i \Gamma_i}, \tag{8.73}$$

as this is merely the ratio of the total power absorbed by the system to the total power incident on the aperture, as required.

The aperture efficiency calculated above is not the figure measured by the SRON setup, rather the detector efficiency is measured, which is the ratio of the detected power to the power coupled to the pixel from the pinhole aperture. It is not possible in SCATTER to analyse

Cooldown	$\mathbf{D_{BB}}$ [mm]	$\mathbf{L_{BB}}$ [mm]
#46	13.00	26.10
#47	0.90	10.77
#48	0.50	10.20
#49	1.40	10.20

Table 8.2: Dimensions for each cooldown

this portion of the system in isolation, as this would require 'breaking' open the system which would alter the performance of the system. Therefore, in order to make the calculated aperture efficiencies comparable with the measured efficiencies, it is necessary to couple to the SRON measurement system, allowing a comparison to be made with the direct levels of power detected in the measurement setup.

The measurement system is assumed to consist solely of the pixel and the conical blackbody radiator. The surface area of the blackbody is taken to correspond to that of the pinhole aperture and is positioned at the appropriate axial distance in front of the pixel. A Planck curve is used to represent the power density of this blackbody source per unit wavelength, per unit area, per steradian at some temperature T, according to equation 8.74.

$$M_\lambda = \frac{2hc^2}{\lambda^5} \frac{1}{\exp\left(\frac{hc}{\lambda kT}\right) - 1}. \tag{8.74}$$

By considering the area of the blackbody source, A_{ap}, the unit bandwidth, $\Delta\lambda$, and the solid angle of the blackbody that covers the pixel aperture, $\Delta\Omega$, it is possible to calculate the total power that is incident on the pixel aperture from the blackbody source. This total power is given by equation 8.75,

$$P_{BB} = M_\lambda A_{ap} \Delta\lambda \Delta\Omega = \frac{2hc^2}{\lambda^5} \frac{A_{ap}\Delta\lambda\Delta\Omega}{\exp\left(\frac{hc}{\lambda kT}\right) - 1}, \tag{8.75}$$

where the solid angle in question is found by integrating the angular intersection with the pixel aperture, given by $\gamma = \arctan\left(\frac{D_{\text{pixel}}}{2L_{BB}}\right)$,

$$\Delta\Omega = \int_0^{2\pi} \int_0^{\gamma} \sin(\theta)\, \mathrm{d}\theta \mathrm{d}\phi. \tag{8.76}$$

If $\tau(\lambda)$ is the transmission of the filter stack for a given frequency stack as a function of wavelength, then the power incident on the pixel is clearly given by

$$P_{\text{pixel}} = \tau(\lambda) P_{BB}. \tag{8.77}$$

By definition, the total power absorbed by the pixel (for a given pixel configuration) is the product of the total power that is incident on the pixel and aperture absorption efficiency of the pixel in that given setup, i.e. $P_{\text{absorbed}} = \eta_{\text{pixel}} P_{\text{pixel}}$.

Using this measure of the detected power resulted in excellent agreement with the measured values for detected power, [64], and this method will be used to express the efficiency of the system when the free space gap is implemented. Since the work presented in this thesis examines the leaked power, the presented efficiency will be a leaked power 'efficiency', thus a

higher efficiency equates to a higher amount of lost power.

The pixel will now be examined. The case where no absorber is present will be considered first, in order to examine the impact of the free space gap on its own.

8.4.2 Empty Cavity

Physically, the free space gap is a section of waveguide of length 50 μm of extremely large radius, thus it can be considered to be a section of free space as the spatial extent of the waveguide modes within this section of waveguide are small compared to the radius of the guide section. This section of waveguide will now be replaced by a free space section in SCATTER, implemented by invoking the previously outlined steps. The free space gap is fixed to be of length 50 μm, with the wavelength varied from 35-70 μm. This scales the gap from $\approx$ 1.42-0.71λ, and so as the wavelength increases, the length of the gap effectively decreases. Due to the wavelengths involved, it is necessary in the calculations to consider azimuthal orders of $n = 0, 1, 2$.

Figure 8.13 shows the power lost from the system as a percentage of the power incident on the horn aperture from the various SRON measurement arrangements, as previously outlined. The lost power is entirely due to the free space gap, as without the absorber, this is the only

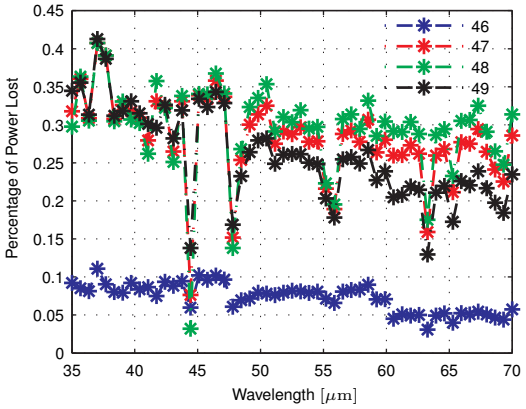

Figure 8.13: Percentage leaked power per cooldown for an empty pixel (simulated)

mechanism through which power can be lost. The spectrum shows many resonant features, likely due to the free space gap forming a resonant cavity, with the effects of this being amplified by the lack of absorber, as all power must leave the integrating cavity and return to the remainder of the pixel, opening up the possibility that it will be lost through the free space gap.

It is clear that the overall trend is for a decrease in the amount of power that is leaked, with increasing wavelength. This is as expected, due to the fact that with increasing wavelength, the effective length of the gap in terms of wavelength, decreases. When this occurs, the modes

will not diffract as much. For the same amount of diffraction to occur for all wavelengths, the gap would have to be increased as wavelength increases. The result of the reduced diffraction is that the power spreads off axis to a lesser degree, and so more power is retained within the pixel radius and so is not lost, giving less leaked power.

In the simulations where the gap was modelled as a section of waveguide, [64], cooldown 46 yielded the smallest amount of leaked power, and cooldown 48 gave rise to the most leaked power. This trend is reproduced when the free space gap is included, as shown in figure 8.13. The average % power lost across the band for each cooldown with no absorber, is shown in column two of table 8.3. The indicated averages further illustrate that cooldown 48 gives rise to the most leaked power.

8.4.3 Combining the Gap and Absorber

The entire pixel will now be considered, including the ohmic sheet absorber. The impedance of the absorber was set to be 368 Ω/sq.

When an absorber is added to the pixel, a new mechanism for losing power is present. The total power loss will not be a sum of the two mechanisms in isolation, but rather a non-linear combination of the mechanisms, as by adding the absorber the system has been fundamentally changed. As power can now be lost within the cavity, less power will re-emerge and enter the free space gap where it will be subject to diffractive losses. As a result of this, the power lost through the free space gap will decrease. Due to the method by which the lost power is calculated, only the total lost power can be found, thus the figures quoted relate to the total power lost due to both the free space gap **and** the absorber. Therefore, although the power lost via the gap will potentially decrease in the presence of the absorber, it will not be possible to verify this with SCATTER in its current form. In order to do this, it would be necessary to implement an additional port to monitor either the absorbed or leaked power, without 'breaking' the system open. The absorber should also have a dampening effect on the resonances noted in figure 8.13.

The total percentage power lost across the band, when the absorber is included, is shown in figure 8.14. Whilst there are still several resonant dips and peaks in the spectrum, they are damped relative to the case without the absorber, as expected. The average % power lost across the band for each cooldown with an absorber is shown in column three of table 8.3. It is clear from this that cooldown 48 is still the one giving rise to the highest leaked power, and that adding the absorber gives higher leaked powers for each cooldown relative to the same figure calculated for an empty pixel, as expected.

Figure 8.15 shows a comparison for each cooldown for a pixel with and without the absorber. It is clear that in most cases the lost power is greater when an absorber is included, due to the additional mechanism for losing power, although the distribution of this total power between

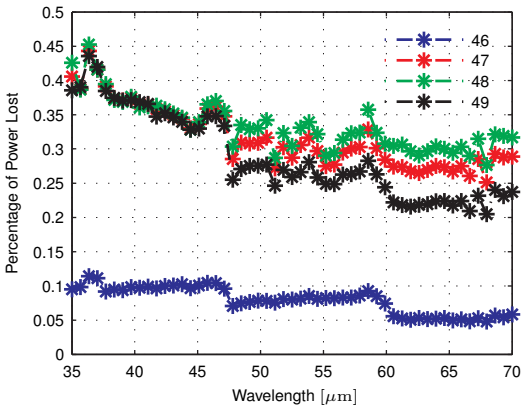

Figure 8.14: Percentage lost power per cooldown for a pixel with absorber

leaked and absorbed power is not known. This gives rise to the higher average occurring when the absorber is considered. The dampening effect of the absorber is also immediately apparent from figure 8.15. Of note, is that as the wavelength increases and the gap gets smaller in terms of wavelength, the lost power in both cases agree to a higher degree. This is likely due to the fact that as the gap is now smaller in terms of wavelength, the propagating modes cannot diffract by the same amount, and so more power is retained within the pixel.

Figure 8.16, [64], compares the actual pixel efficiency for each cooldown of a pixel with just an absorber (waveguide of radius equal to the pixel width replaces the free space gap), a pixel with the free space gap but an empty cavity, and a pixel with both a free space gap and the absorber in the cavity. In the three cooldowns of interest, #47, #48 and #49, the efficiency of the sealed pixel with the absorber is the highest. The lowest efficiency in each case occurs for the pixel with the free space gap where the cavity is empty, with the efficiency of the pixel that includes both the gap and the absorber being marginally higher than this. This is not as might be expected, as it might be assumed that the system which contains both loss mechanisms would give rise to the lowest pixel efficiency despite the fact that the absorber may prevent as much from re-entering the gap region (as it would be absorbed on the first pass) and therefore would reduce the power leaked from the gap. It is also possible that resonant effects with the absorber would result in higher amounts of power being retained in the pixel, and so removing

Cooldown Number	**No Absorber**	**With Absorber**
46	0.0730	0.0794
47	0.2831	0.3172
48	0.2956	0.3340
49	0.2606	0.2884

Table 8.3: Average % power lost across the band for each cooldown, with and without an absorber, both with a free space gap. These figures were calculated as outlined in section 8.4.1

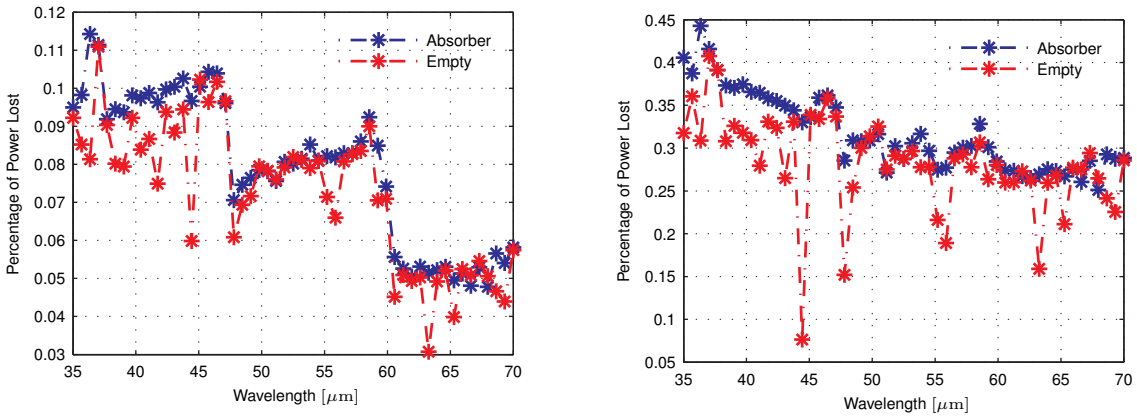

(a) Cooldown 46

(b) Cooldown 47

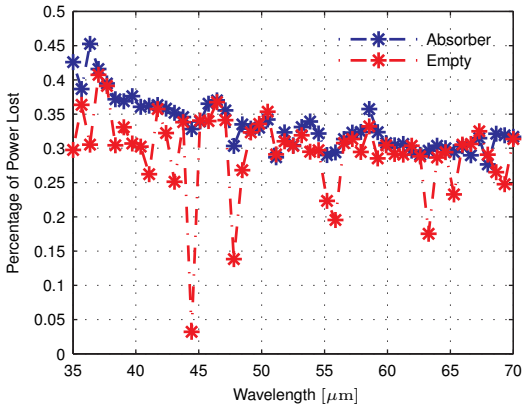

(c) Cooldown 48

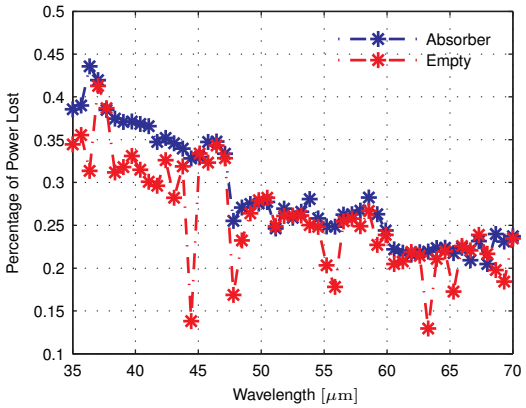

(d) Cooldown 49

Figure 8.15: Comparison of percentage power lost for each cooldown, with and without an absorber

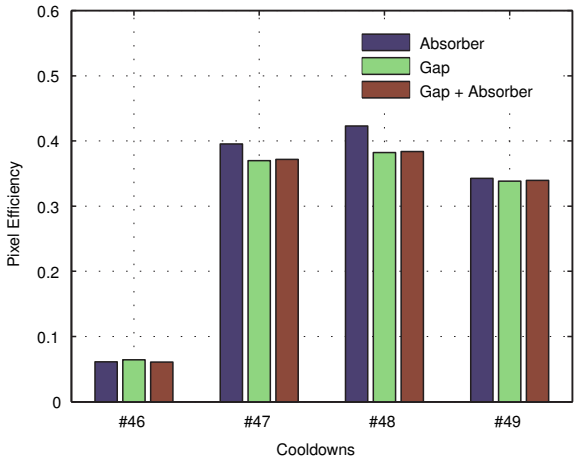

Figure 8.16: Comparison of the pixel efficiency (integrated across the band) for each cooldown for three system configurations

the absorber would result in a higher amount of lost power and a lower pixel efficiency. As stated previously however, it is currently impossible in SCATTER to differentiate between the power radiated away via the free space gap and the power that was absorbed. It is only possible to detect the total power missing. Thus, by equation 8.73, the system with the most overall loss will result in the highest 'efficiency', even though the free space gap will contribute to this figure. Being able to decompose the lost power into leaked and absorbed components would aid in examining the process that yields this result, and allow one to subtract the power lost via radiation from the free space gap and in doing so yield the actual pixel efficiency (the amount of incident power absorbed), whilst including the effects of the free space gap. To do this would involve the implementation of a third port within the SCATTER framework to monitor the power absorbed by the ohmic sheet. This could be implemented as a passive port, so excitation of this port need not be considered. The total lost power would be calculated as normal in this case, and the power that arrives at port 3 could be subtracted from this to yield the power lost via the free space gap.

The argument above raises a possible anomaly in the results presented in figure 8.16. Once the absorber is removed, or both loss mechanisms are simultaneously used, it is not correct to view this as the pixel efficiency in the way in which this quantity is normally considered, but rather a measure of the total power lost from the pixel as a percentage of the total power incident on the pixel. In the original model it is correct to interpret this as the pixel efficiency as it is a measure of the amount of power removed from the system by the absorbing sheet, i.e. the bolometer. In order to investigate the reason for the apparent anomaly, the modes supported by the pixel were examined directly, without coupling to the blackbody source. Two situations were considered, one in which the complete pixel was simulated, including the free space gap, and the other situation considered the entire pixel, but with the gap simulated as a section of waveguide of radius 0.251 mm, the same as the entrance aperture of the cavity

(which is the case referred to in figure 8.16 as absorber). Since in the latter case the pixel is closed by continuous sections of waveguide, this will be referred to as the closed case, and the open case will refer to the model that includes the free space gap.

In each case, the SVD of the S_{11} matrix was calculated. As before, modes with singular values of greater than 0.99 and less than 10^{-6} were rejected. Rejecting the modes with singular values of greater than 0.99 serves two purposes. Modes with singular values in the range 0.99-1 represent modes which although supported at the excitation aperture of the horn, do not enter the cavity. They are effectively rejected by the system. Modes with singular values of slightly above unity (to three or four decimal places) are evanescent, so must be removed. These can be suppressed by increasing the length of the first section of the horn and not considering the additional phase shift, however rejecting modes with singular values of greater than 0.99 provides an additional method to ensure that these modes are not considered. Modes with extremely low levels of power are rejected as they do not contribute to the system in a meaningful manner. The system was analysed for azimuthal orders of $n = 0, 1, 2$, and table 8.4 shows the total number of modes supported by each of the two systems at each wavelength that was considered (35-70 μm).

Wavelength [μm]	**Open Pixel**	**Closed Pixel**
35	11	5
40	12	5
45	9	3
50	9	3
55	9	2
60	9	2
65	8	2
70	7	1

Table 8.4: Total number of modes supported at each frequency for both the open and closed pixels

As expected, in both cases the number of modes supported by the system decreases as the wavelength increases. The only exception to this is in the case of the open pixel, where 11 modes are supported at 35 μm but 12 modes are supported at 40 μm. Examining the equivalent wavelengths for the closed pixel, it can be seen that 5 modes are supported in each case, implying that there is no great difference in the number of modes supported between these two wavelengths. Since the open pixel represents a physically more complicated system involving an additional mode set, and that the free space gap could potentially act as an additional resonance cavity, it is entirely possible that at 40 μm the open pixel could have an additional mode excited relative to a shorter wavelength, and so this observation is of no concern. At each wavelength, significantly more modes are excited in the case of the open pixel. This is due to

the presence of the free space section. The transition between the mode sets results in power being scattered into higher order waveguide modes that are not normally otherwise excited in the pixel. Physically, this represents the behaviour of the effective resonance cavity that the free space gap acts as. Its impact on the system is to excite these additional modes.

Many of the additional modes excited in the case of the open cavity have singular values which although above the cut-off value of 10^{-6}, are still of a very low value. It is possible that they are excited by only a small amount of power scattering into them, however it is not possible to determine if this is the case or if they were absorbed, without breaking the system open. It is not valid either, due to the fact that the pixels are different systems, to simply compare the open pixel to the closed pixel and deduce how much of each mode was excited by the presence of the free space gap and to calculate from this how much power was absorbed from the mode. Thus, it is assumed that each mode supported by the system was excited with unity power. Due to the fact that $1 - \sigma_i^2$ power is assumed to be lost for each mode, this could potentially lead to an overestimation in terms of the power lost, however since the same cut-off criteria is applied to both systems, the overestimation should be consistent and so relative comparisons are still valid.

Figure 8.17 shows the total power lost in each pixel model as a function of wavelength, as a percentage of the power that excited the total number of modes present in each case. It should be noted that this analysis is based on direct use of the S_{11} matrices and so the figures are not expected to be the same as in the case where an efficiency is calculated by considering the coupling of the pixel to a blackbody source in the farfield, as was previously described. Additionally, this will affect any trends that were observed when the coupling to the blackbody source was considered, as this represents an entirely different metric. The overall

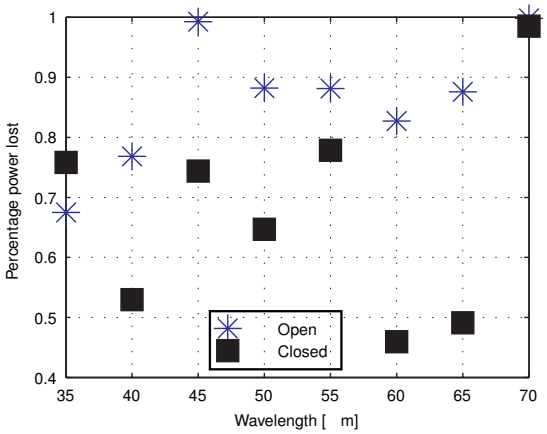

Figure 8.17: Power lost due to the absorber only (closed) and the absorber and the gap (open) as a percentage of the total power contained in the modes excited

trend shown in figure 8.17 is for the lost power in each pixel to increase as wavelength increases. This corresponds to the amount of power being lost increasing as the gap becomes smaller in

terms of wavelength. The system depicted in section 8.3.2 predicts the opposite trend, however this system did not include an absorber. The inclusion of an absorber completely alters the behaviour of the system and so it is possible that the effectively shorter gap means that more power enters the cavity and is removed by the absorber. Thus, even though less power is lost via the free space gap as the gap decreases in effective length, more power is removed by the absorber as there is more power incident upon it.

Considering figure 8.17, it can be seen that in general more power is lost when an open pixel is considered. This is expected, as there are two loss mechanisms when the pixel is open. This is opposite to the results presented in figure 8.16, which predict (using coupling to the farfield blackbody) that the case where only the gap is included gives rise to higher losses when compared with the case where both the gap and the absorber are included. Where the coupling to the farfield source is used in the efficiency calculation, the S_{11} matrix for the pixel is calculated. The farfield of the pixel is then calculated by carrying out an SVD on this matrix and using the modes whose singular values are in the range 10^{-6}-0.99. The coupling between these modes in the farfield and the source is then calculated and an efficiency calculated on the basis of this. When only the S_{11} matrix is analysed to determine the lost power, the opposite trend is observed when comparing the results for both pixels. Since the S_{11} approach returns the expected trend, this implies that the model is functioning correctly, and so the cause of the unexpected result when the pixel efficiency is considered, lies in the coupling to the blackbody source.

To analyse this situation, it is necessary to again consider equation 8.73. Since Γ_i represents the amount of power that is excited at the entrance to the pixel in the i^{th} eigenmode of the system, with the source of the power being the farfield blackbody source, this expression, for a given i, gives a measure of the total power lost from the i^{th} mode of the system relative to the amount of power that mode was excited with at the pixel entrance. Summing for all values of i gives the total power lost as a percentage of the total power excited in the pixel. Since in the open case there are more losses, it is expected that the $(1 - \sigma_i^2)$ term will be greater for the open pixel than for the closed pixel, and this was confirmed in figure 8.17. This should result in the pixel 'efficiency' of the open pixel being higher than the closed case, which is not observed. It must be true then that the coupling of the modes to the farfield source, Γ_i is less in the open case.

The blackbody source is situated behind a pinhole aperture which is located a certain distance from the pixel entrance, co-linear with the central axis of the pixel. Due to this, the coupling of the blackbody source via the aperture to the farfield of the pixel will be maximised when the farfield of the pixel contains a maximum on-axis contribution. For a certain pixel giving rise to a particular value of coupling, this coupling will be reduced if a different pixel is used which has less on-axis power, for example a tooth shape on-axis or more power contained off-axis, as a percentage of the total power in the field. Figures 8.18 and 8.19 show the farfield

patterns of both the open and closed pixels at various wavelengths across the band of interest.

It is clear that in each case the open pixel contains a higher degree of off-axis power as a

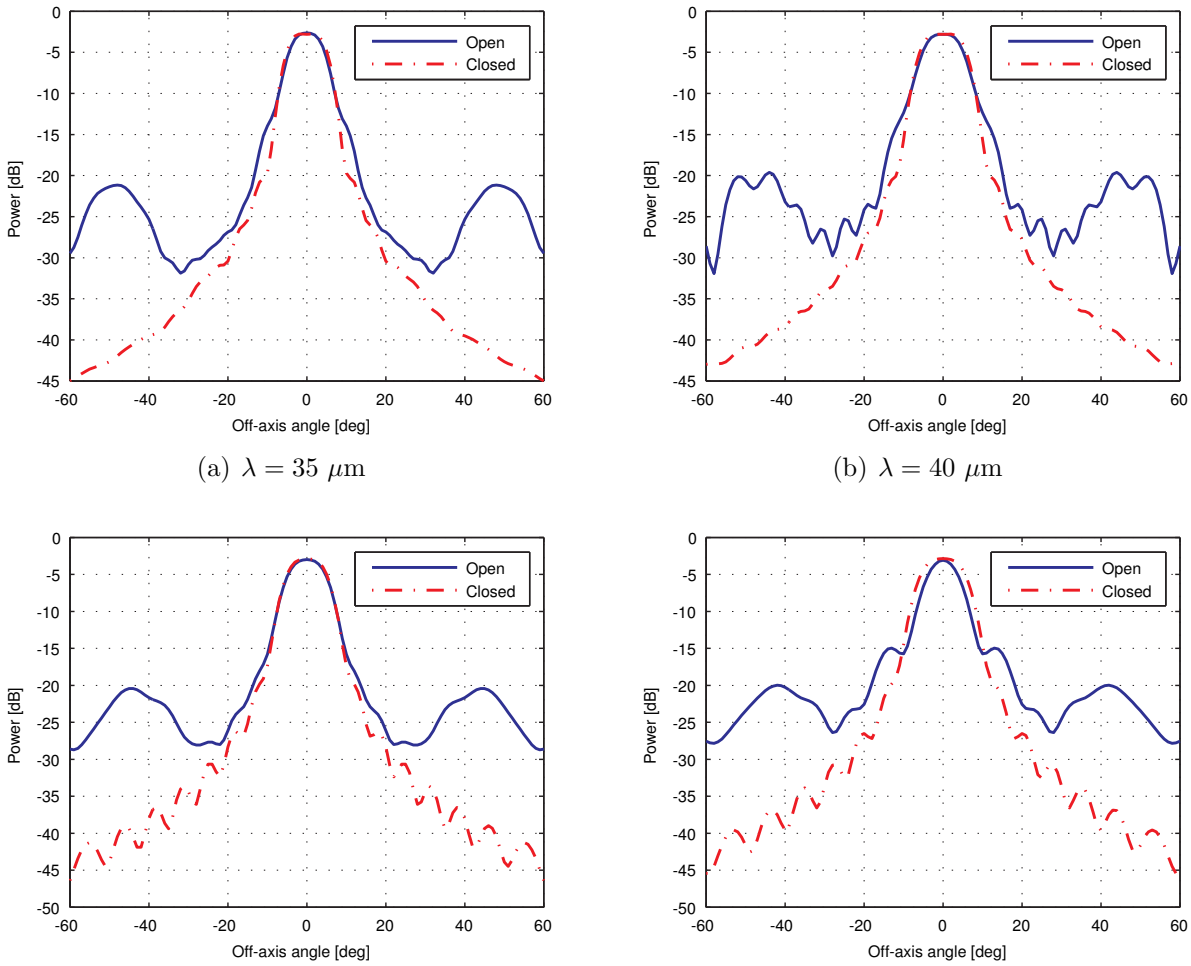

(a) $\lambda = 35$ μm

(b) $\lambda = 40$ μm

(c) $\lambda = 45$ μm

(d) $\lambda = 50$ μm

Figure 8.18: Comparison of the farfields for the open and closed pixels at wavelengths in the range 35-50 μm.

percentage of the total power contained in the field. This occurs due to the excitation of many additional higher order modes relative to those excited in the closed pixel case, as shown in table 8.4. These modes give rise to additional off-axis power in the system. This means that the coupling of the open pixel to the farfield source will be less than that of the equivalent closed pixel. This accounts for the fact that the total power lost from the open pixel is less when calculated by coupling to a farfield source, despite the direct analysis of the scattering matrices indicating that the intuitive result of the lost power being greater in the case of the open pixel is true.

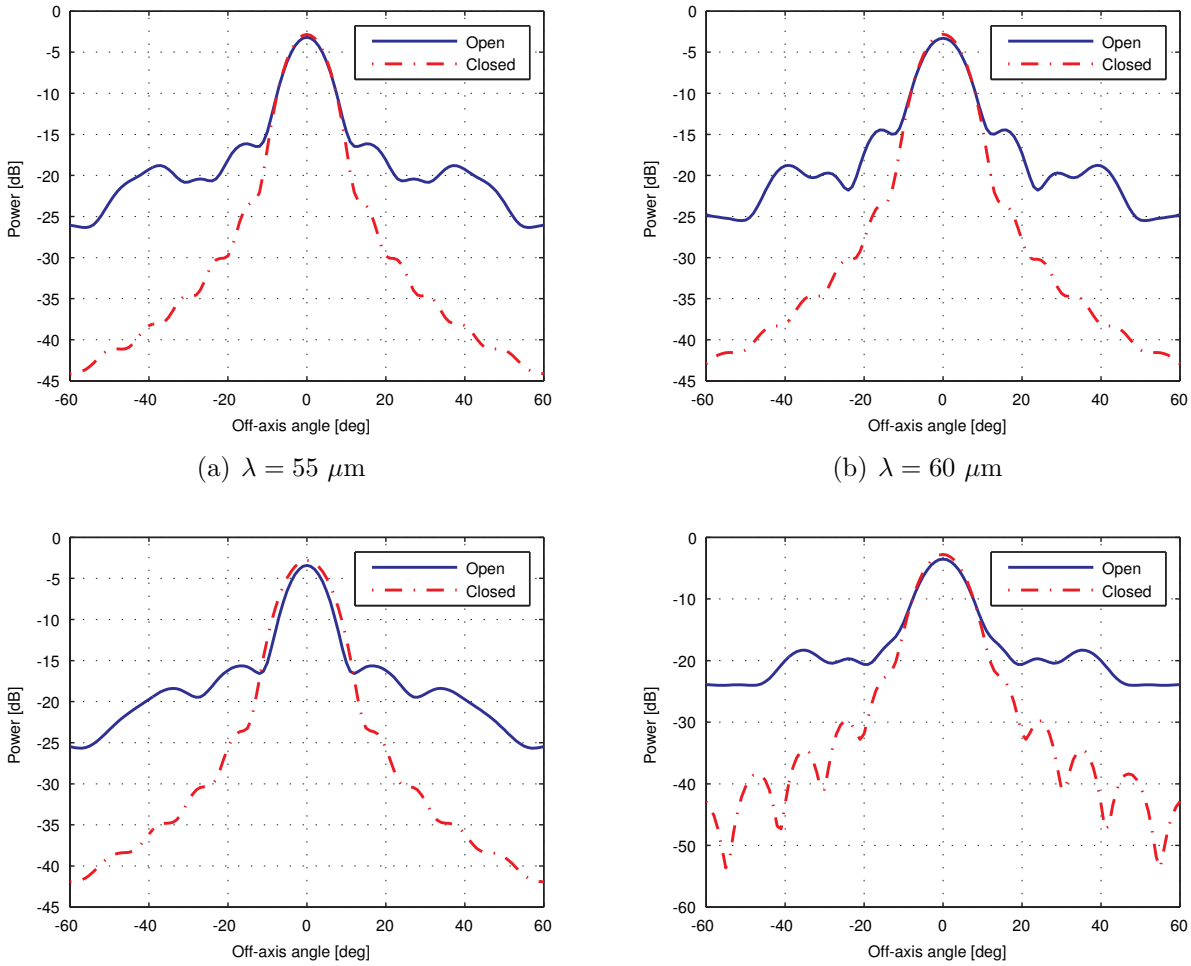

(a) $\lambda = 55$ μm

(b) $\lambda = 60$ μm

(c) $\lambda = 65$ μm

(d) $\lambda = 70$ μm

Figure 8.19: Comparison of the farfields for the open and closed pixels at wavelengths in the range 55-70 μm.

8.5 Conclusions

In this chapter, the analysis of a SAFARI pixel was considered. The steps required to implement such a pixel in SCATTER were detailed. The capability to model a SAFARI pixel (considering the gap between the throat of the horn and the cavity to be a section of waveguide) already existed, however to make the model more realistic it was necessary to model the gap in the system as a free space gap. The steps necessary to achieve this within the framework of the mode matching technique were outlined, including the modelling of the pixel backshort and the ability to radiate power away from the pixel within the free space gap. These elements were integrated in order to model a full SAFARI pixel within SCATTER.

Initially a simple waveguide fed cavity was considered, and the power lost via the free space gap from such a model was calculated. An equivalent system was implemented in CST, resulting in excellent agreement with the SCATTER model, verifying the approach taken. This model was then applied to the SAFARI pixel in order to investigate the effects of adding the free space

gap so as experimental measurements carried out in SRON on a prototype pixel could be better understood. The results of this study indicated that although as expected, the introduction of the gap resulted in additional power being lost from the system, the method used to model the SRON measurement system indicated that the pixel 'efficiency' did not behave as might be expected with the addition of the gap (less power was lost with both the absorber and the gap in place relative to the case with just the absorber). This was reconciled by noting that the excitation of additional modes in the open pixel case resulted in poorer coupling to the system source. It was also noted that in order to quantify the true pixel efficiency, including the impact of the free space gap, it would be necessary to implement an additional waveguide port within SCATTER. The results obtained from the model implemented by the author were found to be physically sensible and self consistent.

Chapter 9

Conclusions

The search for a better understanding as to our understanding of the Universe and its origins is currently a major driving force in astrophysics and cosmology. The key to answering many questions and furthering our knowledge of the Universe may potentially lie in analysing the cosmic microwave background radiation (the radiation that currently permeates the entire Universe that forms the footprint of the Big Bang) and other far infrared radiation. An analysis of this radiation will reveal many pieces of information regarding the formation and evolution of our Universe, galaxies, the stars and the planets. Chapter one of this thesis considers the origin of this radiation and the insights that can be gained from studying it. The methods used to examine the measurements taken were also explored, with a discussion on the study of the polarisation properties of the CMB, as this is a key area of investigation for future missions. A brief description of several past CMB polarisation experiments was also given, in addition to a discussion of several proposed missions in order to provide context for the work presented in this thesis. In order to continue making progress in this area of research, future missions will need to be more sensitive with lower angular resolutions, and be sensitive to polarisation. This presents a significant technological challenge.

One of the main topics addressed in this thesis was the design of smooth-walled horn antennas (a cheaper and more physically acceptable alternative to corrugated horns in terms of weight and thermal penalties) for future CMB and far IR missions. Such horns demand exceptionally high performance in terms of co-polar symmetry and sidelobe/shouldering levels, return loss and cross-polar levels. It was also required that the design process be carried out on the order of hours, on a standard desktop PC using only standard software. In order to achieve this it was necessary to examine the mode matching approach that is used to model horn antennas. Chapter two of this thesis provides an introduction to the background theory of the TE and TM waveguide mode analysis of waveguide structures and the scattering matrix method of modelling such systems, allowing the field patterns of these structures to be predicted. Consideration is also given to the implementation of this modelling technique in the in-house developed software, SCATTER. A brief description is also given of MODAL

(Maynooth Optical Design and Analysis Laboratory) which is also being developed in-house to model quasi-optical systems using a combination of physical optics, Gaussian beam mode analysis and the mode matching technique.

Since it was the aim of the first part of the thesis to design a horn antenna by applying an optimisation algorithm to the mode matching technique, it was necessary to make the code more efficient in order to reduce the computational load and execution time of the optimisation process. Chapter three presented two modifications that were made to SCATTER in order to make the modelling process more efficient. Traditionally, SCATTER considers the number of waveguide modes that is necessary to model the widest part of the horn antenna under consideration at the frequency of interest and uses this number of waveguide modes throughout the horn. This is inefficient, particularly when the process must be repeated thousands of times, as is the case when the optimisation algorithm is applied to it. A modification was presented and verified which allows a number of waveguide modes appropriate to each section of the horn to be used, meaning that no unnecessary modes are considered during the calculation. This was shown to significantly reduce the computational effort required with no loss of accuracy. The other modification that was presented related to the field calculation portion of the code, which is necessary when calculating various performance metrics. The method of singular value decomposition was rigorously applied to the scattering matrices of the horn in order to generate the beam patterns of the horn in a much more efficient manner compared with the traditional method. This modification was verified and shown to significantly reduce the computational time required per calculation.

Chapter four of this thesis was concerned with using the modified, more efficient, version of SCATTER to investigate the limitations of various optimisation techniques. Two global optimisation techniques were identified and presented in this chapter, the genetic algorithm and simulated annealing. These were selected as they are readily available within Matlab, as required, and they are suited to complex functions with large search spaces, as is the case with horn optimisation problems. Two horn types were also identified as candidates for the skeleton of the optimisation process, a few sectioned spline horn (referred to in this thesis as a piecewise conical profile horn) and a Gaussian profile horn antenna. These horns were selected as they have been shown to provide good performance in the areas required for CMB/far IR missions, as outlined above. The adjustments to SCATTER that were necessary in order to interface it with the optimisation algorithms and these horn designs were outlined, along with the results that were obtained when both algorithms were applied to each horn with respect to three performance metrics; co-polar symmetry, return loss, and minimising the cross-polar power. It was found that the best overall results, in terms of the performance requirements outlined above, were obtained when the piecewise conical horn was optimised with respect to minimising the cross-polar power levels using the genetic algorithm. This also met the requirement of only taking of the order of hours to complete the optimisation process on a

standard PC. The simulated annealing was found to take excessively long to produce a result, even when the number of variables was heavily reduced.

In chapter five of this thesis, sample requirements for a horn antenna for a future CMB mission in the 85-115 GHz band were outlined. Using the findings of chapter four, a piecewise conical horn was designed using the genetic algorithm by minimising the cross-polar power. An analysis of the resulting horn using both SCATTER and MODAL revealed that the horn largely met the required performance across the band. This is encouraging as the horn was only designed for optimum cross-polar performance, but it also met the other requirements. The optimisation was also carried out at a single frequency, 100 GHz, but the design performs well across the band. This is impressive for such a simple design that was so heavily customised for 100 GHz operation with respect to one performance metric. The horn that was designed was manufactured by Rutherford Appleton Laboratories in the UK and was rigorously tested using the VNA testbed at NUI Maynooth. It was found to agree very well with the simulated predictions showing that the horn would be suitable for future CMB missions. Consideration was also given to the fact that the manufactured horn contained a transition section to match the cylindrical geometry of the horn to the rectangular geometry of the source.

In order to measure signals as faint as those encountered by CMB/far IR instruments, it is necessary to couple as much power as is possible to the detectors. One method to do this is to use multi-moded systems. In order to fully exploit such systems it is necessary to develop an in-depth understanding of their behaviour and performance. This is currently a challenge. In order to better understand the performance of multi-moded systems, it would be advantageous if the scattering matrices of the systems, with information about the individual waveguide modes, could be recovered from measurements. To do this it would be necessary to selectively excite subsets of the modes supported by the system and analyse the resulting field patterns. In the case of multi-moded horns, a convenient way of doing this would be to illuminate a back to back configuration of the horn under consideration with a plane wave from various angles. This would excite subsets of the modes supported by the horns, in varying proportions. In order to use the measured data obtained from such setups it is necessary to have a model that simulates this. Chapter six of this thesis presents the background theory for implementing off-axis plane wave illumination of a waveguide structure in SCATTER by means of establishing the modes supported by the horn by using an SVD analysis. Each mode is then coupled to a plane wave and excited by an amount that is related to the degree of coupling calculated. This was applied to a back to back arrangement of a corrugated horn in both single and few-moded states. The results were found to be in line with what would be expected.

In chapter seven of this thesis, the development of a 100 GHz Golay cell scalar measurement testbed was presented. The system was tested using horn antennas whose performance is well documented and understood, with the measurements agreeing with the expected results, verifying the correct functioning of the testbed. A design was also presented for a smooth-walled

conical horn that is few-moded at 100 GHz. This horn was manufactured at the University of Manchester and was tested in a back to back arrangement using the Golay cell testbed, illuminated by a conical corrugated horn at various angles. The results are presented in chapter seven, along with the simulated results that were obtained using the model developed in chapter six. The agreement was excellent, considering that this was an initial effort at modelling off-axis systems. This result is encouraging and serves as validation of the code developed in chapter six. It also shows that measuring back to back horns subject to varying angles of illumination allows subsets of the modes supported by the system to be excited in varying proportions which makes this method suitable as a candidate for solving the problem of recovering the scattering matrices of the system, as outlined above.

Chapter eight of this thesis considered the inclusion of a free space gap between two sections of waveguide in SCATTER and applied this to the modelling of a SAFARI like pixel. A prototype of a SAFARI pixel (using cylindrical geometries) is currently under test at SRON. This structure, the feed horn, waveguide gap, and cavity, is electrically large and so full electromagnetic solvers such as CST and COMSOL struggle to simulate it. Previous attempts to model the pixel using the mode matching technique and to use this to recover the measured results for the detector efficiency made no account for the fact that the gap between the horn and pixel is sufficiently wide to allow power to be lost from the pixel through it. This chapter represented this gap as being in free space and presented the theory required to transition from waveguide to free space modes and vice versa and to propagate the free space modes. The theory for introducing a loss in the free space section was also introduced, along with the implementation of the other reflective surfaces in the pixel that must be accounted for. These elements were incorporated into SCATTER and the code was verified using a simple waveguide coupled cavity, including a free space gap which was sufficiently small to be modelled in CST. The model was then applied to the prototype SAFARI pixel. Although the results were as expected in that more power was lost from the system when the gap was added, it was found that the amount of power lost from the pixel relative to the amount of power entering it was reduced relative to expectations. The reason for this was investigated and found to be as a result of the pixel beam being less directed and therefore exhibiting poorer coupling to the source. This reduction in directivity was found to be due to the presence of the higher order modes that are excited in the system when the gap is added.

Bibliography

[1] WMAP web site. `lambda.gsfc.nasa.gov/product/map/current/`. Accessed: 21-08-2014.

[2] J.A. Tauber et al. Planck pre-launch status: The planck mission. *A & A*, 520, September-October 2010.

[3] Planck Collaboration, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, and et al. Planck 2013 results. v. lfi calibration. *ArXiv e-prints*, March 2013.

[4] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, and et al. Planck 2013 results. XVII. Gravitational lensing by large-scale structure. *ArXiv e-prints*, March 2013.

[5] Planck collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, and et al. Planck 2013 results. XV. CMB power spectra and likelihood. *ArXiv e-prints*, March 2013.

[6] G. Gamow. Expanding universe and the origin of elements. *Physical Review*, 70:572–573, 1946.

[7] G. Gamow. Evolution of the universe. *Nature*, 162:680–682, 1948.

[8] F. Hoyle. The synthesis of the elements from hydrogen. *MNRAS*, 106:343, 1946.

[9] G.F. Smoot et al. Preliminary results from the COBE differential microwave radiometers-large angular scale isotropy of the consmic microwave background. *Astrophysical Journal Part 2-Letters*, 371:L1–L5, 1991.

[10] G.F. Smoot et al. Structure in the COBE differential microwave radiometer first-year maps. *Astrophysical Journal Part 2-Letters*, 396:L1–L5, 1992.

[11] P. J. E. Peebles and David T. Wilkinson. Comment on the anisotropy of the primeval fireball. *Phys. Rev.*, 174:2168–2168, Oct 1968.

[12] A. Challinor. CMB anisotropy science: a review. In *Proceedings for IAU Symposium 288, Astrophysics from Antarctica*, 2012.

[13] W. Hu and M. White. A CMB polarization primer. *New Astronomy*, June 1997.

[14] S. Clark. Focus: Ripples in space. *Astronomy Now*, April 1995.

[15] D.J. Fixsen, Cheng E.S., J.M. Gales, J.C. Mather, R.A. Shafer, and E.L. Wright. The cosmic microwave background spectrum from the full COBE FIRAS data set. *Astrophysical Journal*, 473:576–587, 1996.

[16] COBE web site. `lambda.gsfc.nasa.gov/product/cobe/`. Accessed: 21-08-2014.

[17] A.A. Penzias and R.W. Wilson. A measurement of excess antenna temperature at 4080 mc/s. *Astrophysical Journal*, 142:419–421, 1965.

[18] Bennett C.L. and et al. First year wilkinson microwave anisotropy probe (WMAP) observations: Preliminary maps and basic results. *Astrophysicl Journal Supplement Series*, 148, 2003.

[19] G. De Zotti and et al. The Planck surveyor mission: astrophysical prospects. In *Proceedings of the Conference: 3 K Cosmology, Rome, Italy*, 1998.

[20] P.A.R. Ade et al. Planck 2013 results. XVI. Cosmological parameters. 2013.

[21] S. Church and et al. QUEST and DASI: A South-Pole CMB polarisation experiment. *New Astron. Revs*, 47:1083–1089, 2003.

[22] QUaD collaboration. First season QUaD CMB temperature and polarisation power spectra. *Astrophysical Journal*, 674:22–28, 2008.

[23] QUaD collaboration. Second and third season QUaD CMB temperature and polarisation power spectra. *Astrophysical Journal*, 692:1247–1270, 2009.

[24] N.W. Halverson and et al. Degree angular scale interferometer first results: A measurement of the cosmic microwave background angular power spectrum. *Astrophysical Journal*, 568:38–45, 2002.

[25] P. Ade and et al. BICEP2 I: Detection of B-mode Polarisation at Degree Angular Scales. *Phys. Review Letters*, 112, 2014.

[26] BICEP1 Collaboration, D. Barkats, R. Aikin, C. Bischoff, I. Buder, J. P. Kaufman, B. G. Keating, J. M. Kovac, M. Su, P. A. R. Ade, J. O. Battle, E. M. Bierman, J. J. Bock, H. C. Chiang, C. D. Dowell, L. Duband, J. Filippini, E. F. Hivon, W. L. Holzapfel, V. V. Hristov, W. C. Jones, C. L. Kuo, E. M. Leitch, P. V. Mason, T. Matsumura, H. T.

Nguyen, N. Ponthieu, C. Pryke, S. Richter, G. Rocha, C. Sheehy, S. S. Kernasovskiy, Y. D. Takahashi, J. E. Tolan, and K. W. Yoon. Degree-Scale CMB Polarization Measurements from Three Years of BICEP1 Data. *ArXiv e-prints*, October 2013.

[27] BICEP2 collaboration. BICEP2. experiment and three-year data set. *Astrophysical Journal*, 792, 2014.

[28] COrE Consortium. COrE; Cosmic Origins Explorer; a satellite mission for probing cosmic origins, neutrinos masses and the origin of stars and magnetic fields through a high-sensitivity survey of the microwave polarization of the entire sky. Proposal Document.

[29] COrE Consortium. COrE; Cosmic Origins Explorer; a white paper. White Paper.

[30] COrE web site. `www.core-mission.org`. Accessed: 21-08-2014.

[31] The PRISM Collaboration. The polarised radiation imaging and spectroscopy mission-an extended white paper. White Paper.

[32] SPICA Consortium. SPICA at MPE. Mission Overview, Science Potential, and Prospective Contribution and Role of MPE. White Paper for Discussion, December 2011.

[33] ESA. M class internal review report. SPICA. technical and programmatic report. Internal Review, December 2009.

[34] ESA/SRE. SPICA study team collaboration. SPICA Assesment Study Report, 2009.

[35] H. Dole and et al. The cosmic infrared background resolved by spitzer. *Astronomy and Astrophysics*, 451:417–429, 2006.

[36] Ruth Colgan. *Electromagnetic and Quasi-Optical Modelling of Horn Antennas for Far-IR Space Applications*. PhD thesis, NUI Maynooth, 2001.

[37] Emily Gleeson. *Single and Multi-Moded Corrugated Horn Design for Cosmic Microwave Background Experiments*. PhD thesis, NUI Maynooth, 2004.

[38] David White. *The Development of Efficient CAD Software for Terahertz Optical Design and Analysis*. PhD thesis, NUI Maynooth, 2006.

[39] C. O'Sullivan, J.A. Murphy, G. Cahill, M.L. Gradziel, N. Trappe, D. White, V. Yurchenko, S. Withington, and W. Jillema. Developments in quasi-optical design for thz. In *Millimeter and Submillimeter Detectors for Astronomy II.*, 2004.

[40] M.L. Gradziel, C. O'Sullivan, J.A. Murphy, G. Cahill, C. Curran, G.S. abd Pryke, W. Gear, and S. Church. Modelling of the optical performance of millimeter-wave instruments in MODAL. In *SPIE proceedings, terahertz and gigahertz electronics and Photonics VI.*, 2007.

[41] G.L. James. Analysis and design of TE_{11} to HE_{11} corrugated cylindrical waveguide mode converters. *IEEE Trans. MTT*, 29:1059–1066, 1981.

[42] J.A. Murphy, R. Colgan, C. O'Sullivan, B. Maffei, and P. Ade. Radiation patterns of multi-moded corrugated horns for Far-IR space applications. *Infrared Physics and Technology*, 42:515–528, 2001.

[43] S. Ramo, J.R. Whinnery, and T. Van Duzer. *Fields and Waves in Communication Electronics*. Wiley, New York, 3rd edition, 1994.

[44] I.S. Gradshteyn and I.M. Ryzhik. *Table of Integrals, Series and Products*. Academic Press, 1980.

[45] P.F. Goldsmith. *Quasioptical Systems-Gaussian Beam Quasioptical Propagation and Applications*. IEEE Press, 1998.

[46] P.J.B. Clarricoats, A.D. Olver, A.A. Kishk, and L. Shafai. *Microwave Horns and Feeds*. IEEE Press, 1994.

[47] J.A. Murphy and R. Padman. Radiation patterns of few-moded horns and condensing lightpipes. *Infrared Physics*, 31(3):291–299, 1991.

[48] R. Padman and J.A. Murphy. Radiation patterns of scalar horns. *Infrared Physics*, 31, 1991.

[49] C Balanis. *Antenna Theory, Analysis and Design*. Wiley, 1997.

[50] A.C. Ludwig. The definition of cross-polarisation. *IEEE Trans.*, AP-21:116–119, January 1973.

[51] S. Withington, Michael P. Hobson, and Rachel H. Berry. Representing the behaviour of partially coherent optical systems by using overcomplete basis sets. *Journal of the Optical Society of America*, 21(2), February 2004.

[52] MathWorks-Matlab. *Global Optimisation Toolbox-User's Guide. R2011b*, 3.2 edition, September 2011.

[53] Christophe Granet, Graeme L. James, Russell Bolton, and Graham Moorey. A smooth-walled spline-profile horn as an alternative to the corrugated horn for wide band millimetre-wave applications. *IEEE Trans. Antennas and Propagation*, 52(3), March 2004.

[54] G. Yassin, P. Kittara, A. Jiralucksanawong, S. Wangsuya, J. Leech, and M. Jones. A high performance horn for large format focal plane arrays. In *18th International Symposium on Space Terahertz Technology*, pages 199–210, 2007.

[55] Carlos del Rio, Ramon Gonzalo, and Mario Sorolla. High purity gaussian beam excitation by optimal horn antenna. In *Proceedings of ISAP '96*, pages 1133–1136, Chiba, China, 1996.

[56] Roger Penrose. A generalised inverse for matrices. *Proceedings for the Cambridge Philosophical Society*, 51:406–413, July 1955.

[57] Stephen Doherty. *Optical and Quasi-Optical Design and Analysis of Astronomical Instrumentation including a Prototype SAFARI Pixel*. PhD thesis, NUI Maynooth, 2012.

[58] Enda McLoughlin. Optimisation of a near field scanning system. Master's thesis, NUI Maynooth, 2012.

[59] M.L. Gradziel, N. Tynan, N. Trappe, P. Hargrave, and et al. Critical breadboards manufacture and test report (TN6). Technical Report PR-6331-08, NUI Maynooth, Cardiff University, University College London and TICRA, June 2013.

[60] Mairead Bevan. Electromagnetic analysis of horn antennas in the terahertz region. Master's thesis, NUI Maynooth, 2013.

[61] Mark Whale. *Optical Characterisation of Astronomical Submillimetre Receivers including ALMA Bands 5 and 9*. PhD thesis, NUI Maynooth, 2010.

[62] Neil Trappe. *Quasi-Optical Analysis of the HIFI Instrument for the Herschel Space Telescope*. PhD thesis, NUI Maynooth, 2002.

[63] N. Trappe, J. A. Murphy, S. Withington, and W. Jellema. Gaussian beam mode analysis of standing waves between two coupled corrugated horns. *IEEE Transactions on Antennas and Propagation*, 53(5):1755–1761, May 2005.

[64] Department of Experimental Physics Terahertz Space Optics Group. Cavity-backed multimode feed horn interpretation of test results and design improvements (TN4). Technical Report RFQ 3-6418/11/NL/CBi, NUI Maynooth, July 2014.